Mobile Cloud Gaming: Issues and Challenges

Omar Soliman[1], Abdelmounaam Rezgui[1], Hamdy Soliman[1], and Najib Manea[2]

[1] Department of Computer Science & Engineering, New Mexico Tech
Socorro, NM, 87801, USA
{oms,hss}@nmt.edu, rezgui@cs.nmt.edu

[2] University of New Mexico, Los Lunas, NM 87031, USA
nmanea@unm.edu

Abstract. Recent developments in mobile, cloud, and graphics processing technologies have enabled *mobile cloud gaming*, a gaming model where players use mobile devices to play graphics-intensive games that run remotely on cloud servers. This delivery paradigm is called *Gaming as a Service* (GaaS). GaaS is used to stream computer games across the Internet. It gives rise to various technical, legal, and ethical issues. In this paper, we present the current state of the art in GaaS along with open issues and research challenges.

Keywords: GPU service, cloud gaming, on-demand gaming.

1 Introduction

Gaming as a Service (GaaS) is a relatively new concept in the field of computer science and cloud computing. It involves providing video games on-demand to consumers through the use of cloud technologies. One benefit of cloud technology is the transfer of computation from a relatively weak user device (or thin client) to more powerful cloud servers. Even when a user's device is appropriately powerful, GaaS can reduce power consumption and provide other cloud services (e.g., storage).

1.1 Background

The effects of Moore's Law [9] have led to successively cheaper and more powerful computers in a cyclic fashion since the mid-20th century. Since the advent of the ARPANET and its evolution, the Internet has grown at a similar rate [21]. As Moore's Law allowed computer access to a greater population, demand for Internet service also grew. Dial-up access evolved into broadband to serve more customers who now have more powerful machines that can run more compute-intensive applications. Meanwhile, technical progress led to smaller and more portable computers, which needed portable Internet. This was answered through the development of wireless Internet access technologies. These mobile devices continued to shrink and have become more powerful contributing in millions of users migrating from traditional computing platforms to "smart" phones and

F. Daniel, G.A. Papadopoulos, P. Thiran (Eds.): MobiWIS 2013, LNCS 8093, pp. 121–128, 2013.

tablets. In 2012, more than 800 million such devices were sold and demand is expected to exceed 1 billion in 2013 [18]. These devices now provide advanced multi-media capability, three dimensional graphics, multiple cores, and large touch-screens [2]. Their capabilities have allowed anytime, anywhere usage of traditional desktop applications, including games. To serve the Internet demand of such devices, standards such as IEEE 802.11ad and LTE-Advanced are being developed and deployed [1,13].

Although mobile smart devices and their Internet capabilities have greatly evolved, their abilities are still outmatched by similar progress in their desktop counter-parts. Meanwhile, the rise of cloud computing has enabled computational resource consolidation in order to provide data and processing outsourcing on a massive scale [4]. The effect is to provide economies of scale in terms of computation and cost to unrelated users. Cloud technology amenities are provided through virtualization services. These services include Infrastructure, Platform, and Software as a Service. They provide computational infrastructure (IaaS), middleware and runtime environments (PaaS), or application and data hosting (SaaS).

Users need terminal devices (thin clients) to access cloud services. The ubiquity of mobile smart devices allows widespread cloud access to a large portion of the population. Overall, this has led to users likely carrying mobile touch-screen computers with persistent cloud connection, functioning as a phone, Web browser, gaming device, etc. According to [14], a significant portion of device applications consists of games. For instance, Apple's AppStore has 18% mobile games. Additionally, gaming users are shifting towards mobile devices and away from traditional consoles. In summary, the groundwork has been laid to provide a large user base for cloud services in order to deliver game content anytime, anywhere, on mobile devices without being limited by the device's capabilities or game requirements.

1.2 Video Gaming

Single-player video games started spreading with personal computer adoption in the 1980s [10]. With the growth of the Internet and computer networking technologies, video games increasingly incorporated multiplayer features. As mobile devices evolved to become more useful, mobile games also came into existence. Recently, the growth of cloud computing has accelerated the rise of GaaS. Video games are an important and growing segment in the entertainment industry. In April 2008, the Grand Theft Auto IV video game earned $310 million in 24 hours globally (more than the highest earning movie at $60 million and book at $220 million combined) [20]. The video game industry is predicted to reach $82 billion in size by 2017 [5], and mobile gaming presently accounts for $9 billion [19]. Even in a depressed economy, the industry grew 16% [16], while the largest game publisher Electronic Arts netted 40% of its revenue from online content [7]. This indicates that not only are video games a significant market force, but also that they are heavily involved in online infrastructure.

The major uses of gaming related cloud technologies include social gaming, massive multiplayer online role playing games (MMORPGs) and auxiliary services. Social games are largely based on social interactions between users, especially in the context of a social-enabling platforms (e.g., Facebook). These games involve communication and large-scale user participation, and may include aspects such as cooperation and competition. They may also include performance tracking as well as virtual incentives (e.g., achievements/trophies). MMORPGs are online-hosted video games with a massive user base. The major example is War of Warcraft, which has peaked at 1 million concurrent users from a potential base of 10 million subscribers [6]. The scale of these games requires major computational power to coordinate user interaction and its results but still relies on user devices to render graphics and provide much computation. Auxiliary cloud services (such as Steam and XBOX Live) provide IaaS level cloud services such as limited storage and networking infrastructure. They are only meant to augment the functionalities of user devices with networking, storage, and server access.

2 Gaming as a Service

GaaS can be condensed as *any-device* gaming. This relates to the ubiquity of graphical screen devices present in a user's vicinity. As discussed earlier, users maintain mobile smart devices which are growing in power and functionality, able to utilize the Internet and execute desktop applications, including video games. The greatest shortcomings of these devices however are graphics processing and battery life. Efficient graphics processing requires dedicated GPUs and puts a heavy drain on battery. However, with the shift to mobile gaming, users demand the quality of dedicated GPU gaming with the convenience of mobility. Pre-cloud approaches included creating lower-quality games or constructing desktop-replacement laptops. GaaS aims to provide a solution for high-quality games on *any* device. Utilizing the cloud to provide on-demand games to users can be thought of as a form of SaaS. From an administrative point of view, this is more akin to Platform as a Service, with the game's application data fulfilling the SaaS designation. In general, PaaS requires a cloud infrastructure and a middleware for executing application code. In the case of GaaS, the specific middleware may be vendor-dependent. However, the infrastructure will likely emphasize GPU-centric capabilities. GaaS generally requires broadband access for speed, compression techniques for bandwidth, and encryption for access control. Immediate benefits to the user include re-use of existing hardware, no game downloads, and no software maintenance.

A general model of GaaS can be described in terms of three components: the client, the interface, and the renderer:

– The client consists of the user and their thin client. The user is a consumer who wishes to play a game on their device (thin client). The thin client allows the user to access the GaaS interface and receive audio/video from the renderer. It also provides the user with I/O capability as well as Internet access.

- The interface is a mechanism that authorizes users and receives commands from them. An example would be an application that runs on the user's thin client. Through the interface, the user can login to the service, select games to play, or send input to be processed. The interface also separately controls the renderer.
- The renderer is the processing center of the GaaS service. It interprets user commands, creates frames of game data, and transmits them to the user.

In summary, GaaS provides GPU-centric SaaS which allows the users of smart devices to demand any game anywhere anytime. A major benefit of GaaS is that it allows client outsourcing of complex graphical calculations from a consumer device to the cloud infrastructure. This reduces the user-side computation and, as a result, the total energy required to execute game-related code on a client device. Further energy savings are realized as additional users share GaaS cloud GPUs where virtualization technology (e.g., NVIDIA VGX Hypervisor[1]) can maximize user density, resulting in more effective utilization of GPU cycles. To illustrate, a single GaaS-capable NVIDIA VGX K2 unit requires 38 watts per cloud user[2] whereas a comparable single-user NVIDIA GTX 690 consumer unit requires 300 watts to operate[3]. In this case, GaaS can reduce the overall graphics power consumption by 87%. In addition to power saving, users gain all the benefits of traditional cloud services and save the time necessary to contend with offline games.

2.1 GaaS Providers

GamingAnywhere is the first open-source GaaS development platform[4]. It allows researchers to test ideas on a GaaS testbed, GaaS providers to develop services on it and users to construct their own personal gaming clouds from desktops. CloudUnion[5] and OnLive[6] are GaaS companies that provide interfaces through downloadable applications and incorporate social aspects into the environment. Gaikai[7] delivers GaaS through Web browsers and needs no application download. G-cluster[8], the pioneer of GaaS, specializes in delivering GaaS over IPTV and mobile phone networks.

While it does not offer GaaS directly, NVIDIA makes the hardware necessary to enable GaaS. It offers a repackaged version of the traditional GPU-centric IaaS to GaaS companies. It has developed virtual GPU technology that allows 36 concurrent gaming streams on a single server. This service reduces the barrier

[1] `http://www.nvidia.com/object/cloud-gaming-benefits.html`
[2] `http://www.nvidia.com/object/grid-boards.html`
[3] `http://www.geforce.com/hardware/desktop-gpus/geforce-gtx-690/specifications`
[4] `http://www.gaminganywhere.org`
[5] `http://www.cloudunion.cn`
[6] `http://www.onlive.com`
[7] `http://www.gaikai.com`
[8] `http://www.g-cluster.com`

to entry for new GaaS companies that would otherwise be incapable of acquiring such specialized hardware.

2.2 GaaS Issues

The main concerns of GaaS include user responsiveness, stream quality, service quality, and operating cost:

- *User responsiveness* relates to the visible lag of the delivered video frames or audio as well as the perceived latency in the response to user commands. This is caused by network latency and processing delay. Network latency is the result of the user's access network. It can be addressed through gaining access to higher quality connections or deploying more cloud locations by the service provider. Audio/Video lag can be reduced through the use of better GPU systems although it may be related to network quality as well. User responsiveness is especially crucial because many video games depend on extremely quick action performances from the user and may be the single greatest barrier to widespread replacement of gaming devices with GaaS. However, this latency requirement may not exist for all game types (e.g., turn-based games).
- *Video quality* is a balance of many factors. It is limited by the graphical resolution of the user's device, network bandwidth, and load on the cloud GPUs. Bandwidth can be addressed through the use of compression algorithms. A recent example is when Netflix began utilizing a codec optimization technology to reduce its video streaming bandwidth by 50% [15]. Optimizations such as this and the development of new standards will allow higher quality video with fewer resources for GaaS service. Video quality is important because it provides direct data of the game's current state to the user and serves as a major selling point for GaaS.
- *Service quality* is the overall quality of the GaaS service. This includes game selection space, storage capacity, processing speed, and interface. Game selection is partly non-technical although it relates to the compatibility of the middleware with games originally developed for non-cloud platforms. Storage involves the saved game data of users as well as their settings and paused game sessions. Processing speed involves the loading/buffering times of playing a cloud game. Interface involves the attractiveness and ease of use of the service as well as its compatibility with the controls of the user's mobile thin client. Service quality is important because it determines the level of participation in a particular GaaS service.
- *Operating cost* is a combination of equipment cost and operational costs. Equipment must be maintained and replaced when outdated. Operating cost (such as power and bandwidth) can be optimized through the selection of the cloud's physical location but is a trade-off between low billing rates and user responsiveness. Operating cost is important because it determines the profitability of a GaaS company and the level of any charges users may face.

2.3 Related Legal Issues

Many legal issues arise from the use of GaaS [11]. As more interest in GaaS becomes generated, patents may be formed specifically for this domain, resulting in restrictions in evolution and growth. Ownership also becomes a concern. Since game code resides and is rendered on the cloud, the concept of owning a game may give way to purchasing access to a game for a set amount of time. Also of importance are guaranteed service levels and pricing schemes. Other issues include the requirement of a large number of users to participate in the service, not only to provide multiplayer potential but also to offset the large costs associated with providing a GaaS service. On a beneficial note, piracy and hacking may be missing in GaaS services since no code is hosted on the user's device; only streamed video of the game is delivered.

3 Current GaaS Research

The field of GaaS is currently the focus of significant research. We illustrate that research through some representative recent work. In [8], the authors present a GaaS system that permits the simultaneous support of multiple devices of varying resolutions from varying focal points. This is accomplished through the modularization of the cloud structure into service (DSP), rendering (DRS), and streaming (EQS) systems. The DSP receives user input and interacts with the DRS to model a scene. The EQS then renders the image and streams to the client device. Multiple views are achieved through iteratively creating cameras and varying viewports and binding them to render windows. Performance is maintained through the use of per-thread task queues, highly optimized schedulers, and shared GPU resources. In addition, multiple views are packed as a continuous stream of traffic using the H.264 codec. Tests achieved 15 fps for 8 simultaneous streams on a single GPU.

In [17], the authors present a framework that provides a virtual interface for remote hardware. In this sense, it is a form of middleware for GaaS. It is composed of three components: a Kusanagi plug-in, lobby server, and MPEG-4 client. The plug-in provides interactivity from the client to server and video from server to client. The lobby server acts as both a portal and resource manager. The client is composed of a full multimedia player with networking and interactivity. Tests indicated that the system was appropriately adjusting video bitrate to respond to network congestion and that 92% of test users felt positively about the system.

In [3], the authors examined the suitability of the current network infrastructure in the United States in providing GaaS. The authors measured network conditions of 2500 users of the Amazon EC2 service using BitTorrent and geographical location filtering. They found that the median latency to less than 70% of them was 80 ms or better which they argue is the maximum permissible latency for an acceptable GaaS experience. In addition, 10% of users would be incapable of properly utilizing GaaS due to extreme latency. The authors then evaluate the best way to reduce the overall latency for the country by comparing latency vs. region based data center creation strategies. They found that a prohibitive number of additional centers are required to achieve 90% population coverage. An additional strategy of deploying GPU-equipped CDN edge servers results in a 28% increase of user coverage. Overall, the authors conclude that the existing US cloud infrastructure is currently not suitable for providing GaaS.

In [12], the authors examined the evolution of the software business model from traditional packaged format to SaaS and found that it has forced developers to transform into low-margin, high user formats. Specifically, they perform a case study on G-cluster, the pioneer of GaaS. They examined G-cluster in 2005 and found that it reached 15,000 households by running game source-code modified for their infrastructure. This service was then offered to middle-men who, in turn, sold it to consumers. By 2010, G-cluster had evolved to provide full-fledged GaaS to network operators of IPTV totaling 3 million households. They also released a SDK for game developers wishing to submit their own compatible games. G-cluster found that the main audience for GaaS was moderate users who were passionate enough to pay for games but casual enough to not want to spend money on their own GPU devices. Overall, the authors recommend that GaaS companies provide a minimum of PaaS level in order to remain competitive.

4 Conclusion

GaaS is an emerging cloud-based computing paradigm that is enabling mobile gaming. It is the result of consumer demand for anytime, anywhere gaming on low-power, Internet equipped mobile devices. GaaS is game-specific SaaS that utilizes GPU centric hardware. The game is processed and rendered in the cloud and streamed to the user's device. This allows the use of low-end, mobile devices to "play" high-end games. It also reduces the cost and time necessary to play games by distributing these over the entire user base. This is a relatively young field and many developments are being produced as more attention is given to this emerging service. In the future, GaaS will likely become an important facet of the cloud computing portfolio and may lead to related GPU services such as on-demand visualization or encryption cracking.

References

1. WiGig Alliance. Defining the Future of Multi-Gigabit Wireless Communications (2011), http://wirelessgigabitalliance.org/?getfile=2132
2. ASUS. Eee Pad Transformer Prime TF201, https://www.asus.com/Tablets_Mobile/Eee_Pad_Transformer_Prime_TF201
3. Choy, S., Wong, B., Simon, G., Rosenberg, C.: The Brewing Storm in Cloud Gaming: A Measurement Study on Cloud to End-user Latency. In: 2012 11th Annual Workshop on Network and Systems Support for Games (NetGames), pp. 1–6 (2012)
4. Cohen, R.: The Cloud Hits the Mainstream: More Than Half of U.S. Businesses Now Use Cloud Computing. Technical report, Forbes (2013)

5. Gaudiosi, J.: New Reports Forecast Global Video Game Industry Will Reach $82 Billion by 2017. Technical report, Forbes LLC (2012)
6. Grubb, J.: World of Warcraft: Mists of Pandaria Sees 1 Million Concurrent Players in China. Technical report, VentureBeat (2012)
7. Hing, D.: EA Admits Disappointment With Warfighter Sales. Technical report, Bit-Tech (2012)
8. Kim, S.-S., Cho, C.: Multiscreen-based Gaming Services using Multi-view Rendering with Different Resolutions. In: Proc. of the Second International Conference on Mobile Services, Resources, and Users (MOBILITY 2012), pp. 34–37 (2012)
9. Moore, G.E.: Cramming More Components onto Integrated Circuits. Electronics Magazine, 38(8) (1965)
10. Moreno, C., Tizon, N., Preda, M.: Mobile Cloud Convergence in GaaS: A Business Model Proposition. In: 2012 45th Hawaii International Conference on System Science (HICSS), pp. 1344–1352 (2012)
11. Nicholson, J., Leavitt, J.: To the Cloud! Anticipating the Legal Issues in Cloud-Based Gaming. Technical report (2011)
12. Ojala, A., Tyrvainen, P.: Developing Cloud Business Models: A Case Study on Cloud Gaming. IEEE Software 28(4), 42–47 (2011)
13. Parkvall, S., Dahlman, E., Furuskar, A., Jading, Y., Olsson, M., Wanstedt, S., Zangi, K.: LTE-Advanced - Evolving LTE towards IMT-Advanced. In: IEEE 68th Vehicular Technology Conference, VTC 2008-Fall, pp. 1–5 (2008)
14. Robinson, J.: The Evolution of the Gaming Industry into the Pockets of the Consumer. Technical report, Univ. of Southampton, UK (2012)
15. Roettgers, J.: EyeIO: Netflixs Secret Weapon Against Bandwidth Caps? Technical report, GigaOM (2012)
16. Takahashi, D.: Video Game Industry to Hit $70 Billion by 2015, But Growth Will Slow. Technical report, VentureBeat (2010)
17. Tizon, N., Moreno, C., Cernea, M., Preda, M.: MPEG-4-based Adaptive Remote Rendering for Video Games. In: Proceedings of the 16th International Conference on 3D Web Technology, Web3D 2011, pp. 45–50. ACM, New York (2011)
18. van der Meulen, R.: Gartner says 821 million smart devices will be purchased worldwide in 2012; sales to rise to 1.2 billion in 2013. Technical report, Gartner (2012)
19. Watson, P.: Mobile Gaming Market Grows to $9 Billion as OBJE Readies Debut Gaming App. Technical report, BusinessWire (2013)
20. Mark Whiting. GTA4 Launch Sales Break Two World Records (2008), `http://www.1up.com/news/gta4-launch-sales-break-world`
21. Zhang, G.-Q., Zhang, G.-Q., Yang, Q.-F., Cheng, S.-Q., Zhou, T.: Evolution of the Internet and its Cores. New Journal of Physics 10 (2008)