University of Wollongong
Research Online

Faculty of Informatics - Papers (Archive)

Faculty of Engineering and Information Sciences

2008

Supporting a seamless map in peer-to-peer system for massively multiplayer online role playing games

X. Jiang
University of Wollongong

Farzad Safaei
University of Wollongong, farzad@uow.edu.au

Publication Details

This conference paper was originally published as Jiang, X and Safaei, F, Supporting a seamless map in peer-to-peer system for massively multiplayer online role playing games, 33rd IEEE Conference on Local Computer Networks, LCN 2008, Montreal, Canada, 14-17 October 2008. Copyright Institute of Electrical and Electronics Engineers 2008. Original conference paper available here

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Supporting a seamless map in peer-to-peer system for massively multiplayer online role playing games

Abstract

Massively Multiplayer Online Role Playing Games feature huge maps that can be partitioned into smaller components called zones to achieve scalability. In peer-to-peer (P2P) systems, those zones can be assigned to peers who are willing to take up the role of zone servers. A seamless map requires inter-zone communications to allow interactions between players across zone boundaries. In assigning zones to peers, it is critical to seek approaches to reduce the inter-zone communication cost in P2P system. In this paper, we formulate the zone assignment problem for P2P system and demonstrate that the problem is NP-hard. We hence propose a low-cost heuristic that partitions the physical network into bins, as well as, aggregating the partitions of the virtual world by clustering neighboring zones in a game map. We then demonstrate that by assigning clustered neighboring zones in the virtual map into peers in common physical network bins, we are able to achieve similar inter-zone communication cost for the zone assignment problem as that can be achieved by the optimization model we formulated, while avoiding the prohibitively high computation cost required by the latter.

Disciplines

Physical Sciences and Mathematics

Publication Details

This conference paper was originally published as Jiang, X and Safaei, F, Supporting a seamless map in peer-to-peer system for massively multiplayer online role playing games, 33rd IEEE Conference on Local Computer Networks, LCN 2008, Montreal, Canada, 14-17 October 2008. Copyright Institute of Electrical and Electronics Engineers 2008. Original conference paper available here

This conference paper is available at Research Online: http://ro.uow.edu.au/infopapers/683

Supporting a Seamless Map in Peer-to-Peer System for Massively Multiplayer Online Role Playing Games

Xinbo Jiang, Farzad Safaei
TITR, University of Wollongong, Australia

Abstract – **Massively Multiplayer Online Role Playing Games feature huge maps that can be partitioned into smaller components called zones to achieve scalability. In peer-to-peer (P2P) systems, those zones can be assigned to peers who are willing to take up the role of zone servers. A seamless map requires inter-zone communications to allow interactions between players across zone boundaries. In assigning zones to peers, it is critical to seek approaches to reduce the inter-zone communication cost in P2P system. In this paper, we formulate the zone assignment problem for P2P system and demonstrate that the problem is NP-hard. We hence propose a low-cost heuristic that partitions the physical network into bins, as well as, aggregating the partitions of the virtual world by clustering neighboring zones in a game map. We then demonstrate that by assigning clustered neighboring zones in the virtual map into peers in common physical network bins, we are able to achieve similar inter-zone communication cost for the zone assignment problem as that can be achieved by the optimization model we formulated, while avoiding the prohibitively high computation cost required by the latter.**

I. Introduction

The development and growth of online games have been phenomenal in recent years. Online games such as Lineage [1] , World of WarCraft [2] and EverQuest [3] have attracted millions of players on a global scale and achieved huge commercial success. Among genres of online games, scalability is a particularly critical issue for Massively Multiplayer Online Role Playing Games (MMORPG), in which thousands of players need to be supported concurrently in a large virtual world. In this respect, scalability refers to the number and geographic span of players that can be supported.

Currently, online MMORPGs are predominantly deployed using a client server model, in which a central server has the authority to execute the majority of the game logic and then send a list of objects to the clients for rendering (state update). Client server model achieves scalability by employing server clusters. While having advantages such as the ease of central management and application of security measures, this model is subject to potential single point of failure and the resources have to be over-provisioned to handle peak loads. It should also be noted that in a co-located server cluster, while the computations associated with the simulation of a virtual world can be distributed among the servers in the cluster, the possible bandwidth congestion at the server side cannot be removed.

There is a significant body of work related to a class of scalable designs to support MMORPGs. A detailed survey regarding these designs is conducted in [4] and references therein. While differing in many aspects, the fundamental idea of those designs is to advocate the use of the resources of client computers (peers) including CPU cycles, bandwidth and storage. Ideally, a peer-to-peer (P2P) system can dynamically scale up and down with the number of players, and hence address the inherent scalability issue of the client server model.

Regardless of the supporting models, the general approach for games to achieve scalability is to partition the virtual world into smaller components called zones. In a central server model, the simulation of those zones can be distributed among the servers in a server cluster, whereas in a P2P scenario, these zones can be assigned to the peers who are willing to take up the role of zone servers.

If the zones of a game virtual map are segregated from each other, then there would be no communication between zone servers. In this case, certain techniques are required to give the player the impression that the virtual world is contiguous. For instance, when a player moves from one zone to another, he or she has to walk through a path or flying across a valley to be disconnected from one zone server and connected to another. While these techniques can obviously ease the management of inter-server communications, they also impose constraints on game design, particularly on the scalability of the zones. For instance, a recent study of the player population in a game virtual world [5] has found that a small number of zones could have a high player population of up to 100 players. These are mainly places in the virtual world, such as market places, that encourages player congregation. Given the approach of using segregated zones for the game world, the states of each zone are usually managed by a single server. Hence, the scalability of the zone is largely constrained by the capacity of the server. As a result, when a zone is highly populated, further connection requests may be rejected, which would negatively affect player experiences.

On the other hand, to create a seamless map with transparent boundary between zones, servers of neighbouring zones need to exchange state information of players to enable the cross-boundary interactions between players. We envisage that a map of seamless design would have significant potential to enhance game features for better playability. For instance, a widespread battlefield that engages a large number of players or a social congregation of hundreds of participants, could adopt a seamless map that gives the players the impression of a contiguous virtual space without boundary limitations. It should be noted while the scalability of a zone can be enhanced by joining neighbouring zones into a seamless map, the cost is inter-zone communication between neighbouring zone servers, which is the focus of this paper.

While the cost of inter-zone communication might be of less concern for servers in a central server cluster since the traffic is mainly carried by the Local Area Network (LAN), in

978-1-4244-2413-9/08/$25.00 ©2008 IEEE

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

a P2P scenario, the communication cost could be prohibitively high in supporting a seamless map if neighbouring zones in the game world are assigned to peers with limited bandwidth and far apart from each other in terms of latency. Hence, it is of paramount importance to find an effective approach to assign zones to peers to reduce the inter-zone communication cost. Hereafter we refer to the problem of finding such an approach as the '*zone assignment problem'*.

Our previous works have been mainly on the design of P2P system to support MMORPGs. To that end, in [6] we have developed a game support scenario taking into considerations the traffic trace of a MMORPG. We proposed the use of application-level multicast tree to carry the traffic from zone servers to players, the former would be the root of the multicast tree and the latter the leaves or relay points. We demonstrated that, through the distribution of bandwidth of the multicast tree built on top of a structured P2P overlay (e.g., Pastry), a peer with a typical bandwidth capacity (e.g., ADSL) could potentially support even a crowded zone, while a unicast model would lead to excessive capacity requirements for a typical peer. However, the multicast tree may create extra latency between the zone servers and the players leading to degraded performance for interactions between players. As a result, we proposed in [7] to use a two-level hierarchical overlay: the higher level among nodes located close to network access points, the lower level among nodes in a local domain. We demonstrated that, through fairly low-cost techniques we developed, while bandwidth incurred by the game traffic can be distributed by the multicast tree built on the overlay, the latency performance has improved significantly to allow real-time interactions between players.

In both [6] and [7], we have assumed a game world partitioned into segregated zones. In this paper, we focus on the support of a seamless map and consider inter-zone communication cost that leads to the zone assignment problem. It is a continued effort in building a framework for P2P system to support MMORPGs.

To that end, we propose a heuristic to partition both the physical network and the virtual world (game map) in controlling the communication cost in P2P system for supporting MMORPGs. In partitioning the physical network, we are actually trying to group peers that are topologically close to each other and select them to be potential zone servers. We then demonstrate that the zone assignment problem is NP-hard and incurs a high computational cost even restricting the assignment of zones to a small set of peers selected through the physical network partition. As a result, the proposed heuristic aggregates the zones in a virtual map into larger partitions. In partitioning both the physical network and virtual game map, we are able to assign a group of neighbouring zones to a group of topologically close peers. Through the heuristic, we demonstrate that we are able to achieve a near optimal inter-zone communication cost when deploying a seamless map in a P2P system.

The rest of the paper is organized as follows. In section II, we first discuss the rationale and approach for physical network partition. Then through formulation, we demonstrate the NP-hardness of the zone assignment problem. We hence propose to further partition the virtual world. We finally present the dual-partition heuristic. In section III, we present simulation procedures and results. We finally discuss related work in section IV and conclude in section V.

II. Dual-partition Heuristic

In this paper, we propose a dual-partition heuristic that partitions both the physical network and the virtual map for supporting a seamless map in P2P system. In partitioning the physical world, we are actually exploring the physical network proximity in order to select potential peers as candidate zone servers with coarse granularity. In partitioning the virtual world, we aim to optimize the inter-zone communication cost by exploring proximity of virtual zones in a seamless map.

A. Physical Network Partition

Clustering technique is a widely used technique in the research literature [8-10]. The fundamental idea is to find nodes that are close to each other, and provide this locality information to applications to achieve better performance. For instance, Ratnasamy et al. [10] proposed to use the landmark scheme to have nodes clustered into bins. In this scheme, the landmarks are a number of well known nodes in the Internet. A peer in the overlay can ping those landmarks to obtain its topological coordinate. Peers with the same coordinates are then clustered into the same bin. Simulation through a Transit-Stub [11] topology has shown that the RTT between nodes within the same bin is on average 4.06 times shorter than the RTT between nodes across bins.

In this paper, we rely on available techniques for clustering nodes into bins. However, we note that game distributed design might have specific needs for clustering techniques. For instance, while the landmark binning scheme [10] can effectively group nodes that are close to each other, the number of bins can not be easily controlled. As a result, when the number of landmarks increases, clustering heuristic may generate too many bins with each bin having only a few nodes in it. We will leave it to future work to investigate game specific needs for clustering heuristics.

After clustering nodes into bins, instead of assigning game zones to actual hosts, we assign zones to bins. In a large-scale P2P system, the number of available peers who are willing to take the roles of zone servers could be high. As demonstrated in later sections, the zone assignment problem incurs high computation cost. By assigning zones to bins instead of actual hosts, we are able to significantly reduce the search space. For instance, instead of trying to find an optimal assignment scheme among hundreds of available peers, we cluster those peers into a small number of bins. Effective clustering ensures that assigning zones to different hosts in a common bin would incur similar inter-zone communication cost.

Furthermore, the issue of churn [12] is a critical concern for any P2P system design. Churn refers to the phenomena that in a P2P system, peers or client computers come online and go offline unpredictably. Hence, any optimization scheme or heuristic for zone assignment problem that seek reduced inter-zone communication cost by finding the optimal assignment of game zones to peers would be impractical. This is particularly true for MMORPG given the high churn rate of peers [13] in online games. However, by assigning zones to bins instead of particular hosts, if churn occurs, for

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

instance, a zone server goes offline, another backup peer in the same bin can take up the role of the zone server without affecting much the inter-zone communication cost that might have impact on the game quality for cross-boundary interactions between players.

B. Formulation

After clustering nodes into bins, we formulate the zone assignment problem as follows.

The general known variables are:

d_{st} : The delay between bin s and bin t.

m_{ij} : The number of avatars in the boundary of zone i and zone j.

c_{ij}^{st} : The delay cost of state information exchange between zone i and zone j if assigning zone i and zone j to bin s and bin t respectively.

$$c_{ij}^{st} = m_{ij} \text{ x } d_{st} \qquad (1)$$

N: The number of zones in a seamless map.
M: The number of physical bins.

b_s : The number of peers who are willing to take the role of zone servers in bin s, 1≤ s ≤ M.

The decision variable is

$$x_i^s = \begin{cases} 1 \text{ if zone i is assigned to bin s} \\ 0 \text{ otherwise} \end{cases} \qquad (2)$$

The objective function is

$$\text{Minimize} \sum_{\forall i,j} \sum_{\forall s,t} c_{ij}^{st} x_i^s x_j^t \qquad i \neq j \qquad (3)$$

subject to

$$\sum_{s=1}^{N} x_i^s = 1 \quad \forall \text{ i: } \quad 1 \le i \le N \qquad (4)$$

$$\sum_{i=1}^{M} x_i^s \le b_s \quad \forall \text{s: } \quad 1 \le s \le M \qquad (5)$$

To transform (3) into a linear programming problem, we introduce a binary variable y_{ij}^{st} such that

$$y_{ij}^{st} = \begin{cases} 1 \text{ if } x_i^s = x_j^t = 1 \\ 0 \text{ otherwise} \end{cases} \qquad (6)$$

The objective function (3) then becomes

$$\text{Minimize } \left(\sum_{\forall i,j} \sum_{\forall s,t} c_{ij}^{st} y_{ij}^{st} \right) \qquad i \neq j \qquad (7)$$

subject to

$$\sum_{s=1}^{N} x_i^s = 1 \quad \forall \text{ i: } \quad 1 \le i \le N \qquad (8)$$

$$\sum_{i=1}^{M} x_i^s \le b_s \quad \forall \text{s: } \quad 1 \le s \le M \qquad (9)$$

$$y_{ij}^{st} \le x_i^s \qquad (10)$$

$$y_{ij}^{st} \le x_j^t \qquad (11)$$

$$x_i^s + x_j^t \le 1 + y_{ij}^{st} \qquad (12)$$

The constraints in (8) ensure that each zone in the map can only be assigned to one bin. The constraints in equation (9) ensure that the number of zones assigned to a bin is less than or equal to the number of peers in that bin who are willing to take up the role of the servers for the zones. The constraints in (10) (11) and (12) are to make sure that equation (6) is satisfied.

C. NP-Hardness of the Zone Assignment Problem

Given the objective function in (3), we assume t is a constant T. We further assume $x_j^t = x_j^T = 1$. In other words, for each pair of neighboring zones, one has to be assigned to a particular bin T. An example is given in Figure 1 to demonstrate the feasibility of such an assignment scheme.

Then (3) becomes

$$\text{Minimize} \sum_{\forall i,j} \sum_{\forall s} c_{ij}^{sT} x_i^s \qquad (13)$$

subject to (4) and (5).

This model has been proved in [14] to be equivalent to the well-known Bin Packing Problem[14], which is NP-hard. Hence, the Zone Assignment Problem is NP-hard.

N	T	N	T
T	N	T	N
N	T	N	T
T	N	T	N

Figure 1: T – zones assigned to bin T, N-zones assigned to any bin but T.

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

D. Virtual World Partition

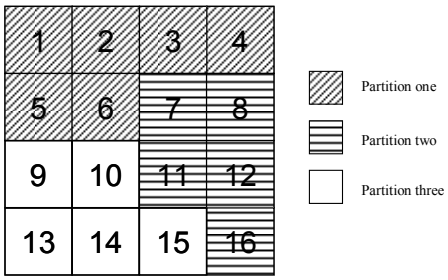

Figure 2: Partitioned Map

The computational cost could be prohibitively high for zone assignment problem, even if we choose to assign zones to a limited number of bins using the model formulated above.

For an equally divided square map as shown in Figure 2, each zone has at least two neighboring zones. Let the number of zones be N, there are at least 2N pair of neighboring zones. Let M be the number of network bins partitioned, there will be M(M-1) possibilities to assign each pair of neighboring zones to the bins. Hence, the number of y_{ij}^{st} decision variables is at least NM(M-1) or $O(NM^2)$. The number of constraints is N + M + 2NM(M-1) or $O(NM^2)$. Given the complexity of the model, as demonstrated in section IV, the formulation can only solve the zone assignment problem in real time for a seamless map having a limited number of zones.

As a result, in addition to physical network partition as discussed previously, we propose to judiciously aggregate the zones in a virtual world map into large partitions to reduce the cost of the zone assignment problem. While a seamless map has multiple zones, neighboring zones can be grouped together and regarded as a single partition to be assigned to a bin. For instance, if a seamless map has sixteen zones as demonstrated in Figure 2, neighboring zones 1 – 6 could be grouped together to form a partition and assigned to peers in a common bin.

The zone grouping policy is a design issue that should consider performance requirement of a seamless map. Zones grouped in one partition, when assigned to peers in a common bin, would have better capacity for cross-boundary interactions between players since the inter-zone communication cost incurred between peers in a common bin would be low. In other words, as an example in Figure 2, the map design for zones in partition one could allow more cross boundary interactions. However, since zones in partition two are assigned to another bin, interactions across zone 3 and 7, for instance, should be kept low since the inter-zone communication cost could be high. As an example, the map for partition one could be designed to handle crowded spaces like cities, battlefield and markets. On the other hand, for boundary area between zones in different partitions, such as zone 3 and 7, the map should be designed as walls or paths to avoid excessive inter-zone communication cost.

E. Selection of Bins

In selecting candidate bins for assigning the zones, we consider the delay cost incurred by the latency between player hosts to the zone servers. To that end, we sort bins based on the sum of delay cost S_t. Let P_s be the number of peers in bin *s* and d_{st} be the delay between bin *s* and *t*. S_t is defined as

$$S_t = \sum_{\forall s} P_s d_{st} \tag{14}$$

Hence, S_t represents the sum of delay cost incurred by the latency from player hosts to zone servers if all zones are assigned to bin *t*.

Based on discussions so far, we present the dual-partition heuristic in Figure 3. In the following sections, we will compare the results obtained from running the optimization model and the heuristic. When presenting the results of the heuristic, we will discuss the heuristic in more detail.

1. Partition the physical network into M bins.
2. Sort the physical network bins based on S_t such that
   $s_1 < s_2 < \cdots < s_M$
3. Let b_i be the number of available zone servers in bin i.
4. Let N be the number of zones in a map.
5. Select the first n bins such that
   $\sum_{i=1}^{n} b_i \geq N$
6. If $\sum_{i=1}^{n} b_i > N$, reduce b_n such that $\sum_{i=1}^{n} b_i = N$
7. Sort the selected bins such that
   $b_1 \geq b_2 \geq \cdots \geq b_n$
8. for i = 1 to n
9. if feasible
10. group b_i neighboring zones along the edge of the map
11. else
12. group b_i neighboring zone
13. end

Figure 3: Dual-partition Heuristic

III. Simulation Procedures and Results

A. Topology Setup

We used Georgia Tech. GT-ITM to generate a transit-stub (TS) topology for experimental purposes. The topology have nodes grouped into either transit or stub domain. In our topology, there are three transit domains with on average eight nodes (transit nodes) in each domain. Each transit node is connected to three lower level stub domains on average. There are on average eight nodes (stub nodes) in each stub domain. TS topology emphasizes the locality feature of the Internet. To

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

this end, link weight is assigned such that traffic path between two nodes in the same domain will remain within that domain [11]. In [15], a similar topology is used to model the Internet. We then have 15 hosts on average connected to each node in the topology to model an overlay network of 10000 hosts. We use the topology assigned link weight as RTT for the links. The link between an end-host and a topology node is assigned a random RTT value between 3 – 10 ms.

We then take advantage of the locality properties of the TS topology by clustering nodes based on the domain they belong to. Thus, we are able to have 24 clusters of nodes because there are 24 stub domains in our topology model. We consider each clustered stub domains as a candidate bin for the zone assignment problem. Given our topology, we have on average 420 end-hosts (peers) in each bin.

B. Results of Solving the Optimization Model

To instantiate the optimization model presented in section II.B, we randomly select a small number of peers who are willing to take up the role of zone servers. In table 1, we list the number of available zone servers in the first five bins sorted based on (14).

Bin Index	1	2	3	4	5
Number of available zone servers	3	7	4	5	12

Table1: Available zone servers in bins

Note that in table 1, the bins are sorted based on the value of S_t as defined (14) such that, for instance, the model should try to assign zones to bin one first, since the sum of delay cost incurred by assigning zones to bin one would be minimal.

2	2	3	3
2	2	1	1
2	2	2	1

(a) 12 zones

3	3	2	2
3	3	2	2
4	1	2	2
4	1	1	2

(b) 16 zones

5	2	1	3	3
2	2	1	3	3
2	2	1	4	4
2	2	4	4	4

(c) 20 zones

2	2	2	4	4
2	2	1	4	4
2	2	1	1	4
5	5	5	3	3
5	5	5	3	3

(d) 25 zones

Figure 4: Results of Optimization Model

We choose Cplex optimization software to optimize the model in section B, III on a Pentium IV 2Ghz Linux server with 2 GB of RAM. We experimented with a square map and gradually increased the number of zones from 12 to 25. The results are presented in Figure 4. The time taken in solving the optimization model increases as the number of zones increases, and is listed in table 2.

Number of zones	12	16	20	25
Time taken to run the optimization model (seconds)	0.22	22	1386	34709

Table 2: Time taken to run the optimization model

In Figure 4, each partitioned square represents a zone in the map. The number in the square represents the bin number that a corresponding zone is assigned to. For instance, in (1) of Figure 4, three zones are assigned to bin one, seven zones are assigned to bin two and the remaining two zones are assigned to bin three.

From Figure 4, it can be first observed that, zones are assigned to bins in a sorted order as expected. For instance, there are three peers available to be zone servers in bin one, and those peers are populated first before bin 2, 3, 4 and 5 are continually to be populated. This is because when we instantiate the optimization model, we assign bins in a sorted order as discussed in section II.E. In doing so, it is ensured that the sum of delay cost between players and zone servers is minimal.

Secondly, the heuristic groups neighboring zones in the virtual map and assigns them to the common bin to achieve minimal inter-zone communication cost. It can be observed in the figure that, zones assigned to a common bin are placed together, with best efforts, in a continued fashion without being separated by other zones. This is because the inter-zone communication cost is incurred by interactions across shared boundaries of neighboring zones. Grouping neighboring zones together and assigning them to the common bin would be beneficial in reducing the inter-zone communication cost since the inter-zone traffic is constrained within the peers in a common bin, in which peers are close to each other in terms of latency.

Finally, the optimization model is not scalable in terms of computation cost. For instance, as shown in table 2, when the number of bins is above 16, the computation cost will become prohibitively high in solving the model in real time.

C. Results of running the heuristic

3	3	2	2	2
3	3	2	2	2
4	1	1	1	2
4	4	5	5	5
4	4	5	5	5

Figure 5: 25-zone aggregated map

Enlightened by the results in Figure 4, our general intuition for the heuristic is that if we partition the virtual map and cluster the neighboring zones, and then assign the entire partition to a common bin, the inter-zone communication cost should be comparable to the results achieved by the

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

optimization model.

We experimented with a 25 zone map. As shown in Figure 5, we applied the heuristic in Firugre 3 to partition the map by clustering neighboring zones based on available zone servers in each bin (table 1).

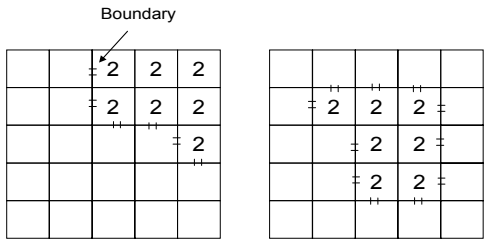

Figure 6: Boundary Comparison

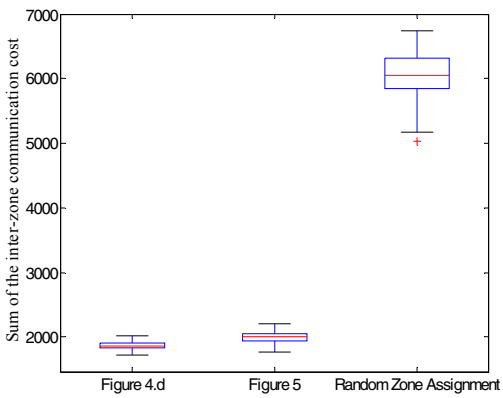

Figure 7: Comparison of the Inter-zone Communication Cost

For instance, in bin one, there are three available peers who are willing to take up the roles as zone servers, we hence cluster three zones to be assigned to bin one as a single partition. When we cluster zones in the virtual map, we try to place the big partition, for instance, partition 2, along the edge of the map to reduce the number of shared boundaries between different partitions (step 6-13 of the heuristic). An example is given in Figure 6.

If we have partition two (all zones in partition two are assigned to peers in bin two) placed along the edge of the map, we will have six boundaries shared with other partitions as shown on the left part of Figure 6. However, if we have the partition placed in the middle of the map, we would have eleven boundaries as demonstrated in the right part of the figure. Since inter-zone communication is incurred by cross-boundary traffic, we would expect less communication cost if we had less shared boundaries when we assign the partition along the edge of the map.

To compare the performance in terms of inter-zone communication cost between Figure 4.d, which is the result of solving the optimization model, and Figure 5, which is achieved by running the heuristic, we assign zones of the map to randomly selected peers in the corresponding bin. For instance, based on Figure 5, we assign the four zones on the top left corner of the map to four selected peers in bin three, whereas based on Figure 4.d, we assign the four zones on the down right corner to four selected peers in bin three. It should be noted that when assigning zones to peers in the same bin, we do not further seek optimization goals beuase of the issue of churn in P2P system as we discussed in section II.A.

To normalize the simulation results, we set $m_{ij} = 1$ (section II.B). Again, we use the network topology as discussed in section III.A. We run the simulation for 10000 times and the results are demonstrated in Figure 7.

Box plots in Figure 7 shows that the sum of the inter-zone communication cost based on Figure 5 is only slightly higher than that of Figure 4.d, which is obtained by solving the optimization model. The mean of the sum of the communication cost of the optimized model (Figure 4.d) is 1847, whereas the mean of the sum of the communication cost of the partitioned map (Figure 5) is 1915 and only 5 percent higher than that of the optimized model.

Given our particular topology, $m_{ij} = 1$ (one player at each boundary of each zone) and a 25 zone map has 45 shared boundaries altogether, the player at the boundary of a zone would be subject to an extra 41ms delay cost on average due the requirement of the inter-zone communication for the optimized model (1847/45 = 41), whereas for the partitioned map, the player would be subject to 42.6 ms delay cost (1915/45 = 42.6). The minor difference of 1.6 ms in delay cost is negligible as far as the game quality of MMORPG is concerned.

For comparison purpose, we finally assign zones of a 25 zone map to peers in the topology randomly and calculate the sum of the inter-zone communication cost. As shown in Figure 7, both the optimized model and the partitioned map can save the communication cost by a factor of three in comparison with a map for which zones are randomly assigned to peers.

To conclude for this section, the computation cost for zone assignment problem is prohibitively high (table 2) when solving the optimization model. We are able to achieve nearly the same level of optimization goal in reducing the inter-zone communication cost by partitioning the zones in the virtual map and assigning the partitions to physical network bins.

D. *Map with Hotspots*

In the previous simulation, we have set $m_{ij} = 1$ across all boundaries. In other words, players are uniformly distributed across the boundaries of the map. However, in MMORGP, player flocking may occur due to the features of the game design. For instance, an interesting occurrence of an event may attract more players to come to a certain part of the game world and create a hotspot in the map. To experiment with a map with hotspots at the boundaries, we selected a number of boundaries of a 20-zone map and set $m_{ij} = 2$ at those boundaries to simulate hotspots at which the interactions between players are twice as much as that of a normal boundary. In Figure 8.a, those hotspots are marked with crosses.

We then run the optimization model for the zone assignment problem and the results are presented in Figure 8.b. It can be observed from the figure that, in addition to what we

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

have discussed in section III.B, the optimization model restricts the hotspots within common physical network bins. In other words, the optimization model assigns the two neighboring zones connecting a hotspot to a common bin. In doing so, the communication cost incurred by the hotspot would be kept low since the two neighboring zones are assigned to peers in a common bin, and hence are close to each other in terms of latency.

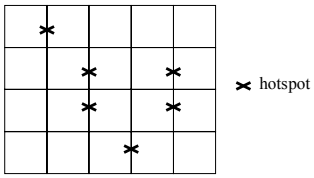

(a) Map with hotspots

3	3	2	2	2
3	2	2	2	2
3	1	1	4	4
5	1	4	4	4

(b) Results of Optimization Model

3	3	1	2	2
3	1	1	2	2
3	4	4	2	2
5	4	4	4	2

(c) Partitioned Map

Figure 8: A map with hotspots at boundaries

Enlightened by the results we discussed above, we again apply the virtual map partition following the heuristic in Figure 3. In addition, while assigning the big partitions along the edge of the map, we restrict the neighboring zones connecting a hotspot to be assigned to peers within a common physical network bin. The partitioned map is presented in Figure 8.c.

The total communication costs of both Figure 8.b and Figure 8.c is compared in Figure 9, as the results of similar procedures to the previous section.

In Figure 9, the mean of the sum of the communication cost of the optimization model (Figure 8.b) is 1699, whereas the mean of the sum of the communication cost as a result of running the heuristic (Figure 8.c) is 1739, less than 3 percent higher than that of the optimized model.

Given that a 20-zone map has 28 shared boundaries in total, the player at the boundary of a zone would be subject to an extra 60.7ms delay cost on average due the requirement of the inter-zone communication for the optimized model (1699/28 = 60.7), whereas for the partitioned map, the player would be subject to 62.1 ms delay cost (1739/28 = 62.1). The minor difference of 1.4 ms in delay cost is negligible as far as the game quality of MMORPG is concerned.

To conclude this section, if games feature hotspots at certain area of the map that could lead to increased inter-zone communication costs at the boundaries, the map that is partitioned based on the location of the hotspots could achieve nearly the same level of communication costs as that can be achieved by the optimization model. We again demonstrated the effectiveness of partitioning both the physical network and virtual map in reducing inter-zone communication costs, while avoiding the high computation costs required by the optimization model.

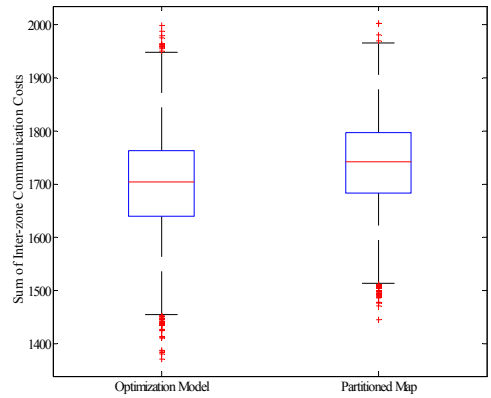

Figure 9: Comparison of the Communication Costs

IV. Related Work

To our best knowledge, there is no existing work that directly addresses seamless maps design in P2P system. Distributed server architecture is usually considered as an alternative to central server model in addressing the scalability and performance bottleneck of the latter. In [14] and [16], heuristics are developed for distributed server system to improve performance. The main focus of those heuristics is to assign players to servers so that the communication cost between the players and the servers (e.g., latency) is low, whereas our focus is to assign zones to servers so that the communication costs between servers is low.

In [16], inter-server communication cost is addressed as part of a design of partitioning scheme for distributed server system. However, the heuristic developed is subject only to a small number of servers (e.g. 8 servers). Given the partitioning of physical network and virtual map, our approach is not limited by the number of servers. This is important since in P2P system, a virtual map might need to be partitioned into much smaller components to be assigned to peers with smaller capacity comparing to servers in a distributed server system. As a result, more peers are required for the zone assignment problem.

It should also be noted that distributed server system is relatively static, whereas P2P system features dynamics of peer participation or churn. As a result, we choose to assign zones to bins instead of to particular hosts.

V. Conclusions

Massively Multiplayer Online Role Playing Games usually feature a huge virtual world for players to explore. To achieve deployment scalability, the game world can be partitioned into smaller components called zones. In a peer-to-peer system, those zones can be assigned to peers who are willing to take up the roles of zone servers to achieve scalability. A seamless map requires inter-zone communications to allow interactions across zone boundaries between players. A seamless map improves game playability by making the boundaries of zones transparent to players. To enable a seamless map in P2P system, it is critical to find effective approaches such that when assigning zones to peers, the inter-zone communications cost can be kept minimum. In

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.

this paper, we formulated an optimization model for the zone assignment problem and demonstrated the NP-hardness of the problem. We then proposed a low-cost heuristic to partition the physical network into bins such that peers in a common bin are topologically close to each other. Furthermore, enlightened by the results of the optimization model, the heuristic partitions the virtual world map such that neighboring zones in the map are clustered together and assigned to a bin as a single partition. In doing so, we demonstrate that we are able to achieve the same level of inter-zone communication costs as those can be achieved by the optimization model we formulated, while avoiding the prohibitively high computation cost required by the latter.

Acknowledgement

This work is supported by the Co-operative Research Centre for Smart Internet CRC (SITCRC) and the University of Wollongong (UOW), Australia.

References

[1] Blizzard. www.blizzard.com.
[2] Counter-Strike. www.counter-strike.com.
[3] Everquest. http://everquest.station.sony.com.
[4] X. Jiang, F. Safaei, and P. Boustead, "Latency and Scalability: A Survey of Issues and Techniques for Supporting Networked Games," presented at International Conference on Networks (ICON 2005), Kuala Lumpur, Malaysia, 2005.
[5] D. Pittman and C. GauthierDickey, "A measurement study of virtual populations in massively multiplayer online games," presented at the 6th ACM SIGCOMM workshop on Network and system support for games, Melbourne, Australia, 2007.
[6] X. Jiang, F. Safaei, and P. Boustead, "An approach to achieve scalability through peer-to-peer network for massively multiplayer role playing games," *ELSEVIER Journal on Computer Communications*, vol. 30, pp. 3075-3084, 2007.
[7] X. Jiang, F. Safaei, and P. Boustead, "Enhancing the multicast performance of structured P2P overlay in supporting Massively Multiplayer Online Games," presented at IEEE International Conference on Networks 2007 (ICON 2007), Adelaide, Australia, 2007.
[8] X. Y. Zhang, Q. Zhang, Z. Zhang, G. Song, and W. Zhu, "A Construction of Locality-Aware Overlay Network: mOverlay and Its Performance," *IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS*, vol. 22, pp. 18-28, 2004.
[9] R. A. Ferreira, S. Jagannathan, and A. Grama, "Enhancing Locality in Structured Peer-to-Peer Networks," presented at Tenth International Conference on Parallel and Distributed Systems (ICPADS?4), 2004.
[10] S. Ratnasamy, M. Handley, R. Karp, and S. Shenker, "Topologically-Aware overlay construction and server selection," presented at INFOCOM, 2002.
[11] E. Zegura, K. Calvert, and S. Bhattacharjee., "How to Model an Internet," presented at IEEE INFOCOM, 1996.
[12] , D. Stutzbach, and R. Rejaie, "Understanding churn in peer-to-peer networks," presented at Proceedings of the 6th ACM SIGCOMM conference on Internet measurement, Rio de Janeriro, Brazil, 2006.
[13] J. Kim, D. Chang, and T. Kwon, "Traffic Characteristics of a Massively Multi-player Online Role Playing Game," presented at Netgames'05, New York, USA, 2005.
[14] D. N. B. Ta and S. Zhou, "A two-phase approach to interactivity enhancement for large-scale distributed virtual environments," *Computer Networks,* vol. 51, pp. 4131-4152, 2007.
[15] C. D. Nguyen, F. Safaei, and P. Boustead, "A distributed server architecture for providing immersive audio communication to massively multi-player online games," presented at IEEE International Conference on Networks (ICON 2004), Singapore, 2004.
[16] J. C. S. Lui and M. F. Chan, "An Efficient Partitioning Algorithm for Distributed Virtual Environment Systems," *IEEE Transactions on Parallel and Distributed Systems,* vol. 13, pp. 193-211, 2002.

Authorized licensed use limited to: University of Wollongong. Downloaded on November 4, 2008 at 22:24 from IEEE Xplore. Restrictions apply.