COMPUTER ANIMATION AND VIRTUAL WORLDS
Comp. Anim. Virtual Worlds 2008; **19**: 129–149
Published online 27 November 2007 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/cav.218

A scalable interest management scheme for distributed virtual environments

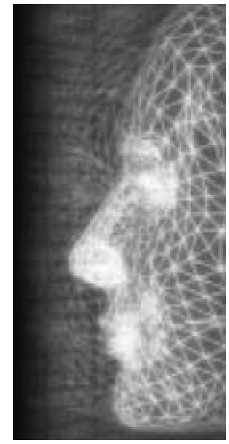

By Seunghyun Han, Mingyu Lim, Dongman Lee and Soon J. Hyun*

With the expansion of the internet and its bandwidth, distributed virtual environment (DVE) applications have become more prevalent. In DVE applications, users frequently crowd in a specific place, and a key aspect to consider is how to provide interactive performance for users. However, existing approaches using multicast require users to receive uninteresting messages. Even though recent works have addressed fine-grained filtering, they still incur other drawbacks in terms of assigning lots of multicast addresses or handling overhead of multicast groups. This makes the system less scalable as the number of users increases. In this paper, we propose a new scalable filtering scheme that reduces not only the number of messages during interaction in a region and among neighboring regions, but also the number of multicast addresses without significant computational overhead. Interest management in a region dynamically creates groups of users with the same interests. While members communicate with each other with high fidelity, a representative sends information to non-members with low frequency. For interaction among neighboring regions, we propose a sub-region concept to select only a subset of users from the neighboring regions based on proximity, the distribution of the users' locations, and the viewing direction of a user.

Received: 25 April 2006; Revised: 2 October 2007; Accepted: 8 October 2007

KEY WORDS: distributed virtual environment; interest management; multicast; scalability; sub-region; user interest group

Introduction

As the internet becomes prevalent and network bandwidth increases, various distributed virtual environment (DVE) applications have emerged. As the number of users in a DVE increases, a large number of interactions are likely to impose a heavy burden on network and computational resources. One of the key aspects to consider for lessening the burden is scalability, that is, a system should keep interactive performance of users from significantly degrading even with an increase in the number of users.[1] Though it has been rapidly improving, network resources still remain very expensive in comparison with computational resources. To overcome this, various relevance-filtering mechanisms[2–4] have been proposed. Filtering means that a user receives messages only from those in whom the user is interested. Users' interests are represented not only by the spatial distance among them but also by social groups such as those who have the same hobbies or those who want to buy the same items in a shopping mall. Filtering mechanisms can be classified as follows: *region-based filtering*,[5–8] limiting the scope that messages are bound into spatial regions; *aura-based filtering*,[9–11] localizing the spatial area of interest (AOI) of the users; *class-based filtering*,[12,13] classifying delivered messages based on object types; *hybrid approaches*,[14–19] combining region-based filtering and aura-based filtering; and *aggregation mechanisms*,[20,21] providing only an abstracted view of target information.

Since the region-based approaches[5–8] use only spatial regions, they require a user to receive all the messages generated by others in the current region in which he participates, even though he is not interested in them. This imposes high communication overhead on users if many users crowd into a region and its neighboring

*Correspondence to: D. Lee, Information and Communications University, 119 Munjiro, Yuseong-gu, Daejeon, Korea.
E-mail: dennis@icu.ac.kr

regions. The aura-based approaches[9–11] have the same problem as the region-based ones when many users gather together in certain areas because a user receives all the messages generated by others whose aura collides with his aura even though he is not interested in them. Class-based filtering[12] is mostly based on object types (classes or organizations) in which users are classified before starting a session. However, it is inappropriate to most DVE applications where it is difficult to fix a role or interest of a user to a pre-defined type or class throughout a session. Although there is another approach which uses dynamic adaptation of users' interests according to user behavior,[13] it is not scalable since a server manages the interest levels and their changes between all the pairs of objects in a virtual world whenever one of them interacts with others. Recently, the hybrid approaches[14–19] have solved the problems mentioned above using fine-grained filtering where a filtering scheme is dynamically adapted to users' interests or QoS requirements. However, these approaches require dynamic region partitioning overhead,[16] multicast group management overhead,[17,18] or generate unnecessary messages when many users with different interest are clumped in a specific place.[14,15,19] Similarly, the aggregation mechanisms either require as many multicast addresses as the number of objects[20] or consider only spatial aggregation.[21] Thus, the existing filtering mechanisms do not scale well in terms of computational overhead and network resource usage for DVEs where users crowd in a specific place and their roles or interests are not fixed to a single one during a session.

To overcome the limitations of the existing approaches, we propose a new interest management scheme that filters unnecessary data transmissions based on user interests and spatial distance without significant use of multicast addresses or high region partitioning overhead. We assume that a virtual world is divided into geographically distinct rectangular regions and that each user has its own spatial interest area (IA). Interaction among users with the same interest occurs only in a user's IA (intra-region interaction), while a user can see others regardless of their interests in the same region and its neighboring regions (inter-region interaction).

For intra-region interaction, we leverage human behavior where users tend to move and crowd in specific places based on their own interests, and interact with those who have similar interests. In the proposed scheme, users with the same interests dynamically form a group when they move closer to each other. Each user in a group multicasts update messages to the rest of the group. On the other hand, when a group is included in or collides with the IA of a user who is not a member of the group, the representative user (RU) of the group sends him updates for the group with low frequency. The member with the lowest ID among the members of a group is elected as the representative.

For further scalability improvement, filtering on inter-region interaction is used to enable users to interact mostly with adjacent users in the neighboring region(s) instead of users from a distance. Our scheme uses a sub-region concept to select only a subset of users from neighboring regions whose members have high probability of interaction with users in the current region. This subset of users forms another multicast group. This enables users in the region to avoid receiving all the update messages from neighboring regions. They receive the update messages regarding only the users in whom they are interested in the neighboring regions. The size of a sub-region scope is dynamically changed according to the distribution of users in the neighboring regions so that inter-region interaction does not deteriorate the interactive performance of a user, regardless of the number of users or their distribution. In addition, only sub-regions in the neighboring regions, which are in the field of the vision of a user, are selected. The simulation results show that the proposed scheme enhances system performance without significantly degrading the interactive performance of users in large-scale DVE systems compared with existing approaches.

This paper is organized as follows: in the following sections, we discuss the various approaches for existing systems and describe a fundamental operational model for a DVE system. Next, we describe our scalable interest management and group management scheme for intra-region interaction, followed by our scalable inter-region interaction management scheme. Following this, we give an overview of our scalable network framework, ATLAS,[22] and its experimental DVE application, and we also describe our experimental results. Finally, we present our conclusions.

Related Work

In this section, we describe the existing approaches proposed for scalable interest management and the aggregation schemes used to reduce the network bandwidth for large DVEs. We analyze these approaches in terms of whether they are suitable for

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

DVE applications where users often crowd in a specific place, and whether they use an excessive number of multicast addresses or not.

Interest Management

Various interest management schemes have been proposed to reduce the network bandwidth by limiting the scope that messages are bound into. These can mainly be divided into three approaches: spatial distance based, class based, and a hybrid of these two approaches.

Spatial distance (proximity)-based filtering mechanisms are mainly divided into region-based filtering and aura-based filtering. NPSNET[6,7] and SPLINE[5] are exemplary systems that use a region-based filtering mechanism. They divide a virtual world into several regions and assign a separate multicast address to each region. In addition, these systems improve scalability by bounding the dissipation of update messages within a region and its neighboring regions if required. However, update messages are transmitted to all users in the target regions regardless of the users' interests. The number of unnecessary messages becomes larger as the number of users with different interests in a region increases. There is another region-based approach which partitions regions according to users' density.[8] However, since partitioning is a static operation, it is not adaptable to dynamic changes of the user distribution in regions.

Unlike with region-based filtering, in aura-based filtering the scope for message dissipation is controlled by each user; that is, each user has his own ''aura'' that represents the spatial area of his interest. When auras collide, a connection between them is established and update messages are exchanged through the connection. The MASSIVE[9] and DIVE[10,11] systems adopt this filtering mechanism. However, a user will still receive all the information from other users when they come together in range of his/her aura, even though they do not share the same interest. In addition, MASSIVE and DIVE overlook interaction among users in neighboring regions.

In class-based filtering, messages are filtered not only by spatial relationship but also by the users' interest. In high-level architecture (HLA),[12] objects in a virtual world are classified based on classes. Users register their interest in any related class before participating in the world, and thus receive messages only from objects of the pre-registered classes. However, this does not work well with DVE systems where the interests of users are not fixed but change dynamically. Ding and Zhu[13] devised another interaction model with various interest degrees. According to a user's behavior, the user changes his/her interest level based on an interest network and an influence rule. However, since this approach is based on the client/server model, and thus a server manages the interests of all the objects in a virtual world, it is not scalable as the number of objects increases.

Some systems attempt to combine the filtering schemes described above for fine-grained data filtering. Liu *et al*. propose a tracking needless grouping scheme combined with region-based and aura-based filtering.[14] Since it defines a cell's vision domain and visible set, a user does not have to frequently find visible cells whenever the user moves one step. Instead, a user in a cell uses the cell's vision domain and visible set until the user goes to another cell. However, since this approach is based on only spatial interest, it still generates unnecessary messages when a lot of users with different interest are clumped in a specific place.

MASSIVE-3[15] extends the ''locale'' concept introduced by SPLINE.[5] It supports not only region-based filtering but also class-based filtering and an aggregation approach using an ''aspect'' concept. However, since it adopts the client/server communication architecture, it incurs significant load on the server as the number of users increases.

SCORE[16] uses dynamic division of cells according to the density of users. If the number of users increases in a cell, the cell is sub-divided into smaller equal-sized cells. This approach, however, incurs management overhead for handling multicast groups. If many users quickly move around a virtual world, they need to repeat a series of join/leave and split/merge of cells, which are notified to all the affected users. In addition, since SCORE equally splits the cell even if many users are clumped in a specific area of a cell, the rest of the cell where few users are located would be split uselessly.

Bamboo[17] provides three tiers of data filtering. In the first tier, a region can be dynamically sub-divided into eight smaller regions using an octree when lots of objects are clumped in the region. The second tier can reduce more data than the first tier with generic information such as the AOI of a user. The third tier allows users to receive only necessary data from others with a protocol which is designed to support a specific situation adapting to the characteristics of a target DVE system. Since the first tier filtering is very similar to SCORE, the three-tiered architecture has the same management overhead for handling multicast groups.

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

VELVET[18] aims to allow real-time adaptation according to the local client's needs using not only spatial distance, but also interests. With the use of revisited multi-level AOI and parallel virtual world, the system supports collaboration of heterogeneous users who have different resource capabilities. Although Bamboo and VELVET use fine-grained filtering based on users' interests, both approaches assign as many multicast addresses as the number of users. This makes the system less scalable since multicast addresses are a limited resource.

Pryce's approach[19] dynamically allocates a hierarchy of multicast groups to volumes of space in which interaction is taking place and uses this hierarchy to support QoS adaptation by using a layered multicast protocol as is prevalent in multimedia distribution. Group subdivision is performed dynamically when members of a group detect that the QoS available to the group has degraded. Each entity receives messages from all enclosing groups in the hierarchy but selectively transmits update messages, that is, they are transmitted most frequently to the innermost groups and less frequently to the outermost groups. However, since this group subdivision is based on only spatial interest, it still generates unnecessary messages when a lot of users with different interests are clumped in a specific place.

Aggregation Mechanisms

Aggregation schemes are proposed to lessen network bandwidth by providing an abstracted view of objects. Aggregation servers or local hosts merge update messages from multiple clients and send an aggregated view to other clients which then update their view accordingly.

Paradise[20] explores a graphical abstraction to provide aggregated views to other users who have relatively low interest in a group of objects. Projection is performed based on the grid (spatial area) and the object type (class or organizations). Objects send high fidelity information to others who are interested in them; otherwise, an aggregated view is transmitted. However, Paradise assigns as many multicast addresses as the number of objects for transmission of high-fidelity data, and has the overhead of managing summary views of projections.

Third party objects[21] as introduced in MASSIVE-II provide an aggregation view of a group of users to other users from a distance; this is called ''secondary sourcing'' performed by third party objects. A third party object gathers information from group members belonging to them, generates an aggregation view of the group and then transmits it to users who are within distance. However, since third party objects generate the aggregation view based only on spatial distance, a user still receives all update messages from other users whom the user is not interested in when users are clumped in a small area.

Table 1 summarizes the features of the existing interest management schemes.

An Operational Model

For designing a scalable interest management architecture, we basically use a hybrid of aura-based and region-based filtering schemes which restrict the user's perception of a virtual world based on spatial area and proximity. That is, a virtual world is divided into rectangular regions (spatial areas) and a user has an IA, which is similar to an aura in the MASSIVE or DIVE systems, within which his interest is bound (proximity). While a user can see others regardless of their interests in the same region as well as those in neighboring regions, the user interacts only with others who are in his IA. The size of an IA is smaller than the maximum range of sight for a user (the current region and its neighboring regions). It increases scalability by localizing the number of messages which a user receives. However, a user is still overwhelmed with numerous messages especially when many people are crowded in a small area where most users are included in each other's IA or if the user distribution in a virtual world is skewed.

For a more scalable design, we leverage real world situations. First, in the real world, people focus more on objects of high interest than those of low interest. Another factor is that people tend to interact more frequently with those that have similar interests than with those that have different interests. Moreover, users with the same interest often gather together in specific locations where the target objects exist.

To exploit these findings into our interest management scheme, we assume that a user wishes to receive high-fidelity data from high interest objects, while low-fidelity data are received from low interest objects. We also assume that users with the same interests are highly likely to be interested in each other compared with users with different interests. A user can specify interests on his IA as a property. Users dynamically form a group if they are bound into each other's IA and share the same interests, and they exchange high-fidelity data. The leader of a group, on behalf of the group members,

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

Approach	Filtering type	Feature	Limitations
NPSNET	Region-based	Hexagonal cells	Not interest-based filtering
SPLINE	Region-based	Locales	Not interest-based filtering
Steed's scheme	Region-based	Region partitioning based on user density	Static region partitioning
MASSIVE	Aura-based	Spatial model of interaction	Not interest-based filtering
DIVE	Aura-based	Spatial model of interaction	Not interest-based filtering
HLA	Class-based	Pre-determined class	Not adaptable to dynamic change of interests
Ding's scheme	Class-based	Dynamic change of interest degree	Client/server model
Liu's scheme	Hybrid	Tracking needless grouping	Not interest-based filtering
MASSIVE-3	Hybrid	Extended locales	Client/server model
SCORE	Hybrid	Dynamic cell division	Overhead of handling multicast groups
Bamboo	Hybrid	Three-tiered filtering	Overhead of handling multicast groups and excessive use of multicast address
VELVET	Hybrid	Multi-level AOI and parallel virtual world	Excessive use of multicast addresses
Pryce's scheme	Hybrid	Dynamic group partitioning based on QoS requirements	Not interest-based filtering
Paradise	Aggregation	Aggregated view	Excessive use of multicast addresses
MASSIVE-2	Aggregation	Third-party object which generates aggregated view	Not interest-based filtering

Table I. Summary of existing approaches

sends an aggregate view of the group members with low fidelity to the rest of users in a region. It results in that users who do not share the same interests receive low-fidelity data even if they are close enough to interact. This helps reduce the number of messages when many people are crowded in a small area. We call this *intra-region interest management*.

Second, while users have high interest in the region that they belong to, they have relatively low interest in neighboring regions because most people perceive objects nearer to them rather than those farther from them. It is unnecessary for a user to continually receive update information of all the users in a neighboring region. For this, we divide a region into sub-regions and users partially receive update information from neighboring regions. Moreover, a problem can occur when users are crowded into a specific place in neighboring regions which results in excessive filtering or no filtering benefit. We solve these problems by dynamically adjusting the scope of a sub-region boundary. We call this *inter-region interest management*.

Our scheme is designed to not only reduce network bandwidth but also to lessen load on servers, since we cannot claim that a system is scalable just by reducing network bandwidth. We adopt the peer-server approach to reduce the workload of servers by distributing the roles of servers to peers as much as possible. Servers such as region managers manage membership information, but actual communication or collision detection of an IA are performed by a region manager which resides in each peer to lessen the workload of servers and remove bottlenecks.

We divide messages generated by users into two types: position-update messages and interaction messages. The former are generated when users move in the virtual world; the latter are generated when users interact with virtual world objects or with other users, which include the motion details for avatars, text chat, audio messages, *etc*. There is no clear definition of high- and low-fidelity data.[17] In this paper, we define the update frequency as the measure of fidelity. When a user wants high-fidelity data on an object that he is interested

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

in, an update message(s) is immediately sent to him whenever there is a change in the object. On the other hand, update messages are transmitted to him with low frequency if he is not interested in an object, that is, he gets low-fidelity data on the object.

In the following sections, we describe the detailed mechanisms for intra-region interest management using user interest-based group creation, destruction, election of RU and data transmission, and inter-region interest management using sub-regions.

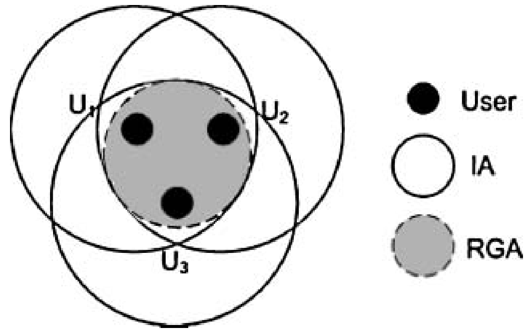

Figure 1. Creation of a user group among users U_1, U_2, and U_3.

Intra-Region Interest Management

In this section, we focus on interest management when a user interacts with others in his IA. We describe how an interest-based group is created, how a group is managed, and how interaction messages are transmitted by this group.

User Interest-Based Group

When a user joins a region, the user receives information about the other users who have already participated in the region. The information contains the addresses, interests, and joined groups of the users. The IA of a user is represented as a sphere and the user specifies his interests as the IA's property. Users can receive selected messages according to the interests specified in their IAs. User interests can be represented as any type of object, such as virtual world objects, as well as other users and specific locations in a region.

As a user navigates a region, objects (users or virtual world objects) come into his IA. When a user goes into the IA of another user with the same interest(s), these two users form a group. We name the boundary of a group the representative group area (RGA) as shown in Figure 1. Group members communicate with each other via a multicast address assigned to an RGA. When all users leave an RGA or change their interests to something different from the current one, the RGA ceases to exist.

If each user has an IA with a different size, the diameter of an RGA becomes the radius of the smallest IA because an RGA is created only when all group members are in the intersection of their IAs as shown in Figure 2.

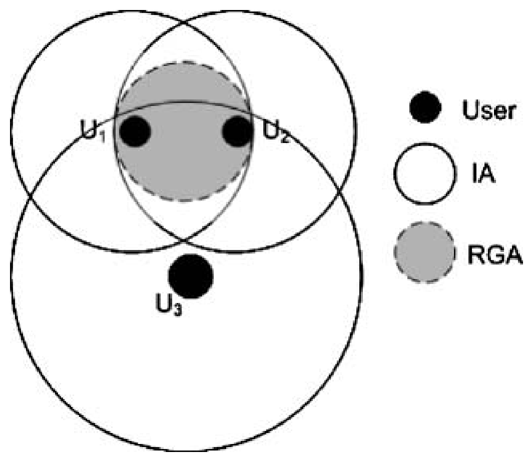

Figure 2. RGA maximum size anomaly: user U_1 and U_2 can form a group while user U_3 cannot with either user U_1 or U_2, since user U_3 is not included in the IA of either user U_1 or U_2.

User with Several Interests

Users can be interested in several types of objects in a virtual world. However, users should not create or join several groups at the same time because this causes not only duplication of user representations, but also duplicated message transmission for the same user from multiple groups. To avoid these problems, users join or create a group based on the priorities of their interests. A user might join a group and communicate with other members. If he gets included in the group with a higher priority interest, a user leaves the current group and joins the group with the higher priority. Figure 3 shows an example that RGA reformation due to a change in priority of interests. A user also can change interest priority. Whenever a user changes the priority,

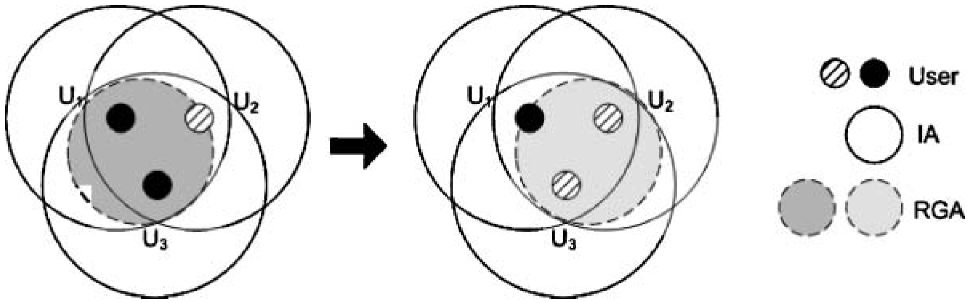

Figure 3. RGA reformation by change in priority of interests: user U_3 changes the priority of his interests and reforms a new group (RGA) with user U_2.

he notifies other users in order to reform a new group with nearby users.

Election of RU

Whenever a group is created, its RU must be elected. The primary role of an RU is to infrequently send update information for group members to users who have low interest in the group. Negotiations between group members are required for election of an RU. An RU is elected based on criteria such as how long it was in the group and the object ID; users multicast their IDs or session time to all the other users in a region whenever their positions are updated. This allows each member to be aware of who has the highest priority. The user with the highest priority can create an RGA and the other users become members. When it leaves the group, the current RU elects a new RU based on the membership list. It notifies in the RU to the next candidate RU, the other group members, and other users in the current region of the change. The new RU then multicasts its identity to all the objects participating in a region to confirm that there has been a change in the RU. RU change notification and confirmation can be piggy-backed to position-update messages for the RU.

Data Transmission

A user uses separate multicast addresses to send and receive messages: one for low-fidelity messages (assigned to a region) and the other for high-fidelity messages (assigned to an RGA). If a user does not belong to any group, the user sends low-fidelity messages via a multicast address assigned to the region which he belongs to. When a user becomes a member of a group, he sends high-fidelity messages through his RGA multicast address. In addition, if a user is an RU, he, on behalf of the group members, also sends low-fidelity messages for his group via the multicast address of his region. As a receiver, a user joins a region which he belongs to as well as neighboring regions for low-fidelity data, and joins his group for high-fidelity data.

Figure 4 illustrates an example of how a user receives update messages from other users depending on its interest. First of all, each user can see others because they are within his/her neighboring regions. For example, user U_1 in region R_1 can see user U_4 in the same region R_1. However, they do not exchange high-fidelity messages even though they have the same interest, because they are distant from each other. There are two

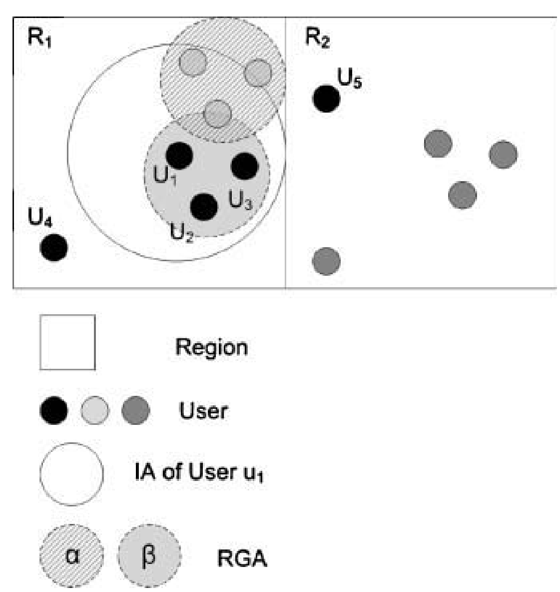

Figure 4. Interactions of user U_1 based on interests.

DOI: 10.1002/cav

different groups of users in the region R_1. Since user U_1 belongs to RGA β, it sends/receives high-fidelity messages to/from the other two group members, U_2 and U_3, via the multicast address assigned to RGA β. On the other hand, the RU U_2 of RGA β sends update messages on behalf of user U_1 and U_3 with low frequency through the multicast address assigned to region R_1, and which are received by users in region R_1 and R_2 receive them.

Inter-Region Interest Management

In this section, we describe interest management between regions; how a sub-region is defined, how update messages are limited by it, and how sub-region boundaries are adapted to various conditions.

Sub-Regions

For interactions between users in adjacent regions, each user must receive periodic updates of the status of other users in neighboring regions such as their positions. While users have high interest in the region that they belong to, they tend to have relatively low interest in neighboring regions. We deal with the differences of interest based on the spatial distance between users. That is, users are interested in the entire region that they belong to, but their interest in neighboring regions is likely to diminish as the distance from them increases.

As a user approaches the boundary of a region, it is likely that he will also become interested in a neighboring region. However, it is unnecessary for a user to continually receive updates for all the users in his neighboring region. We use sub-regions to distinguish the interest of users in a neighboring region. Users in a sub-region that is near a neighboring region tend to have a higher degree of interest in users in the neighboring region than in those in a sub-region that is distant from the neighboring region as shown in Figure 5.

The sub-region concept is similar to that in CyberWalk[22] system in terms of dividing a region. However, there are several significant differences between them. First, in the proposed approach, sub-region boundaries are dynamically adjusted to user distribution for improving interactive performance of users without significant management overhead. Second, CyberWalk system is mainly focused on load balancing among multiple servers when user distribution becomes unbalanced among them, while the proposed approach addresses how to receive only interested messages using multicast communication exploiting the sub-region concept. We tackled the load-distribution issue among multiple servers in another publication[23] and it is out of scope in this paper.

Group Communication Model for Inter-Region Interaction Management

Existing DVE systems assign a multicast address to each region by which users can interact with each other in the same region.[5,6] We distinguish interactions within a region and those with neighboring regions. To enable

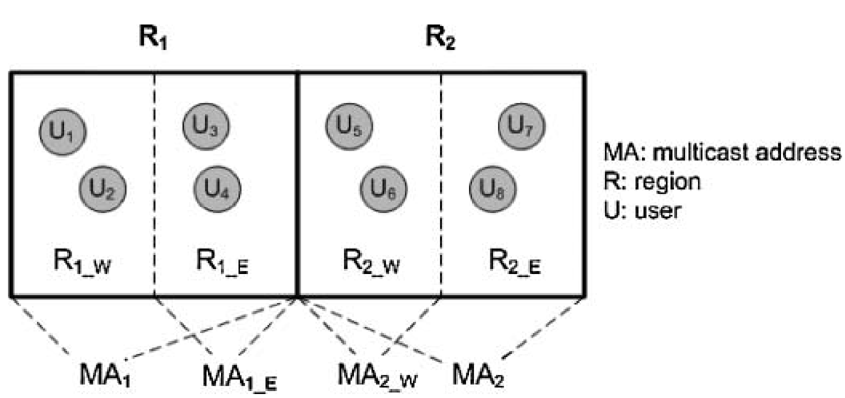

Figure 5. Sub-region concept: region R_1 and R_2 is logically divided into sub-region R_{1_W}, R_{1_E} and sub-region R_{2_W}, R_{2_E}, respectively. While user U_1 in sub-region R_{1_W} only gets update messages for users in sub-region R_{2_W} and R_1, the user U_3 in sub-region R_{1_E} receives all the update messages for users in region R_2, that is, both sub-region R_{2_W} and R_{2_E} and R_1.

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

users to receive only subsets that they are likely to be interested in, we assign a separate multicast address to each sub-region.

Since each user must join at least two multicast addresses, if they send messages using both addresses, then they will receive duplicate messages that they have already received from other address. To avoid this problem, users use different multicast addresses according to whether they send or receive messages based on related sub-regions. A user multicasts update messages to other users who are interested in him/her. On the other hand, a user receives, via a different multicast address, update messages from other users in whom he/she is interested.

We now examine our group communication model for intra- and inter-region interactions in detail. In Figure 5, we can regard sub-regions R_{1_E} and R_{2_W} as the boundary areas of region R_1 and R_2. To support inter-region interaction, we can assign an additional multicast address for each sub-region. First of all, each region, R_1 and R_2, needs a multicast address for intra-region interaction. The boundary areas, sub-regions R_{1_E} and R_{2_W}, need an additional multicast address for inter-region interaction of users who are interested in sub-regions R_{1_E} and R_{2_W}, respectively. If users U_1 and U_2 in sub-region R_{1_W} are interested in sub-region R_{2_W}, they listen to multicast address assigned to sub-region R_{2_W}, MA_{2_W}, for inter-region interaction. Note that the multicast address assigned to a sub-region is only used for inter-region interactions, that is, users U_1 and U_2 use multicast address MA_{2_W} only for receiving update messages from sub-region R_{2_W} while users U_5 and U_6 use it for sending update messages to users in sub-region R_{1_W} as far as Figure 5 is concerned. Table 2 summarizes multicast address assignment based on region and sub-region configuration in Figure 5. Using selective join/leaves to sub-region addresses based on proximity of users, we reduce the number of messages exchanged for inter-region interaction.

We can extend our mechanism to a more general model. Let us consider four adjacent regions which are part of a whole virtual world as shown in Figure 6. Our interest here is inter-region interaction between regions R_1 and R_3 or between R_1 and R_4. The same mechanism is simply applied to these cases. Each region is divided into eight sub-regions (two vertical and horizontal sub-regions, and four diagonal sub-regions). Each sub-region has a multicast address for inter-region interaction with users in all neighboring regions. For example, region R_1 has multicast addresses MA_{1_W} and MA_{1_E} for interaction with neighboring regions located horizontally such as region R_2, multicast address MA_{1_N} and MA_{1_S} for interaction with neighbor regions located vertically such as R_3, and multicast addresses MA_{1_NW}, MA_{1_NE}, MA_{1_SW}, and MA_{1_SE} for interaction with neighboring regions located diagonally such as R_4. For inter-region interaction, users use different multicast addresses based on the mechanism explained above. The vertical and diagonal case is the same as the horizontal one. Table 3 summarizes multicast address allocation according to the different interaction modes for users in region R_1 in Figure 6.

A user should send his update messages through all the multicast addresses for the sub-regions and the region to which he belongs; he should then join the appropriate multicast addresses for the sub-regions or the regions in which he/she is interested. For example, user U_1 in sub-region R_{1_NW} uses multicast address MA_{1_N} to send its update messages to users in sub-region R_{3_NW}, R_{3_NE}, and R_{4_NW}, and uses MA_1 to send its update messages to users in sub-region R_{3_SW}, R_{3_SE}, and R_{4_SW}. It uses multicast address MA_3 to receive all update messages from R_3 and uses MA_{4_W} to receive update messages from users in sub-regions R_{4_NW} and R_{4_SW}. User U_1 also joins MA_{2_W} for inter-region interaction with R_2. Figure 7 illustrates the change of an interested area (sub-regions) according to the position of a user in region R_1.

	Intra-region interaction	Inter-region interaction	
		Send	Receive
U_1, U_2	MA_1	MA_1	MA_{2_W}
U_3, U_4	MA_1	MA_1, MA_{1_E}	MA_2
U_5, U_6	MA_2	MA_2, MA_{2_W}	MA_1
U_7, U_8	MA_2	MA_2	MA_{1_E}

Table 2. Multicast addresses of intra- and inter-region interactions

DOI: 10.1002/cav

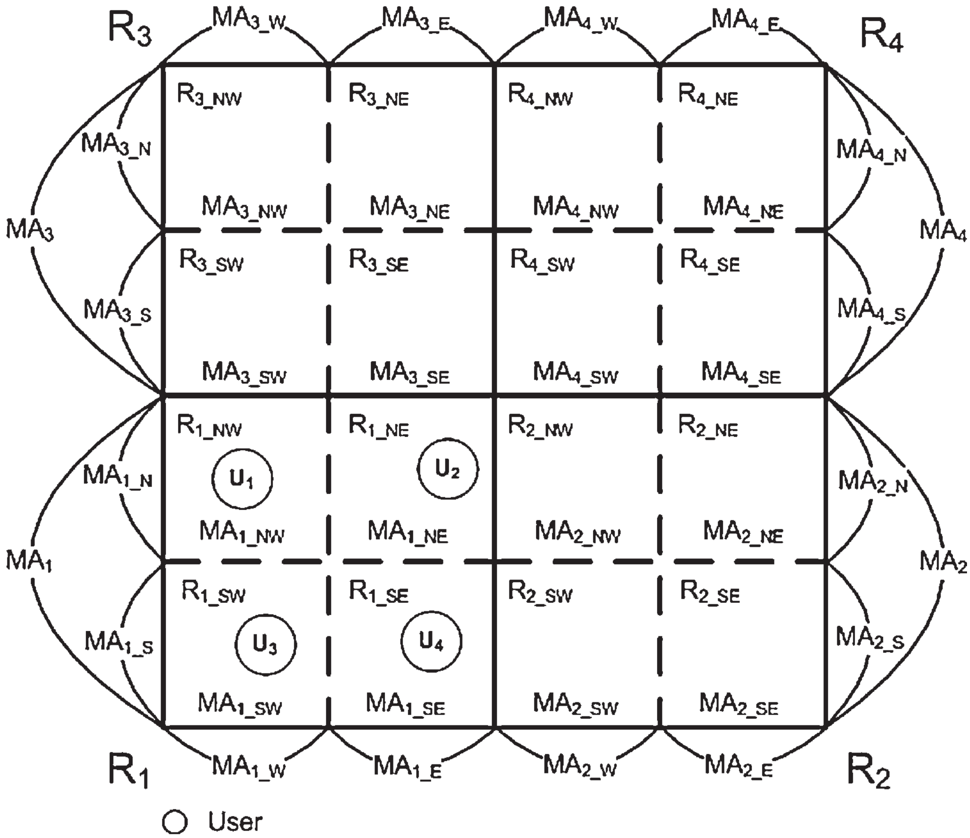

Figure 6. Extension of scalable inter-region interest management.

Interaction with users	In R_2		In R_3		In R_4	
Users in R_I	Send	Receive	Send	Receive	Send	Receive
U_I	MA_I	MA_{2_W}	MA_I, MA_{I_N}	MA_3	MA_I, MA_{I_N}, MA_{I_W}	MA_{4_W}
U_2	MA_I, MA_{I_E}	MA_2	MA_I, MA_{I_N}	MA_3	MA_I, MA_{I_N}, MA_{I_E}, MA_{I_NE}	MA_4
U_3	MA_I	MA_{2_W}	MA_I	MA_{3_S}	MA_I, MA_{I_S}, MA_{I_W}	MA_{4_SW}
U_4	MA_I, MA_{I_E}	MA_2	MA_I	MA_{3_S}	MA_I, MA_{I_S}, MA_{I_E}	MA_{4_S}

Table 3. Multicast addresses of inter-region interactions for users in region R_I

Adaptable Scope of Sub-Regions

If the users are uniformly distributed at all times, it is enough for the sub-region boundary to be fixed. However, a problem can occur when users in neighboring regions crowd in a specific region as shown in Figure 8. In Figure 8(a), a user U_1 in region R_1 cannot interact with others because they are all located outside of the IA. This results in excessive filtering, and is an unrealistic situation. In contrast, Figure 8(b) shows that all the neighboring users are together in the IA and thus sub-regions have no benefit.

Our inter-region interaction mechanism solves these problems by adjusting the scope of a sub-region boundary. Whenever a region manager in a peer-side receives a position update, join, or leave message in its region, it inspects the number of users in four distinct sub-regions. If the difference in terms of the number of

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

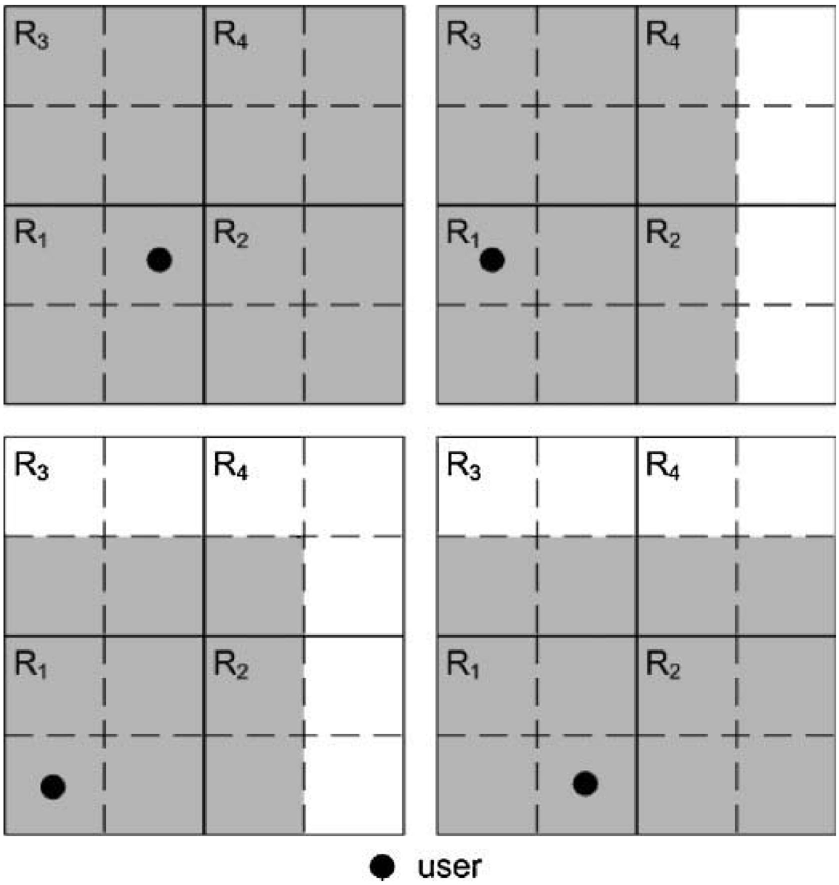

Figure 7. The change in interest to neighboring regions according to a user's position.

users in two or more sub-regions exceeds a specific threshold, the region manager calculates a new intersection point of sub-region boundaries by calculating the average (x, y) coordinate values of all the users in the region (currently we do not care about the height, which is the z axis, but it is simple to incorporate the z axis). This averaged coordinate becomes a new intersection point of the sub-region boundary of the user. Then, it updates the sub-region boundary information. If a user belongs to different sub-regions due to a change in the boundary information, then the user leaves the multicast groups assigned to the current sub-regions and joins new ones. This process is performed in a distributed manner in each peer-side so that it does not incur significant overheads server-side. Figures 8(c) and (d) shows the adjusted scope of the sub-regions in Figures 8(a) and (b), respectively. The threshold value triggering the adjustment of sub-regions is determined according to the application characteristics. If users move fast, the value must be high enough not to frequently resize the sub-regions. If users move slowly, the value becomes low because the difference in terms of the number of users is usually small.

An interest in neighboring regions can be further reduced using the gaze direction of a user. We leverage a limited field of a user's vision: that is, a user does not need to receive update messages from those who are out of his field of view. Table 4 defines the terminology required for describing our algorithm. We use a direction to define the relative position of a region and sub-region. If P(R, R′) is N, region R is the northern neighboring region of region R′. Sub(R, N) means a northern sub-region of region R. A user is then interested in neighboring regions Rv and Rv′. Rv always has higher priority when being selected than Rv′. If both v and v′ intersect with a neighboring region, the region becomes an element of Rv.

The interest level of region R to region r I(R, r), where r∈Rv, is then determined by the following algorithm, which is the same as the group communication model described in the previous section:

DOI: 10.1002/cav

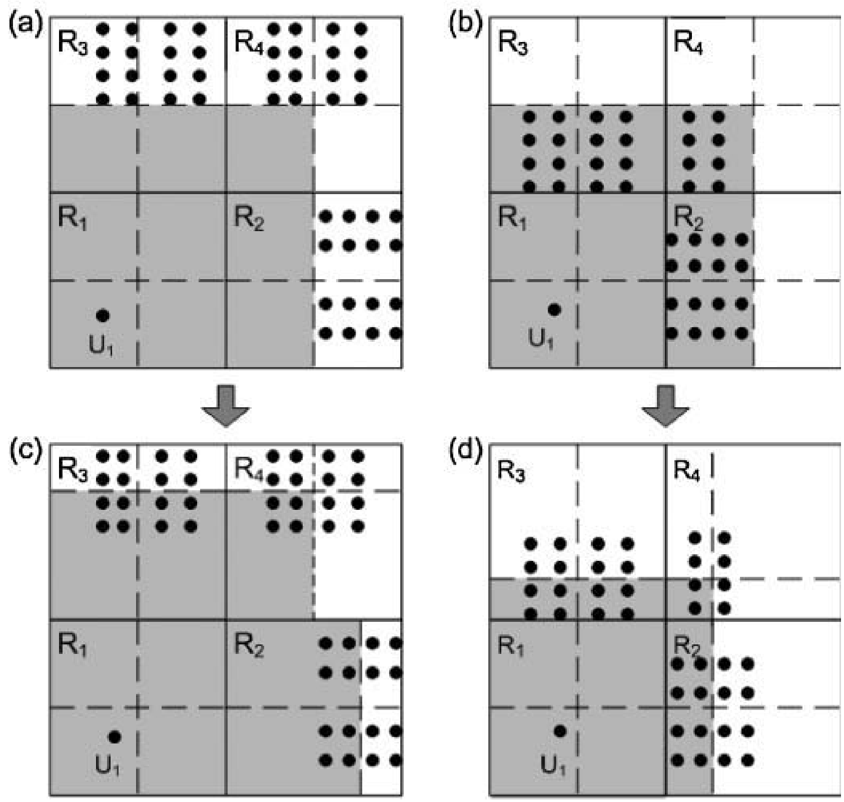

Figure 8. The adjustment of the sub-region boundaries to resolve skewed distribution of users.

Term	Definition
R	A region to which the current user belongs
Rv	A set of neighboring regions of region R that intersect with the gaze direction of the user, v
Rv′	A set of neighboring regions of region R that are within the range of the user's visibility, and that do not intersect with v
P(R, R′)	A position of region R relative to region R′, P(R, R′) $\in$ {N, S, W, E, NW, NE, SW, SE}
Sub(R, d)	A sub-region of region R, where d $\in$ {N, S, W, E, NW, NE, SW, SE}
Sub(u)	A set of sub-regions to which a current user u belongs
I(R, R′)	The interest level of region R to region R′, I(R, R′) = R′ or Sub(R′, d)

Table 4. Terminology

If Sub(u) ^ P(r, R) ! = NULL, I(R, r) = r
else I(R, r) = Sub(r, P(R, r))

If a user is close to a neighboring region, he is assumed to be interested in the entire neighboring region. If a user is far from a neighboring region, his interest in that region is limited to a sub-region of the neighboring region, which is determined by the relative position between the neighboring region and the region to which he belongs.

The interest level of region R to region r′, where r′∈Rv′, is also determined by the following algorithm:

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

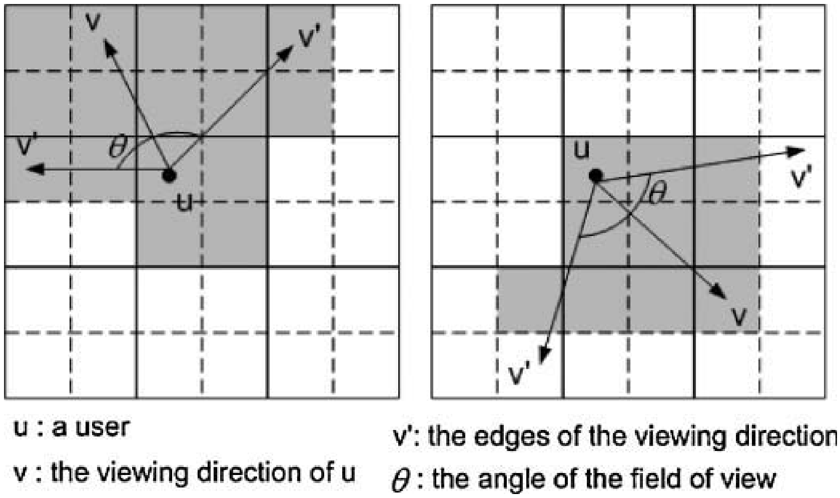

Figure 9. Change in interest to neighboring regions according to a user's viewing direction.

found = FALSE
While (r where r∈Rv) && (!found)
 If (Sub(u) ^ P(r, R)! = NULL) && (P(r, r′)∈{N, S, E, W})
  I(R, r′) = Sub(r′, P(r, r′)), found = TRUE
If (!found)
 I(R, r′) = Sub(r′, P(R, r′))

If a user's viewing direction is toward a nearby neighboring region, r, and r′ is one of the four (northern, southern, eastern, and western) neighboring regions of r, then the interest level of region R to region r′ is determined by the relative position between r and r′. Otherwise, the interest level is determined by the relative position between the current region R and r′. Figure 9 illustrates two examples of interesting sub-regions according to a user's viewing direction.

Experiments

In this section, the proposed scheme is evaluated by simulations in terms of communication and computational overheads. We also describe how the proposed scheme with varying experimental parameters is applicable to users' experiences in DVEs.

Experiment Environment

Our scalable sub-region based inter-region interest management scheme and user interest group based intra-region interest management scheme are simulated as follows. We implement existing inter-region (simple region-based filtering) and intra-region interaction mechanisms (aura-based filtering) and the proposed method on top of our DVE network framework, ATLAS,[24] using C++. ATLAS provides various scalable schemes in four scalability aspects: communication architecture, interest management, concurrency control, and data replication.

We construct a web-based virtual shopping mall application, the ICU/ETRI virtual shopping mall,[25] which uses ATLAS as a communication infrastructure and the RTV X3D viewer[26] as a browser as shown in Figure 10(a). To enter a virtual shopping mall, a user first connects to the web server, and downloads an applet and an X3D browser to set up a user interface. As a user selects a session and an avatar type which he/she is interested in, the static world data and world rule information are downloaded from the web server in order to initialize the virtual world. Using the world rule information, an ATLAS peer initializes dynamic objects and their properties obtained through EAI.[27]

We ran the experiment on three Windows XP machines (2.4 GHz Pentium IV with 512 MB of RAM) connected by 100 Mbps ethernet. The experiments give similar results to those performed on systems with even lower capacity (Pentium II Windows 2000 with 256 MB of RAM). Two machines run an ATLAS server and a web server, and another runs an ATLAS peer. The ICU/ETRI virtual shopping mall application emulates multiple users generated as non-playable characters by the ATLAS server. The users and their interest values are

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

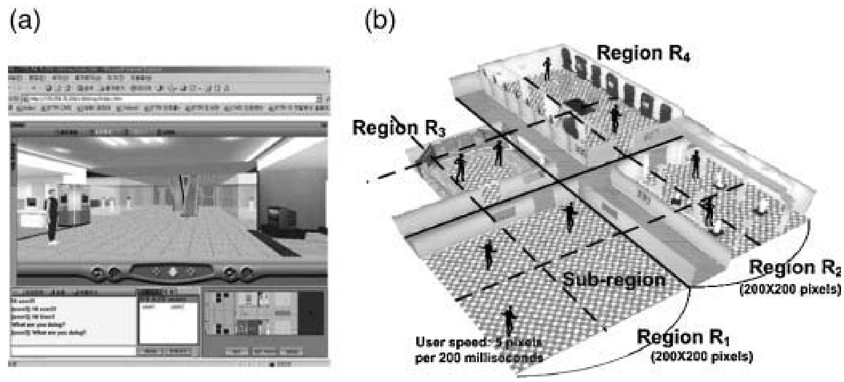

Figure 10. Experimental environment. (a) ICU/ETRI virtual shopping mall peer; (b) Virtual world configuration.

randomly created and move in random directions. We assume that the speed of users is fixed at five pixels per 200 milliseconds (we denote a pixel as the logical unit of distance in the virtual world). A user generates a maximum of five events per second lest messages should be flooded. To support scalability, we divide a virtual shopping mall into four logical regions as shown in Figure 10(b). The application then applies the intra-region and the inter-region interest management scheme together. The proposed schemes are compared with a combination of aura-based and region-based filtering schemes where a user receives messages only from others in the regions which collide with his/her aura. The hybrid approaches are not included because they propose specific policies, not mechanisms, for fine-grained filtering. We compared our schemes with the hybrid ones only in terms of the number of multicast addresses.

Communication Overhead

First, we evaluate the scalability in terms of the number of messages exchanged among various interest management schemes as the number of users increases. Figure 11 shows the average number of messages that a user receives per second as the number of users increases. In intra-region interaction, our proposed scheme results in fewer messages than the existing aura-based approach. The proposed scheme exploits the behavior of a user in the real world for reducing required communication resources in DVEs. When users are clumped in a specific place, the user is not interested in what all the others around him do and say, rather each user concentrates on others with the same interest. In inter-region interaction, as the number of users increases, our proposed scheme requires fewer messages than the existing aura-based scheme. Even if a user can see objects 200 m away from him, the user may not see them because other objects in front of them are densely populated. In addition, a user cannot see objects behind him. The proposed inter-region interest management scheme takes advantage of this to reduce the bandwidth requirement.

Users in DVEs may be in a heavily or sparsely populated place. When users crowd in a shop, it is hard for a user to see even other users nearby him. The user can then shrink his/her IA. On the other hand, if there are few users, a user can expand his/her IA which may span several shops. Figure 12 shows the relationship between the number of received messages and various IA radii. Since the region size (200×200 pixels) is fixed, the larger an IA radius becomes, the more messages a user receives. In Figure 12, the proposed scheme reduces more number of messages than the existing aura-based scheme even when users change the size of their IA according to various populations. When the IA radius becomes larger than the segment of a region in the existing approach, the number of received messages does not increase, because the aura always collides with all the neighboring regions.

Our intra-region interest management scheme affects the number of received messages according to the rate of low-fidelity message transmission as shown in Figure 13. A user can select the rate according to his/her intention in a shopping mall. If a user wants to see detailed behavior for other users (such as the detailed motion of avatars) who have different interests from him/her, a user sets a high rate for low-fidelity messages. For example, when the rate is 100% of the

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

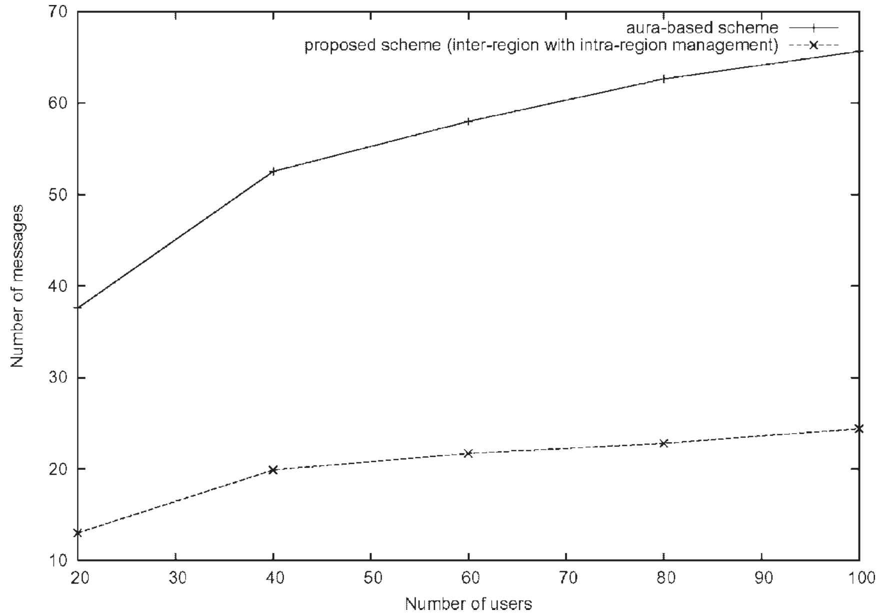

Figure 11. Average number of messages received per 1 second (the rate of a low frequency message = 30%, an IA radius = 100 pixels, and the number of interests = 5).

high-fidelity message rate, it means that a user wants to see all the interaction among others in his/her IA with high fidelity. However, in this extreme case, the user receives all the messages without any filtering. On the other hand, if a user does not want to see the detailed behavior of others with different interests, a user sets a low rate. When the rate is 0%, it means that a user wants to interact only with others who have the same interest with him. However, a user never knows where other users with different interests are, which breaks percep-

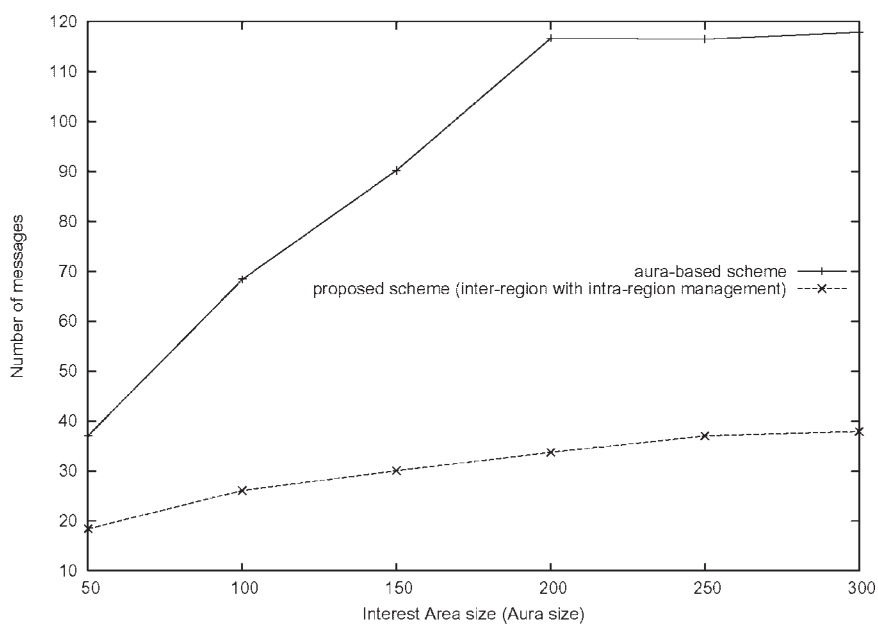

Figure 12. Average number of messages according to the IA radius (the rate of a low frequency messages = 30%, the number of users = 100, and the number of interests = 5).

DOI: 10.1002/cav

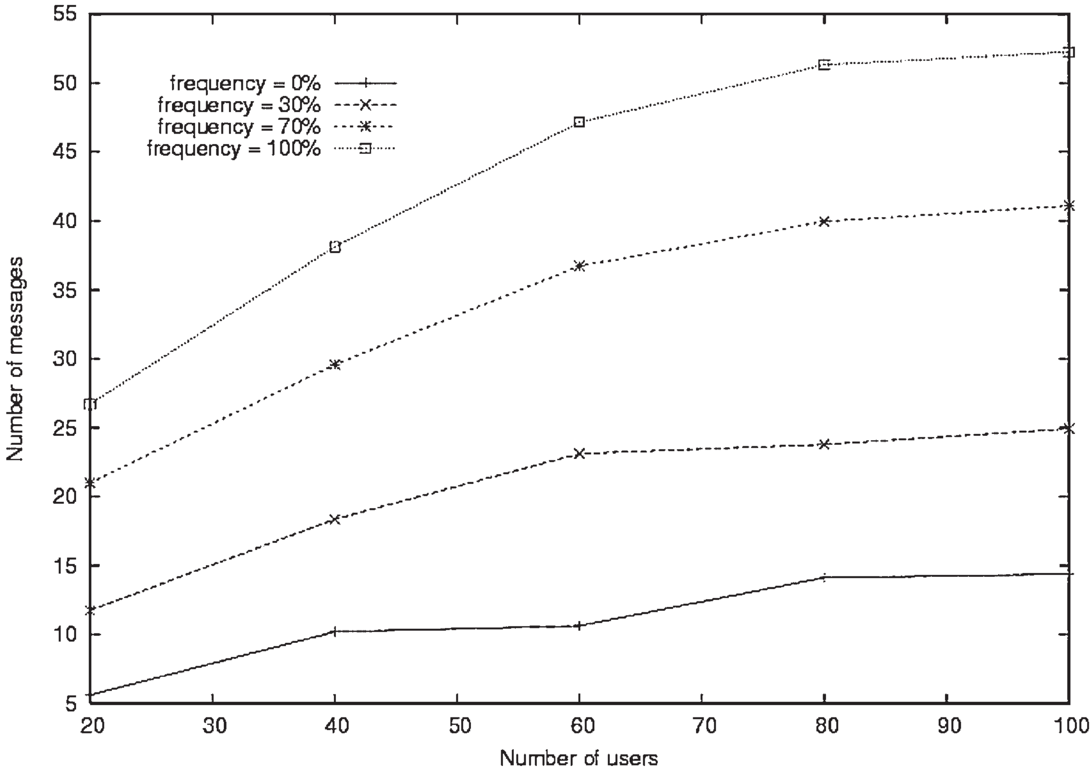

Figure 13. Average number of messages with the different rates for low frequency messages (an IA radius = 100 pixels and the number of interests = 5).

tual consistency among users with different interests. Therefore, the transmission rate of low-fidelity messages is an important factor which controls the tradeoff between perceptual consistency and bandwidth usage. If a user does not want to see the detailed interaction of others with different interests but rather their rough behaviors, he selects a low frequency rate between the above two extremes. We set the low frequency rate to 30% of the high frequency rate, since it is low enough for a user to recognize in which direction other users move. If the rate is lower than 30%, a user misperceives other users, who seem to appear and disappear arbitrarily.

Another factor which affects the number of received messages of the proposed scheme is the total number of interests as shown in Figure 14. As the number of interests increases, the probability of high-fidelity interaction becomes low since few users have the same interests. If a virtual shopping mall provides many products like a department store, users have different interests. On the other hand, if a shopping mall sells only one kind of product, all users have the same interest. The number of interests does not have an effect on the existing aura-based schemes because they perform only proximity-based filtering. Compared with the existing scheme, the proposed scheme shows better performance because it filters messages based not only on proximity, but also on users' interest, population density, and viewing direction.

The proposed schemes can also be optionally used to support various filtering levels from coarse-grained to fine-grained as in a multi-level filtering scheme. If a user wants an overall view of a virtual world (which means that a user is not interested in a specific place), only the inter-region interest management scheme is required because it filters unnecessary messages in the current region and its neighboring regions. Without this scheme, a user receives all messages from neighboring regions even though the user cannot see them. For example, when a user enters a shopping mall and wants to know which shop is heavily populated, the user does not need to be hampered by all the traffic generated from those who are clumped in a distant shop. On the other hand, if a user visits a populated appliance shop and wants to buy a television, the user is interested in televisions and

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

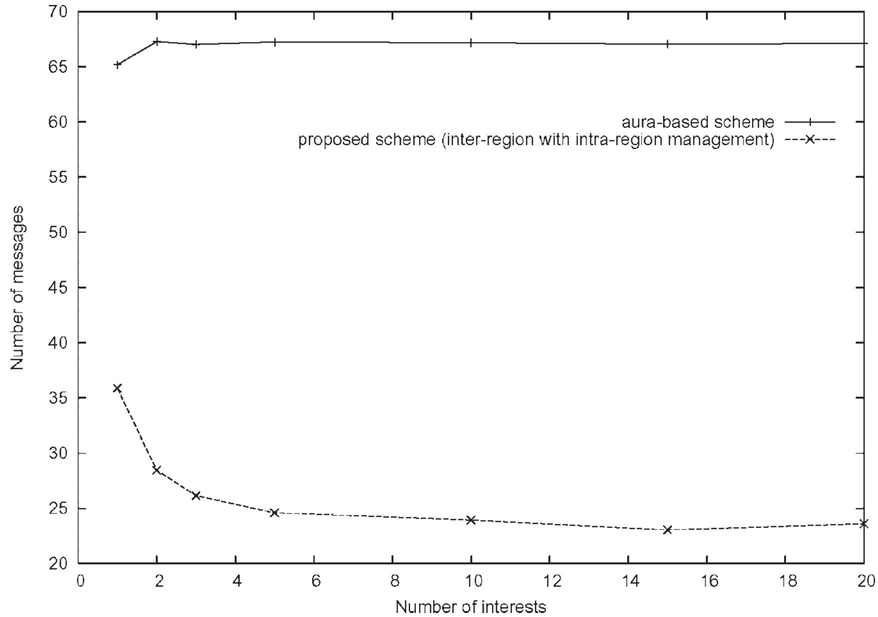

Figure 14. Average number of messages according to the number of interests (the rate of a low frequency messages = 30%, the number of users = 100, and an IA radius = 100 pixels).

which models other shoppers in the shop see and buy. Because there are lots of users and appliances in a relatively small place, it is helpful for the intra-region interest management scheme to reduce uninteresting messages generated from those who have different interests. Without this mechanism, far from buying a television, a user may have difficulty even in moving just one step. Another strategy for supporting various filtering levels can be to simply adjust the level depending on the number of users.

Figure 15 shows the results from the experiments which apply the proposed scheme with different filtering levels. If we use only the inter-region interaction management without adaptable scoping, messages are filtered only by sub-regions in neighboring regions and the filtering level is low. When the adaptable scoping is added to inter-region interest management and both inter-region and intra-region interest management are combined, the filtering level becomes high. If the proposed scheme uses only the inter-region interest management, the number of received messages is similar to but still higher than that of the aura-based scheme, because our scheme in this level is performed only by sub-regions and adaptable scoping, and does not consider the collision between an IA and regions, and neither does it consider users' interests. However, the most fine-grained filtering version of the proposed scheme (intra-region interest management combined with inter-region interest management) outperforms the existing approach due to the separation of high and low fidelity interaction based on users' interests.

Computational Overhead

In the second experiment, we evaluate the computational overhead incurred by various interest management schemes. While the proposed scheme reduces the number of messages, it may appear to incur more local processing overhead than the existing schemes. The overhead results from the sub-region check, switching from one multicast address to another for inter-region interaction, and from dynamic grouping of users based on IA and interests in the intra-region interaction. Figure 16 shows the differences of the local processing time of an update message on average. The computational overhead is less than a few milliseconds regardless of the number of users.

Resizing sub-regions and joining sub-regions in neighboring regions based on a user's viewing direction are local operations within a peer in our inter-region interest management, so there is no need to re-configure regions, which would require changing the total number of regions, reforming multicast groups,

DOI: 10.1002/cav

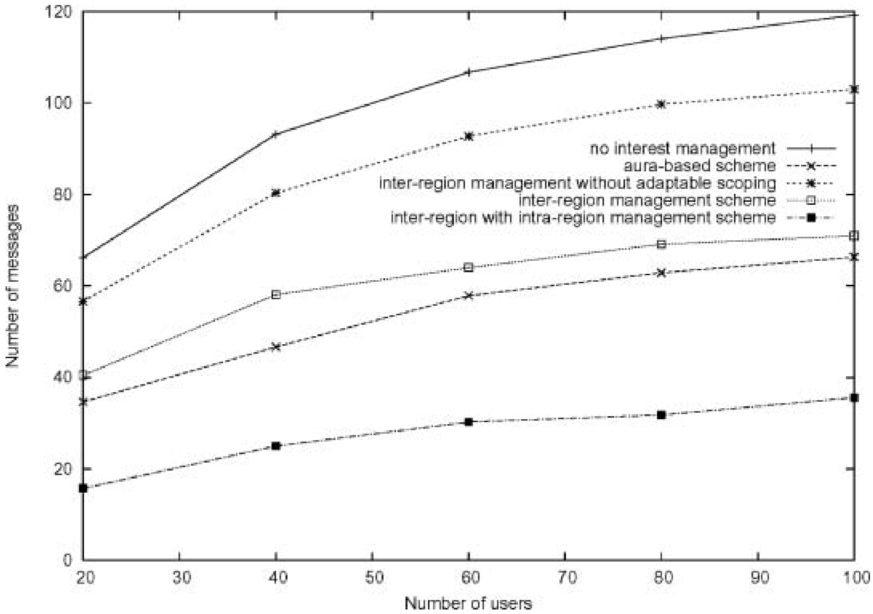

Figure 15. Average number of messages according to the filtering level (the rate of a low frequency messages = 30%, an IA radius = 100 pixels, and the number of interests = 5).

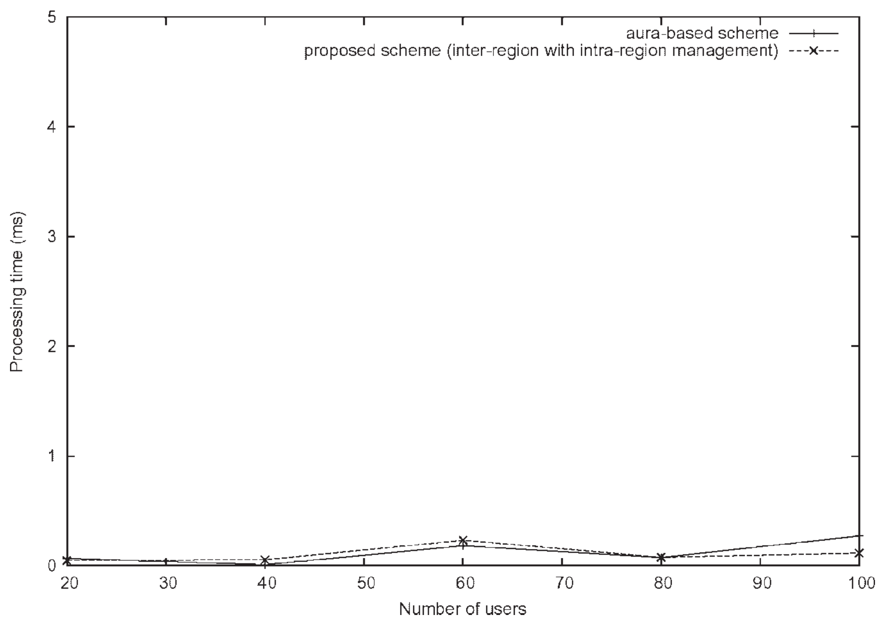

Figure 16. Average computational overhead per message in peers (the rate of a low frequency messages = 30%, an IA radius = 100 pixels, the number of interests = 5).

and notifying all affected users. Therefore, there is no significant difference in the computational overhead between our proposed scheme and the existing aura-based scheme.

Usefulness of Filtering

In this sub-section, we analyze how much the proposed scheme is useful from the viewpoint of users. We focus

DOI: 10.1002/cav

on how many interesting messages users miss, and how many uninteresting ones they receive. A user generally sees objects far from him/her although the user cannot hear any sound from them, and cannot see objects behind him/her. On the other hand, a user hears what others around him say. From this, we assume that a user has a fanwise view range and a circular audio range. Matching the interest range with that in DVEs, we compare the proposed scheme with existing one in terms of the usefulness of received messages.

In the proposed scheme, inter-region interest management can be applied to the view range of a user, and intra-region interest management to the audio range. However, the existing scheme specifies that a user is interested in regions with which his IA collides, which is broader than that of the proposed scheme. Figure 17 shows examples of the comparison. The dark gray area means the range from which a user cannot receive messages even if the user is interested in them, and the light gray area shows the range from which a user

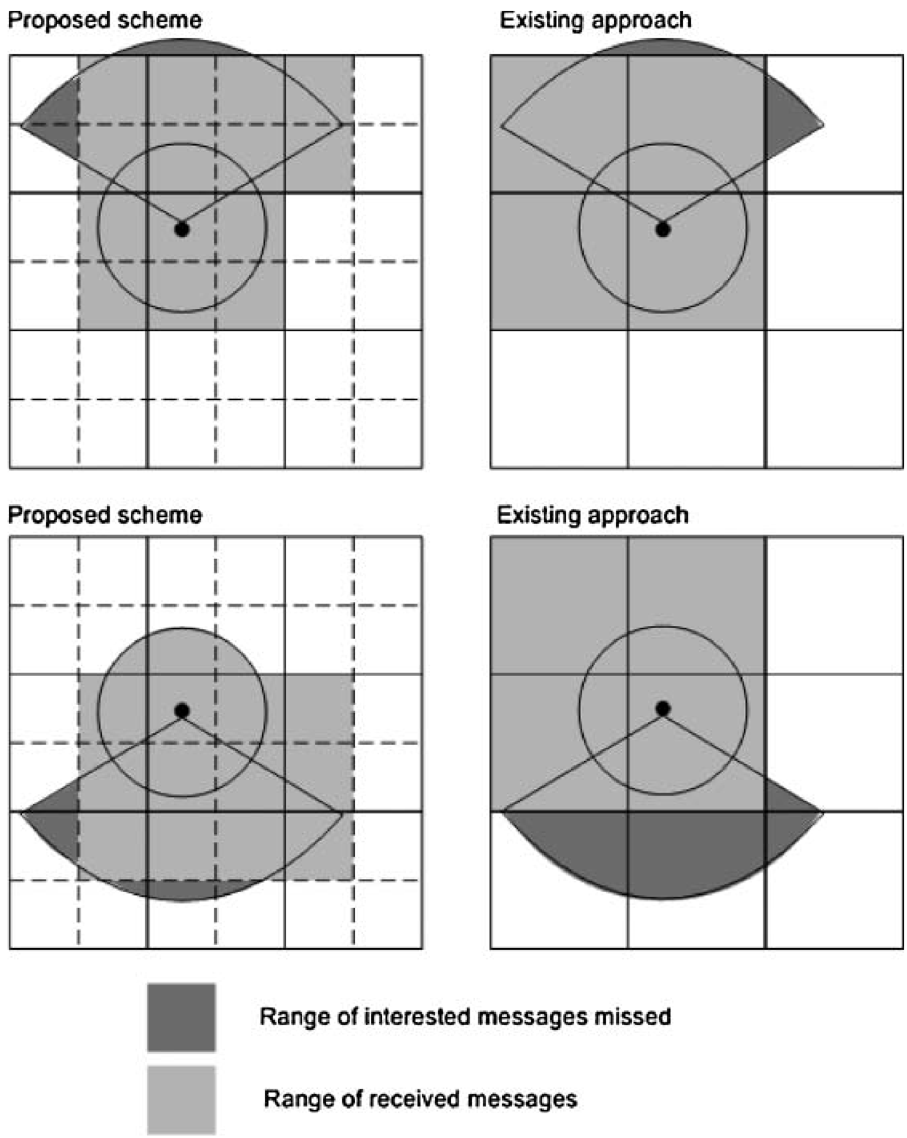

Figure 17. Excessive and missed messages.

receives messages including excessive ones. Since the proposed scheme is approximately adaptable to the change of users' IAs using sub-regions and low-fidelity transmission, the number of excessive and missed messages is less than that of the existing scheme.

One of ways to eliminate all excessive and missed messages is to make each user have his/her own multicast address, which enables the most fine-grained filtering. However, assigning a multicast address per user is wasteful because it is a limited resource. The proposed scheme is an approximation to fine-grained filtering, while it reduces the number of multicast addresses as discussed in the following sub-section.

Use of Multicast Addresses

Lastly, we analyze the usage of multicast addresses. For simplicity of comparison, we assume that a region is a 100×100 rectangle, that a user's IA radius is 10, that a hundred users are located in the region, and that all users have the same interest type. In the region-based filtering scheme, one multicast address is assigned to a region since the region forms one multicast group. Although it uses a small number of addresses, a user joining a multicast group must receive unnecessary messages. The hybrid approach is another extreme case such that it must assign 100 multicast addresses to a region because each user has his/her own address while it filters messages in a most fine-grained manner. The aura-based filtering scheme is somewhere between the above two extremes. The number of multicast addresses is determined by the size of an IA. If the radius is 0, there is no one who is interacting with the others, and each user has his/her own multicast address. If an IA covers a whole region, there is only one multicast group because all the users in the region are located in the IAs of others. According to the above assumptions (a region is a 100×100 rectangle, and a user's IA radius is 10), the radius of an IA is one-tenth of a segment of a region. Twenty-five multicast groups are then formed in a region in the worst case. Our scheme is comparable to both the region-based and aura-based methods in terms of address usage. The proposed inter-region interest management scheme assigns multicast addresses per region and sub-region. Because a region has eight sub-regions, our inter-region interaction management scheme requires nine addresses per region. The proposed intra-region interest management scheme forms an interest group only when users are located in each other's IA with the same interest. If all users have different interests, no additional address is required since no interest group is formed. In the worst case, if there are 50 interests and each two of a 100 users have the same interest, our scheme forms 50 interest groups and requires the same number of additional addresses. This means that the required number of addresses in the worst case is half the number of users. Although it assigns more addresses (59) than the aura-based approach (25), the proposed schemes use a much smaller number of addresses than the hybrid approach (100).

In summary, the proposed scheme reduces the number of messages with no significant computational overhead compared with the existing region-based and aura-based schemes since the upper bound of acceptable interactive performance is 100 ms[2]. This implies that a user receives required messages within a time bound. In addition, the proposed schemes use multicast addresses more economically than the hybrid approaches.

Conclusion

We have proposed a new interest management scheme that enhances the scalability of DVE applications where many users with different interests often crowd in a specific place. It consists of two parts: user interest group-based intra-region interest management and sub-region-based inter-region interest management.

The key idea of the proposed intra-region interest management scheme is that users with the same interests are dynamically grouped when they are included in each other's IA. An RU is elected among them by priority based on characteristics like user ID or how long they have stayed in the region. The group members communicate with each other with high fidelity. The RU, on behalf of the group members, transmits updated member's interaction data with low frequency to other users who have different interests. The proposed inter-region interest management scheme improves scalability through dividing a region into several sub-regions, and by sending update messages from these sub-regions to users in neighboring regions. It also makes the scope of sub-regions adaptable to the distribution of various users and a user's viewing direction in order to support consistent interest in the part of neighboring regions. The proposed scheme improves scalability by reducing the number of messages without significant computational overhead, perceptual inconsistency, or extravagant use of multicast addresses.

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav

ACKNOWLEDGEMENTS

This research is supported by the ubiquitous Computing and Network (UCN) Project, the Ministry of Information and Communication (MIC) 21st Century Frontier R&D Program in Republic of Korea.

References

1. Macedonia MR, Zyda MJ. A taxonomy for networked virtual environments. *IEEE Multimedia* 1997; **4**(1): 48–56.
2. Singhal S, Zyda M. *Networked Virtual Environments Design and Implementation*. Addison Wesley: New York, NY, USA, 1999.
3. Van Hook DJ, Rak SJ, Calvin JO. Approaches to relevance filtering. In *Proceedings of 11th Workshop on Standards for the Interoperability of Distributed Simulation*, September 26–30, 1994; pp. 367–369.
4. Morse K, Bic L, Dillencourt M. Interest management in large-scale virtual environments. *Presence: Teleoperators and Virtual Environments* 2000; **9**(1): 52–68.
5. Barrus JW, Waters RC, Anderson DB. Locales: supporting large multiuser virtual environments. *IEEE Computer Graphics and Applications* 1996; **16**(6): 50–57.
6. Macedonia M, Zyda M, Pratt D, Brutzman D, Barham P. Exploiting reality with multicast groups. *IEEE Computer Graphics and Applications* 1995; **15**(5): 38–45.
7. Capps M, McGregor D, Brutzman D, Zyda M. NPSNET-V: a new beginning for dynamically extensible virtual environments. *IEEE Computer Graphics and Applications* 2000; **20**(5): 12–15.
8. Steed A, Abou-Haidar R. Partitioning crowded virtual environments. In *Proceedings of ACM Symposium on Virtual Reality Software and Technology*, Osaka, Japan, October 1–3, 2003; pp. 7–14.
9. Greenhalgh C, Benford S. MASSIVE: a collaborative virtual environment for teleconferencing. *ACM Transactions on Computer-Human Interaction* 1995; **2**(3): 239–261.
10. Hagsnad O. Interactive multiuser VEs in the DIVE system. *IEEE Multimedia* 1996; **3**(1): 30–39.
11. Frécon E, Smith G, Steed A, Stenius M, Sthål O. An overview of the COVEN platform. *Presence: Teleoperators and Virtual Environments* 2001; **10**(1): 109–127.
12. IEEE Standard for Modeling and Simulation [M and S]. High level architecture [HLA] – federate interface specification. *IEEE Std, 1516.1-2000*, 2001.
13. Ding D, Zhu M. A model of dynamic interest management: interaction analysis in collaborative virtual environment. In *Proceedings of ACM Symposium on Virtual Reality Software and Technology*, Osaka, Japan, October 1–3, 2003; pp. 223–230.
14. Liu C, Wen C, Chen H. Tracking-needless grouping: an efficient and scalable grouping scheme in networked virtual environments. In *Proceedings of The First IEEE Consumer Communications and Networking Conference*, Las Vegas Nevada, USA, January 5–8, 2004; pp. 477–482.
15. Greenhalgh C, Purbrick J, Snowdon D. Inside MASSIVE-3: flexible support for data consistency and world structuring. In *Proceedings of ACM Collaborative Virtual Environments*, San Francisco, CA, USA, September, 2000; pp. 119–127.
16. Léty E, Turletti T, Baccelli F. SCORE: A scalable communication protocol for large-scale virtual environments. *IEEE/ACM Transactions on Networking* 2004; **12**(2): 247–260.
17. Abrams HA. Extensible interest management for scalable persistent distributed virtual environments. *Ph.D. dissertation*. Naval Postgraduate School, Monterey, California, United States, 1999.
18. Oliveira J, Georganas N. VELVET: an adaptive hybrid architecture for very large virtual environments. *Presence: Teleoperators and Virtual Environments* 2003; **12**(6): 555–580.
19. Pryce N. Group management and quality of service adaptation in distributed virtual environments. In *Proceedings of 4th UK VR-SIG Conference*, November, 1997.
20. Singhal SK, Cheriton DR. Using projection aggregations to support scalability in distributed simulation. In *Proceedings of 16th IEEE International Conference on Distributed Computing Systems*, Hong Kong, May 27–30, 1996; pp. 196–206.
21. Benford S, Greenhalgh C. Introducing third party objects into the spatial model of interaction. In *Proceedings of Fifth European Conference on Computer Supported Cooperative Work*, Lancaster University, UK, September 7–11, 1997; pp. 189–204.
22. Ng B, Si A, Lau R, Li F. A multi-server architecture for distributed virtual walkthrough. In *Proceedings of ACM Symposium on Virtual Reality Software and Technology*, Hong Kong, China, November 11–13, 2002; pp. 163–170.
23. Lim M, Lee D. A fine-grained load distribution scheme for multi-server-based distributed virtual environment systems. In *Proceedings of CollabTech2005*, Tokyo, Japan, July 15–16, 2005; pp. 34–39.
24. Lee D, Lim M, Han S. ATLAS – a scalable network framework for distributed virtual environments. In *Proceedings of ACM Collaborative Virtual Environments*, Bonn, Germany, September 30–October 2, 2002; pp. 47–54.
25. Han S, Lim M, Lee D, *et al.* Scalable network support for 3D virtual shopping mall. In *Proceedings of International Conference on Virtual Systems and Multimedia*, Gyeongju, Korea, September 25–27, 2002.
26. Real Time Visual Co. X3dbrowser: http://www.realtimevisual.com.
27. Carey R, Bell G, Marrin C. The Virtual Reality Modeling Language (VRML) - Part 2: External Authoring Interface (EAI). *ISO/IEC 14772-2*, 2001.

Comp. Anim. Virtual Worlds 2008; **19**: 129–149
DOI: 10.1002/cav