Leveraging Fog to Extend Cloud Gaming for Thin-Client MMOG with High Quality of Experience

Yuhua Lin and Haiying Shen
Department of Electrical and Computer Engineering
Clemson University, Clemson, South Carolina 29634
Email: {yuhual, shenh}@clemson.edu

***Abstract*—With the increasing popularity of Massively Multiplayer Online Game (MMOG) and fast growth of mobile gaming, cloud gaming exhibits great promises over the conventional MMOG gaming model as it frees players from the requirement of hardware and game installation on their local computers. However, as the graphics rendering is offloaded to the cloud, the data transmission between the end-users and the cloud significantly increases the response latency and limits the user coverage, thus preventing cloud gaming to achieve high user Quality of Experience (QoE). To solve this problem, previous research suggested deploying more datacenters, but it comes at a prohibitive cost. We propose a lightweight system called *CloudFog*, which incorporates "fog" consisting of supernodes that are responsible for rendering game videos and streaming them to their nearby players. Fog enables the cloud to be only responsible for the intensive game state computation and sending update information to supernodes, which significantly reduce the traffic hence the latency and bandwidth consumption. Experimental results from PeerSim and PlanetLab show the effectiveness and efficiency of CloudFog in increasing user coverage, reducing response latency and bandwidth consumption.**

Keywords*—*Cloud gaming; P2P network; Online gaming; Quality of experience

I. INTRODUCTION

Mobile gaming is seeing the fastest growth [1] in the gaming industry, and cloud gaming, as a flourishing gaming model, is a solution for such thin-client Massively Multiplayer Online Game (MMOG). Cloud gaming frees players from hardware requirement and game installation on their local computers. In cloud gaming, games are stored and run on remote servers, and game videos are streamed to end-users through broadband Internet connections. Cloud gaming also saves the cost of game service providers. They can buy cloud resource based on the actual demands in the large-scale system. Also, game service providers do not have to develop multiple versions of the same game to meet different operating systems (e.g., Linux, Windows, Mac), and spend money on software piracy protection.

The advantages of cloud gaming makes it a very promising model to cater to the dramatically rapid growth of MMOG and online mobile gaming considering their very large user scale and thin clients. Though the advantages of cloud gaming makes it a very promising model to cater to thin-client MMOG, it currently faces severe challenges (i.e., latency, network connection, user coverage and bandwidth cost) that prevent it from becoming a leading gaming model. First, response latency is a critical factor in user quality of experience (QoE). By offloading computation to a remote host, cloud gaming suffers from long *response latency*; the delay in sending the user action information and game video between the end-user and the cloud. Second, cloud gaming services post a strict requirements of high-speed network connection for a relatively high constant downlink bandwidth (e.g., 5Mbit/s recommended by OnLive). Third, the shortage of datacenters limits user coverage. Players begin to notice a response delay of 100ms [2]; 20ms attributed to playout delay on client side and processing delay on the cloud, 80ms attributed to the network latency. The playout delay of a client includes the time to send action information, receive and play the game video. Choy *et al.* [3] found that Amazon's EC2 (with 13 datacenters) can provide a median latency of 80ms or less to only fewer than 70% of their 2500 tested end-users in the US. They also found that substantial increase in the total number of datacenters is required to significantly increase user coverage. Existing cloud infrastructure is not sufficient for hosting cloud gaming, as a sizeable portion of the population would experience significantly degraded QoE. Fourth, besides server time, bandwidth costs represent a major expense when renting on-demand resources. Considering the MMOG's huge user scale, these costs can significantly affect the feasibility of thin-client MMOG [4] on the cloud.

II. LEVERAGING FOG TO EXTEND CLOUD GAMING

The great promises of cloud gaming and the obstacles it faces motivate us to explore approaches to efficiently handle the challenges. Though previous study suggested deploying more datacenters [2], building and maintaining a large number of datacenters is cost-prohibitive. In this paper, we propose a lightweight system called CloudFog. We introduce a concept called "fog", formed by powerful supernodes, that are close to end-users and connected to the cloud. In CloudFog, the intensive computation [5] of the new game state of the virtual world is conducted in the cloud. The cloud sends update messages to supernodes, the supernodes update the virtual world, render game videos for different players and stream videos to the players. Thus, users without high speed network connection to cloud or out of the coverage of the cloud can be supported by nearby supernodes, and the cloud does not need to transmit entire game videos to far-away users. This strategy can increase user coverage, shorten response latency, ensure relatively high-speed network connection for high QoE and reduce bandwidth cost.

Previous studies [2], [3] revealed that the uploading from the players to the cloud does not seriously affect the response latency, and downstream latency is an important factor for QoE [2], which is affected by the game video streaming rate. Thus, we aim to reduce the downstream latency by

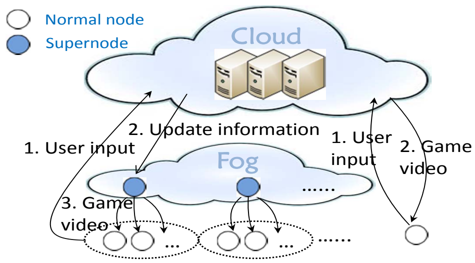

Fig. 1: Fog-assisted cloud gaming infrastructure.

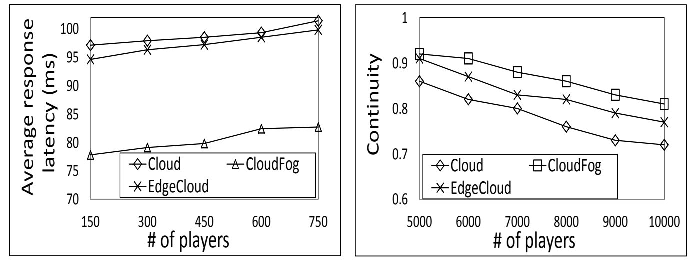

(a) Response latency. (b) Playback continuity.

Fig. 2: Performance Evaluation in PlanetLab

reducing the traffic transmitted from the cloud. In our design, game videos are streamed from nearby supernodes to players, instead of from remote game servers. As the computation of a virtual world for MMOG has a very high demand on server capacities [5], cloud is responsible for this task. Figure 1 shows our fog-assisted cloud gaming infrastructure. The fog is formed by supernodes, and normal nodes are connecting to their nearby supernodes. The normal nodes that cannot find nearby supernodes directly connect to the cloud.

We use n_i to denote a normal node, and sn_j to denote a supernode in the system. When each supernode is initially deployed, it is pre-installed with the *game client*. During the game playing, when node n_i makes an action (e.g., launching a strike or moving to a new place), this information is sent to the cloud server. The server collects action information from all involved players in the system and performs the computation of the new game state of the virtual world (including the new shape and position of objects and states of avatars). The cloud then sends the update information to the supernode of n_i (sn_j), which updates its virtual world accordingly. sn_j then renders game video for n_i based on $n_i's$ viewing position and angle. sn_j finally encodes the game video and stream it to n_i. As a player is close to its supernode in network distance, and the traffic from the cloud is significantly reduced, so the game video transmission delay is much shorter than that of downloading game video directly from the cloud as in the current cloud computing systems.

III. Preliminary Results

We conducted experiments on the PlanetLab [6] real-world testbed to evaluate the performance of CloudFog in comparison with other systems. *CloudFog* denotes the proposed fog-assisted cloud gaming infrastructure, and we compared *CloudFog* with the current cloud gaming model [7] (denoted by *Cloud*) and *EdgeCloud* [8]. *EdgeCloud* deploys a number of powerful servers to increase user coverage. The difference between *EdgeCloud* and *CloudFog* lies in the responsibility of newly added servers. *EdgeCloud* simply adds powerful servers to takeover all the cloud's tasks (including storing and computing game status and rendering new game videos), while in *CloudFog*, the supernodes only need to receive updates from the cloud to render new game videos and stream them to the players. The default number of main datacenters is 2 for all methods, and *EdgeCloud* has additionally 8 randomly distributed servers. We used 750 distributed nodes nationwide, and 300 of them have the capacity to be supernodes.

Figure 2(a) shows the average response latency per player in different systems. We see that *EdgeCloud* generates slight shorter response latency than *Cloud* due to the use of scattered servers, and users are more likely to connect to servers within a short distance. In *CloudFog*, users are supported by supernodes that are physically close to them. As the game video is streamed from supernodes to the users, instead of from servers that are physically far away. Thus, *CloudFog* is able to reduce the response latency for users. When there are not enough packets in the cache, the player suffers from an playback interruption. We measured continuity by the proportion of packets arrived within the required response latency over all packets in a game video. Figures 2(b) shows the average playback continuity of game videos when different number of players are playing games concurrently. We see that *Cloud* yields the lowest playback continuity because there are only a small number of cloud servers, which may locate far away from some players. So most game videos need to be transmitted from remote servers to clients, thus large portion of packets cannot be received within the required response latency. *EdgeCloud* produces higher continuity than *Cloud* because players in *EdgeCloud* are supported by their nearby servers. *EdgeCloud* generates smaller continuity than *CloudFog*, because not all users in *EdgeCloud* are able to connect to a nearby server due to the shortage of servers. So game video packets need to travel longer distance than that in *CloudFog*.

Acknowledgements

This research was supported in part by U.S. NSF grants NSF-1404981, IIS-1354123, CNS-1254006, CNS-1249603, and Microsoft Research Faculty Fellowship 8300751.

References

[1] N. Bilton. Video Game Industry Continues Major Growth, Gartner Says. *The New York Times*, 2011.

[2] M. Jarschel, D. Schlosser, S. Scheuring, and T. Hobfeld. An Evaluation of QoE in Cloud Gaming Based on Subjective Tests. In *Proc. of IMIS*, 2011.

[3] S. Choy, B. Wong, G. Simon, and C. Rosenberg. The brewing storm in cloud gaming: A measurement study on cloud to end-user latency. In *Proc. of NetGames*, 2012.

[4] E. Carlini, M. Coppola, and L. Ricci. Integration of P2P and Clouds to support Massively Multiuser Virtual Environments. In *Proc. of NetGames*, 2010.

[5] C. Bezerra and C. Geyer. A load balancing scheme for massively multiplayer online games. *Multimedia Tools Appl*, 2009.

[6] PlanetLab. http://www.planet-lab.org/, [Accessed in Nov 2014].

[7] C. Huang, C. Hsu, Y. Chang, and K. Chen. GamingAnywhere: An Open Cloud Gaming System. In *Proc. of MMSys*, 2013.

[8] S. Choy, B. Wong, G. Simon, and C. Rosenberg. EdgeCloud: A New Hybrid Platform for On-Demand Gaming. Technical Report CS-2012-19, University of Waterloo, 2012.