Cell-based Multicast Grouping in Large-Scale Virtual Environments

[Extended Abstract] *

Emmanuel Léty
INRIA Sophia-Antipolis, FRANCE
Emmanuel.Lety@inria.fr

Thierry Turletti
INRIA Sophia-Antipolis, FRANCE
Thierry.Turletti@inria.fr

François Baccelli
INRIA-Ecole Normale Superieure, FRANCE
Francois.Baccelli@inria.fr

1. INTRODUCTION

This paper describes the motivations and the design of a communication architecture for Large-Scale Virtual Environments (LSVE) on the Internet. Such Virtual Environments (VE) include massively multi-player games, Distributed Interactive Simulations (DIS), and shared virtual worlds. We consider only many-to-many applications, where each participant is both source and receiver. Using IP multicast solves part of the scalability problem by allowing each source to send data only once to all the participants. However, with a large number of heterogeneous users, transmitting all the data to all the participants dramatically increases the probability of congestion within the network and particularly at the receiver side. Moreover, [4] shows that in a group communication setting, as the number of flows of data and the number of users increase, the percentage of content that participants are interested in receiving decreases. We argue that the superfluous received traffic has to be filtered out before it reaches the end-host. The main difficulty in this filtering mechanism comes from the heterogeneity and the dynamicity of the receivers, not only in terms of bandwidth and processing power but also in terms of data interest, virtual and physical locations.

In this paper, we propose a filtering mechanism done at the transport-layer, using multiple multicast groups and agents. Our approach involves the dynamic partitioning of the VE into spatial areas called *cells* and the association of these cells with multicast groups. We describe a method, based on the theory of planar point processes, to determine an appropriate *cell-size* so that the incoming traffic at the receiver side remains with a given probability below a sufficiently low threshold. We then propose mechanisms to dynamically partition the VE into cells of different sizes, depending on the density of participants per cell, the number of available multicast groups, and the link bandwidth and processing power available per participant.

2. THE CELL-SIZE TRADEOFF

In this paper, we focus on the *cell-based grouping strategy* which basically consists in partitioning the VE into cells and assigning to each cell a multicast group. During the session, each participant identifies the cell it is currently "virtually" located in, and sends its data to the associated multicast group. To receive the data from the other participants included in its area of interest, each participant has to join the multicast groups associated with the cells that intersect its area of interest. Similarly, when a participant moves, it needs to leave the multicast groups associated with the cells which do not intersect anymore its area of interest. The cell-based grouping strategy is particularly suitable on VE that can easily be partitioned into virtual areas (e.g. virtual Euclidean spaces). However, the main difficulty in this partitioning is to find the appropriated cell-size. Decreasing the cell-size increases the overhead associated with dynamic group membership whereas increasing the cell-size increases the unwanted information received per participant [2].

Two approaches are possible to estimate the best cell-size in a LSVE. The first approach requires the pre-calculation of a static cell-size parameter, which remains the same during the whole session. The second approach consists in dynamically re-estimating the cell-size during the session, taking into account dynamic parameters. To motivate the choice of one of these two approaches, we identified the parameters involved in the cell-size calculation and then, we examined the impact of the dynamicity of these parameters on the appropriate cell-size. The **number of available multicast groups** is an important parameter to take into account for the cell-size calculation because it gives a lower bound of the cell-size. As the number of multicast groups used is inversely proportional to the size of the cell, a small set of available multicast groups will lead to a bigger cell-size. The **receivers capabilities** identify the link capacities and the processing power available per receiver. Assuming each user roughly generates the same amount of traffic, the incoming traffic per receiver grows linearly with the total number of sources contained in the multicast groups to which it has subscribed. Nevertheless, some of these participants may be located outside the area of interest but inside a cell that

*A full version of this paper is available at `www.inria.fr/rodeo/elety` includes this area of interest. The ratio between the corresponding number of unwanted participants and the total number of sources received represents the percentage of *superfluous* traffic received. So, the cell-size and more particularly the ratio between the cell-size and the size of the area of interest, have a direct impact on the amount of unwanted traffic. The **density of participants** represents the ratio between the number of participants and the size of the VE. In the cell-based grouping strategy, the area of interest is approximated by the smallest set of cells covering the area of interest. In the rest of the paper, we refer to the difference between these two areas as the *superfluous area*. So, the density of participants in a VE not only has an impact on the average number of participants located in the area of interest, but also on the superfluous area. A smaller cell-size could allow a better approximation of the area of interest and a significant reduction of superfluous area and its corresponding traffic. The **participant velocity** can be used in a cell-based grouping strategy, to estimate the bandwidth overhead generated when participants cross cells, and the ratio between the join and leave latency, and the mean time that the participant stays per cell.

3. MODELS AND SIMULATIONS

We introduce a model of area of interest and of participants based on random point processes which is inspired of the stochastic geometry approach proposed in [1]. We make the assumption that the participants are static and located on the plane according to a random homogeneous Poisson point process[3]. We show from simulations based on our model, that for a given intensity of participants, it is possible to find the largest cell-size (i.e., the smallest number of multicast groups), so that the incoming traffic remains bellow a sufficiently low threshold with a high probability. This probability reflects the tradeoff between the satisfaction of the users and the required number of multicast groups. These results argue in favor of a dynamic partitioning of VE into cells of different sizes.

We also introduce a model of mobility which is compatible with the assumption that the point process of participants is Poisson at any time. From this model, we perform simulations, which show the mean residence time per cell as a function of the participant mean velocity. We observe that the mean residence time decreases exponentially as the mean velocity approaches 1 cell-size per second. This result argues in favor of a limited velocity in LSVE, so that the residence time per cell remains higher by orders of magnitude than the join latency. Indeed, a participant needs to anticipate its join request by subscribing to the multicast groups which map the cells where it can go during the time corresponding to a join latency. Hence, its velocity and the cell-size impact on the number of multicast groups it needs to join by anticipation, and therefore on the IGMP traffic generated.

4. FRAMEWORK FOR A SCALABLE COMMUNICATION ARCHITECTURE

In this section, we describe a framework for a scalable communication architecture for LSVE, taking the simulation results into consideration. The goal of this architecture is to make VE scalable with thousands of heterogeneous users on the Internet. Moreover, we claim that this solution works with a limited number of available multicast groups. The aim of our work is not to propose a new IP multicast model neither to come up with a network-layer approach, adding new mechanisms in the routers. Alternatively, we present a transport-layer solution with multiple agents, assuming that all the users are capable of receiving multicast transmissions. We define *agents* as servers or processes running at different parts of the network (e.g., on a campus LAN, hosted by an ISP or by LSVE developers). Administrators of LSVE are responsible for deploying such agents on the Internet and for positioning them as close as possible to their potential users. Agents have to dynamically determine appropriate cell-size values in order to maximize the users satisfaction. During the session, four successive operations are required: First, to partition the VE into several *zones*, according to the distribution of users, and the VE structure (e.g., rooms, walls, etc.). Second, to compute the appropriate cell-size for each zone, according to the parameters listed in Section 2. Third, to divide each zone into cells, according to its computed cell-size, and assign a multicast group address to each cell of each zone. And last, to inform the participants of which multicast groups they need to join in order to interact with participants located around them.

Each participant is connected to a single agent, using a unicast connection. Agents communicate with each other on a dedicated multicast group. Throughout the session, agents periodically compute the average density of participants per multicast group, by dividing the number of connected participants with the number of available multicast groups for the application. This average density is then compared with the different densities of participants in each zone. Exploiting these comparisons, agents can decide to increase or decrease the cell-size in different parts of the VE. [5] describes the different messages exchanged between participants and agents, and details the different mechanisms involved in our architecture. In order to evaluate the performance of our dynamic mapping strategy, we compared it with a static partitioning strategy using simulations. We show its adaptive capacity to a dynamic VE, even with a limited number of available multicast groups.

5. REFERENCES

[1] F. Baccelli and S. Zuyev. Poisson-Voronoi spanning trees with applications to the optimization of communication networks. *in Operations Research*, 1999.

[2] D. J. V. Hook, S. J. Rak, , and J. O. Calvin. Approaches to relevance filtering. In *Proceedings of 11th DIS Workshop*, September 1994.

[3] J. Kingman. *Poisson processes.* Oxford Studies in Probability, No3. Clarendon Press, 1993.

[4] B. Levine, J. Crowcroft, C.Diot, J. Garcia-Luna-Aceves, and J. Kurose. Consideration of receiver interest in content for ip delivery. In *Proceedings IEEE INFOCOM 2000*, March 2000.

[5] E. Léty and T. Turletti. Issues in designing a communication architecture for large-scale virtual environments. In *Proceedings of the 1st Workshop on Networked Group Communication*, November 1999.