

Do Cloud Games Adapt to Client Settings and Network Conditions?

Minzhao Lyu
UNSW Sydney, Australia
minzhao.lyu@unsw.edu.au

Yifan Wang
UNSW Sydney, Australia
wangyifan.frank@student.unsw.edu.au

Vijay Sivaraman
UNSW Sydney, Australia
vijay@unsw.edu.au

Abstract—Cloud gaming relieves gamers from acquiring highly configured gaming hardware to smoothly play high-profile games. To this objective, cloud gaming platforms operate cloud servers that receive inputs from client devices, render gaming graphics, and stream the gaming scenes back to the client devices in real-time. While significantly expanding the gaming industry, this business model imposes high demand on the Internet that transports gaming video for good user experience (QoE) with a typical requirement of more than 10-20Mbps bandwidth, 100ms or less latency, and less than 5% packet drops. In this paper, focusing on two leading platforms, Nvidia's GeForce NOW and Microsoft's Xbox Cloud Gaming, we systematically profile and compare how cloud games adapt their streaming characteristics to client settings (including game streaming frame rate and graphic resolution) and network conditions like bandwidth, latency, and packet drop. Key insights are obtained such as Nvidia's GeForce NOW optimizes its game streaming frame rate and graphic resolution for a smooth user experience, particularly for users using its native application; Xbox does not reactively adapt and shows better tolerance for poor network conditions on PC browsers compared to its proprietary gaming console.

I. INTRODUCTION

Cloud gaming has become a USD \$3.34B industry worldwide and is projected to have a further expanded market value of USD \$18.71B in 2027, as predicted by [1]. Major technology companies that hold advantages in cloud computing, graphic processing, or digital entertainments are developing, releasing and operating their cloud gaming platforms to seize this fast-growing marketplace for both desktop and mobile gamers such as Nvidia's GeForce NOW, Microsoft's Xbox Cloud Gaming, Sony's PS5 cloud streaming, Amazon's Luna and Meta's Facebook Cloud Gaming.

This emerging business model aims to reduce or eliminate the hardware barriers for casual and regular players to access high-profile games that often require highly configured graphic cards and large volume of storage space. In this model, instead of hosting tens of gigabytes of gaming contents on local devices and having (poor- or medium-quality) graphic scenes rendered by local hardware, users have their gaming storage, logical processing, and graphic rendering all on cloud platforms. User actions and gaming scenes are synchronized between clients and cloud servers in real-time. The operational mechanism offloads heavy computational tasks from

local devices to cloud clusters, which inevitably introduces significant demands on network quality-of-service. As reported by F. Aumont *et al.*, a generally good user experience on cloud gaming requires a minimum of tens of Mbps bandwidth, less than 100ms latency, and less than 5% packet drop rate [2].

Given the important role of Internet service providers (ISPs) in cloud gaming user experience, the first step ISPs should take to support a satisfying cloud gaming user experience is to understand the appropriate levels of network QoS (*e.g.*, bandwidth, latency, and packet drop rate) needed for various levels of cloud gaming QoE across a diverse range of user setups, such as browsers or dedicated software applications on mobile or PC devices. Such understandings of network demand can come with two broad questions: (i) How do cloud games (*i.e.*, game sessions on cloud gaming platforms) adapt their network streaming characteristics to different options of client settings such as frame rate and graphic resolution? (ii) How do cloud games adapt their streamed user experience to various levels of network constraints like bandwidth, latency, and packet drop?

Prior works have studied cloud gaming from various aspects including detecting cloud games and measuring user experience metrics [3], video encoding and decoding on client and server hardware [4], detecting cloud gaming sessions in edge networks [5], and specific network anatomy of cloud games served by certain platforms like Google's Stadia [6] and Nvidia's GeForce NOW [7]. While prior works [8], [9] analyzed the changes of network behaviors (*e.g.*, streaming throughput and packet rate) of cloud games under constrained network conditions, the two aforementioned questions regarding the adaptability of cloud games, particularly in their streamed user experience, have not been fully studied.

In this paper, we systematically profile how cloud games adapt its streaming characteristics to various client settings and constrained network conditions, with a particular focus on two major cloud gaming platforms (*i.e.*, GeForce NOW and Xbox Cloud Gaming) that together take over 90% global market. Using our controlled lab setup (in §II) that allows us to enforce traffic control policies (*e.g.*, limiting bandwidth) and measure network traffic before and after the shaping point, we make two specific contributions.

For the **first contribution** (in §III), we investigate into how cloud games adapt their network streaming characteristics to client settings including frame rate and graphic resolution.

This project is funded by the Australian Government's Cooperative Research Centres Projects (CRC-P) Grant CRCPXIV000099.

ISBN 978-3-903176-63-8 ©2024 IFIP




Fig. 1. Our lab setup.

By adjusting available client settings on different user setups while playing games of various genres, we measure the network characteristics of game streaming. We observe that GeForce NOW actively adapts streaming throughput to client settings on frame rate and graphic resolution while Xbox does not. Additionally, GeForce NOW does not exhibit different streaming characteristics across game titles, whereas Xbox streams at a low frame rate for certain game titles, likely due to less-optimized graphics computation.

For the **second contribution** (in §IV), we study how cloud games adapt their streamed user experience (*i.e.*, frame rate and graphic resolution) to network constrained such as bandwidth, latency and packet drop. By enforcing different levels of network constraints on cloud games played via various user setups, we measure the user experience along with their network streaming characteristics. Our analysis reveals unique adaptations used by the two platforms. For example, GeForce NOW prioritizes frame rate over graphic resolution for a smooth streaming experience with limited bandwidth conditions, while Xbox does not. GeForce NOW adapts well to high latency on its native application compared to the browser, whereas Xbox offers a better streaming experience on the browser than on its proprietary hardware console.

II. MEASUREMENT SETUP AND METHODOLOGY

We now describe our lab setup for measuring cloud gaming network traffic with various client settings and network condition constraints.

A. Lab Setup

We set up our lab environment to collect network traffic traces of cloud gaming sessions exchanged between client devices and cloud gaming servers. The logical schematic is shown in Fig. 1. The client devices consist of mobile (Android phone and iOS tablet), PC (macOS PC and Windows PC) and Xbox gaming console. They are wired/wireless connected to the access gateway (*i.e.*, a home router with 1Gbps bandwidth capacity) to the Internet. All traffic exchanged between the Internet and the access gateway are shaped by a traffic control (TC) proxy that runs Linux TC commands to enforce controlled network QoS constraints such as bandwidth limitation, latency and packet drop rate. In this paper, the client devices in our university lab are communicating with regional cloud gaming servers in our city operated by GeForce NOW and Xbox Cloud Gaming, thus, are under nearly ideal network QoS conditions of 1Gbps available bandwidth, less than 10ms latency, and nearly 0% packet drop rate when additional constraints by our traffic control proxy are enforced.

TABLE I

CONTROLLED PARAMETERS IN OUR CLOUD GAMING MEASUREMENT.

	GeForce NOW	Xbox Cloud Gaming
User setup	mobile/PC; browser/app	mobile/PC/Xbox console; browser/app
Client setting	graphic resolution; video frame rate	graphic resolution
Network constraint	bandwidth; latency; packet drop rate	bandwidth; latency; packet drop rate
Game title	CS:GO, Cyberpunk 2077, Monster Hunter, Flight Simulator	GTA V, FIFA 23, Sword & Fairy 7

As indicated by the green and red arrows in Fig. 1, traffic traces of cloud games are captured at two vantage points before and after the traffic control proxy, so that we can analyze the network characteristics before and after the QoS constraints are introduced, to understand the streaming profiles at both near client- and near server-side.

B. Measurement Methodology

User setup: We have collected network traffic traces (*i.e.*, PCAP files) of cloud gaming sessions served by both GeForce NOW and Xbox Cloud Gaming via their supported user setups. As indicated in Table I, we used both browsers (Chrome and Safari) and the native GeForce NOW application on both mobile and PC devices for GeForce NOW cloud games. For Xbox Cloud Gaming, we used browsers on mobile/PC devices and the Xbox proprietary hardware gaming console.

Game selection: To capture variations introduced by game genres and titles that have different availability on each platform, we played popular titles representing diverse genres, such as CS:GO (first-person shooting), Cyberpunk 2077 (action role-playing), FIFA 23 (sports), etc. Additionally, different games require unequal amounts of computational resources on cloud servers, which may impact the user experience. We included games with high demands on hardware configuration (Microsoft Flight Simulator and Sword & Fairy 7) as well as games that reportedly have low configuration requirements (GTA V and CS:GO).

Client settings: GeForce NOW allows users to select their gaming graphic resolution and streaming frame rate to better suit their display hardware, while Xbox Cloud Gaming only provides the flexibility in graphic resolution. Therefore, as indicated by the second row of Table I, we collected traffic traces by varying the graphic resolution from SD (480p) to UHD (4K) and frame rate from 30 to 60 fps on their supported platforms. The results will be discussed in §III.

Network conditions: Both platforms support an “auto” option for client settings. In this mode, the platforms will automatically configure gaming graphic resolution and frame rate for the user with respect to the client hardware specifications and network conditions. We tune three common network QoS constraints (*i.e.*, bandwidth, latency and packet drop rate) to investigate how each platform adapts game experience (QoE) accordingly. The results will be discussed in §IV.

III. CLOUD GAMES ADAPT TO CLIENT SETTINGS

In this section, we profile how cloud games adapt their network streaming characteristics when client selects different




Fig. 2. Time series plots for frame rates and streaming throughput of cloud games with different client resolution settings.

levels of QoE settings (*i.e.*, graphic resolution and streaming frame rate) on two major platforms, namely Nvidia’s GeForce NOW (§III-A) and Xbox Cloud Gaming (§III-B). Key observations are summarized in §III-C.

A. Nvidia’s GeForce NOW

Fig. 2(a) shows the downstream throughput and game-play streaming frame rate of First Person Shooting (CS:GO) gameplay sessions on Windows PC console application with Ultra-High Definition (4K), Full High Definition (1080p), High Definition (720p) and Standard Definition (480p) graphic resolution settings. We note that GeForce NOW also allows users to configure their preferred maximum frame rate for gameplay streaming as either 30 or 60 fps. Therefore, we fixed the maximum frame rate option as 60 fps in Fig. 2(a). Our other experiments with frame rates set to 30fps reached consistent conclusions, thus, they are not shown here for simplicity. To measure the actual frame rate received by the client, we use the number of frame markers per second in the streaming RTP flows [10] as a ground-truth indicator, as introduced in [3].

From Fig. 2(a), we can clearly observe that the maximum streaming throughput of GeForce NOW gameplay sessions is directly determined by the specified graphic resolution. In the beginning phase of each gameplay, when we are in the waiting room for an upcoming game match, the throughput stays at a relatively dynamic but low level. It reaches the maximum level for the respective graphic resolution band (35Mbps, 28Mbps, 16Mbps, and 13Mbps for UHD, FHD, HD, and SD, respectively) after the actual game match begins.

In addition to the CS:GO cloud game sessions with a 60fps frame rate, we also compare the distribution of streaming throughput for another game title (*i.e.*, Cyberpunk 2077) with either a 60fps or 30fps maximum frame rate. Unsurprisingly, with the same graphic resolution, a lower frame rate at 30fps results in lower throughput.

We do not observe differences in streaming throughput among various game titles (*e.g.*, Monster Hunter and Flight Simulator) and genres. Similar insights are obtained for other supported user setups, including PC browsers and mobile console applications.

B. Microsoft’s Xbox Cloud Gaming

Xbox does not allow users to choose the game streaming frame rate. In Fig. 2(b), we first show the streaming throughput and frame rate for the four resolution levels of GTA cloud games on the Xbox hardware console. Unlike GeForce NOW, the throughput of Xbox Cloud Gaming does not adapt to different resolution settings, as evidenced by the constant

level of peak bandwidth usage (while higher variations are observed) for the four resolution bands on the Xbox hardware console.

The second observation we made for Xbox Cloud Gaming platform is that, the streaming throughput reduces for game titles that are offered with low frame rates by the cloud servers, perhaps due to high graphic rendering costs. One such example is given in Fig. 2(c) where we play Fairy and Sword, a role-playing game with reportedly high requirements [11] for graphic rendering. The frame rate under all resolutions (*i.e.*, UHD to SD) stays around 30fps and the maximum streaming throughput remains constantly at 13Mbps, in contrast to the 60fps games such as GTA, also shown in Fig. 2(c).

In Xbox cloud game sessions on different user setups, we observed changes only in the numerical values for the maximum streaming throughput. The conclusions from our analysis remain consistent.

C. Highlights

We now draw two key highlights from our analysis and comparison of the two platforms.

Firstly, our observation reveals that GeForce NOW optimizes streaming throughput for clients requesting different levels of **graphic resolution**, maintaining the same maximum frame rate across various game titles. In contrast, Xbox exhibits a constant streaming throughput irrespective of user setup types and the requested resolution. Consequently, for cloud gaming platforms like GeForce NOW, Internet service providers can leverage bandwidth as a robust indicator of the user’s preferred graphic resolution, facilitating a better understanding of user preferences. On the other hand, platforms like Xbox Cloud Gaming, which maintain a constant streaming quality, require ISPs to provide a unified minimum bandwidth for ensuring a consistently good cloud gaming experience.

Secondly, GeForce NOW dynamically adapts its streaming throughput for clients choosing different **frame rate settings**. In contrast, Xbox Cloud Gaming maintains a fixed frame rate for various game titles based on their well-optimized or less-optimized graphic rendering costs. Frequent instances of low frame rates (*e.g.*, 30fps) observed from a user on the GeForce NOW platform are likely a voluntary choice, considering the user’s display device’s refresh rate. However, similar observations from users on Xbox Cloud Gaming can be attributed to the specific games being played. Importantly, users on Xbox Cloud Gaming or similar platforms may not be aware of the underlying reason, potentially attributing (mistakenly) the network QoS provided by their Internet Service Providers as the root cause.



(a) Good resolution. (b) reduced resolution. (c) Unplayable resolution.
Fig. 3. Three types of graphic quality of cloud games a user can perceive: (a) good, (b) reduced and (c) unplayable.

IV. CLOUD GAMES ADAPT TO NETWORK CONDITIONS

In this section, we profile how GeForce NOW and Xbox cloud games adapt their streaming QoE to constrained network conditions (§IV-A), including limited bandwidth (§IV-B), increased latency (§IV-C) and high packet drop rate (§IV-D).


A. Clients' Perception on Cloud Game User Experience (QoE)

Before delving into the measurement results from the network traffic, we first visually examine different levels of graphic qualities that a user can perceive, as depicted in Fig. 3. The first type of quality is defined as **good resolution** (Fig. 3(a)). With this level of quality, users can clearly discern the graphic details matched with the capability of their displays, including UHD monitors, FHD laptop screens, or HD mobile screens. When the network QoS cannot fully support the good resolutions for the respective display settings, users may experience **reduced resolution** (Fig. 3(b)). The graphics are downgraded to a lower resolution level, depending on the current level of network QoS, such as FHD, HD, or SD on a UHD display. Despite the reduction in visual quality, users can still play the game, albeit with a degraded visual experience. The third type of graphic experience, which occurs when the users are not able to see complete gaming scenes, is named **unplayable resolution**. As shown in Fig. 3(c), the game scene is fragmented into multiple blocks that are inconsistently synchronized with each other, as opposed to just being blurry, as is the case with “reduced resolution”. Users find it challenging to proceed with gameplay due to such a distorted graphic experience.


In addition to the graphic resolution, the cloud gaming user experience is also directly dependent on the streaming frame rate. A high value (*e.g.*, near 60fps) indicates smooth transitions between video frames and accurate synchronization of user input with the cloud servers. A very low frame rate (*e.g.*, below 30fps) can result in discrepancies in game scenes and unresponsiveness to user motions. As discussed in §III, the streaming frame rate can be quantitatively and directly measured by counting the frame markers in the respective RTP flows for our analysis.

B. Bandwidth

Clients who subscribe to broadband network services often have an expected bandwidth as part of their service level agreements. However, in practice, such available bandwidth may vary (*e.g.*, become less than what is advertised) due to factors such as the time of day (*e.g.*, busy or idle hours) and bottlenecks in the routing path to the accessed service. Therefore, we consider the maximum available bandwidth as a network QoS constraint to study its impact on cloud gaming



(a) Throughput sent by the cloud server.



(b) Frame rate and color-coded graphic resolution.

Fig. 4. Frame rates and streaming throughput of **GeForce NOW** cloud games under limited bandwidths.

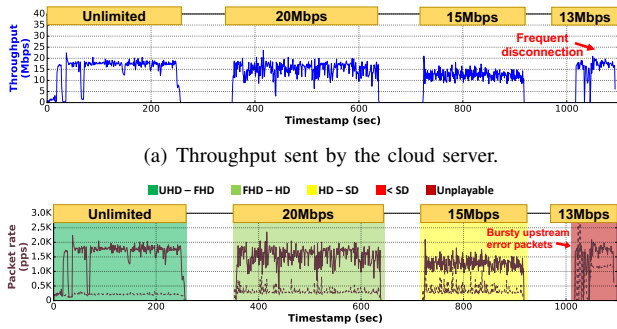
user experience, measured by graphic resolution and game streaming frame rate.

In this set of experiments, we adjust the available bandwidth of the client using our traffic control (TC) proxy, as discussed in §II. For each game title considered in this study, we incrementally reduce the bandwidth from unlimited (*i.e.*, 1Gbps) until reaching the unplayable level, characterized by either frequent disconnection from the server or broken graphics on both cloud gaming platforms. We capture cloud gameplay sessions lasting more than five minutes for each level of bandwidth limitation to draw conclusions from the stable states of gameplay. From the traffic traces captured at the vantage point before the traffic control proxy (indicated by the green arrow in Fig. 1), we verify that the downstream traffic throughput is consistently capped at the configured levels.


1) **GeForce NOW**: For the GeForce NOW platform, we tune the maximum available bandwidth to the client device from unlimited to 20Mbps, 15Mbps, 10Mbps, 5Mbps and 3Mbps, respectively. Fig. 4(a) shows the downstream traffic throughput of game streaming content sent by the server, as measured at the post-shaping vantage point (indicated by the red arrow in Fig. 1). Fig. 4(b) shows the streaming frame rate received by the client device, as measured at the pre-shaping vantage point (indicated by the green arrow in Fig. 1).

It is clear from the plots that GeForce NOW cloud gaming servers adapt the traffic throughput they stream to the client based on network bandwidth limits. With the default settings on our Windows PC connected to a UHD display, the servers sent a maximum 36Mbps traffic throughput to the client under unlimited available bandwidth. The throughput from cloud servers dropped to 18Mbps when a “20Mbps” bandwidth limit is imposed on the client device. Similar observations can be made from Fig. 4(a) for bandwidth limits of “15Mbps”, “10Mbps”, “5Mbps” and “3Mbps”, respectively.

As shown in Fig. 4(b), to adapt the traffic throughput for a lower bandwidth limit, GeForce NOW cloud servers tend to reduce graphic resolution as the first choice before decreasing the frame rate. The frame rate remains at 60fps under unlimited, 20Mbps, and 15Mbps available bandwidth conditions, while the graphic resolution drops from UHD to FHD and HD, respectively. A further reduction of available bandwidth from 15Mbps to 10Mbps results in a drop in frame



(a) Throughput sent by the cloud server.



(b) Upstream packet rate, downstream packet rate, and graphic resolution.

Fig. 5. Streaming throughput, upstream/downstream packet rate of Xbox cloud games under limited bandwidths.

rate from 60fps to 30fps, while the resolution remains HD. Similar observations are obtained for other game titles and user devices. In this figure and other figures in this paper, we use color-coding to represent different graphic qualities in the gameplay regions, associating dark green with UHD, light green with FHD, yellow with HD, red with SD, and dark red with unplayable quality. Notably, in the gameplay sessions discussed above, we encounter an unplayable quality (*i.e.*, graphic fragmentation and discontinuity) when the bandwidth limit is set to 3Mbps. The frame rates for such unplayable scenarios are also unstable and consistently below 30fps.

2) **Xbox Cloud Gaming**: The Xbox Cloud Gaming platform exhibits a significantly different behavior when the available bandwidth to client is insufficient to support the default client QoE settings. We reduce the bandwidth limit from unlimited (1Gbps) to 20Mbps, 15 Mbps, and 13 Mbps. The downstream throughput sent by the cloud gaming servers is shown in Fig. 5(a), and the upstream/downstream packet rate measured at client/server sides is shown in Fig. 5(b).

Unlike GeForce NOW cloud servers, which adapt their streaming throughput according to available bandwidth on the user side for a smooth user experience, Xbox Cloud Gaming servers exhibit less optimization with limited bandwidth. From Fig. 5(a), it is evident that while the streaming throughput from the server drops when bandwidth limits are introduced, its pattern becomes quite unstable. Notably, the streaming throughput from the server remains at a level much higher (around 16Mbps) than what can be delivered to the client, even when the imposed limit (set at 13Mbps) leads to unplayable scenarios with frequent disconnections and graphic fragmentation. Apparently, the required throughput for a smooth Xbox cloud gaming session has not been reactively reduced for limited bandwidth conditions.

The upstream packet rate from the client to the cloud server increases from 300pps under the “unlimited” condition to 400pps and 500pps under 20Mbps and 15Mbps conditions, respectively. Moreover, this value goes up to 1300pps when the game becomes unplayable with a 13Mbps bandwidth limit. After investigating the packet traces, we found that these extra packets are for error acknowledgment in streaming frames.

In terms of user experience, we could not observe any systematic adaptation for graphic resolution and frame rate

when bandwidth is limited. Adding to the inadaptability of Xbox cloud servers, the high volumes of error/lost packets, as just discussed, further overwhelm the already congested network, causing the frame rate to be highly unstable and gaming graphics to be desynchronized and fragmented.

3) **Key Takeaways**: When faced with limited available bandwidth on the client side, GeForce NOW cloud servers optimally adapt the streamed graphic resolution as a primary measure to lower bandwidth demands while keeping frame rates at a relatively high level for a smooth user experience. This technical implementation provides more tolerance to the level of bandwidth ISPs can allocate to cloud gaming users, particularly during peak hours, without subjecting them to an unplayable experience.

In contrast, Xbox Cloud Gaming does not effectively support the dynamic adjustment of the streaming throughput sent from cloud servers based on the client-side bandwidth limitation. In other words, the cloud servers stream the same amount and quality of gaming content to the clients regardless of their available bandwidth. Consequently, users perceive unplayable graphic quality and/or frequent disconnection when the expected client bandwidth cannot be achieved.

C. Latency

Latency is another critical network QoS metric that can directly impact the user experience of networked applications. We now investigate how cloud games served by GeForce NOW and Xbox adapt their streamed user experience to increased latency. Using our traffic control proxy, we incrementally add extra latency between our client device and the cloud gaming servers until the user experience becomes unplayable.

1) **GeForce NOW**: Without any additional latency, our cloud games on the GeForce NOW platform have latency between the client and server sides smaller than 5ms, a value recognized as ideal by the platform. We then introduce extra latency for our gameplay sessions on PC device, ranging from 15ms to 30ms, 50ms, 100ms, 200ms, 300ms until 1000ms, when we experience frequent disconnections.

The streaming throughput sent by the sever is visually presented in Fig. 6(a). Compared to the sessions with moderate levels of extra latency (*i.e.*, 15ms to 200ms), the ideal case (labeled as “0ms” extra latency) has relatively lower throughput (less than 30Mbps) as it can stream more efficiently (*i.e.*, demanding fewer bytes on the wire). As clients, we perceive that the delay of our user input is linearly correlated with the latency, whereas the graphic quality and video frame rate do not change with increased latency up to 300ms. When the extra latency is set to 500ms and more, the cloud server actively reduces the graphic quality from UHD to HD or lower levels for the smooth delivery of gaming frames. This active reaction also results in lower streaming throughput, below 25Mbps.

The sessions via other user setups behave similarly when extra latency is introduced. However, we would like to note that cloud gameplay via browsers (shown in Fig. 6(c)) is less optimal to extra latencies in terms of maintaining graphic

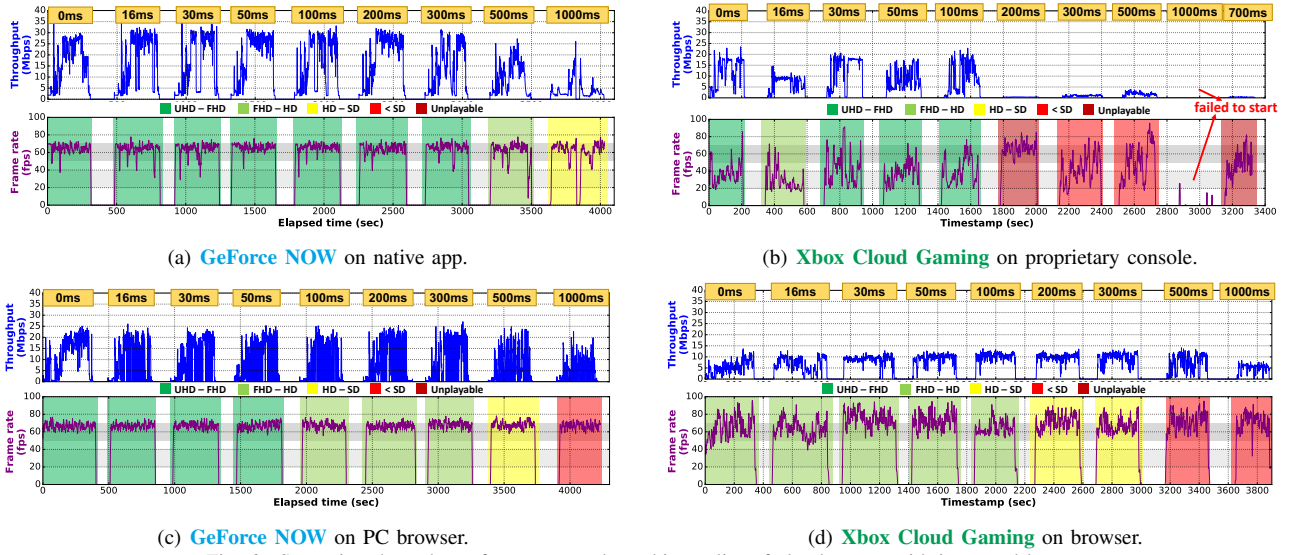


Fig. 6. Streaming throughput, frame rate and graphic quality of cloud games with increased latency.

resolution while having smooth frame delivery. In the shown example, the resolution starts dropping after 100ms.

Compared to limited bandwidth, as discussed in §IV-B, a poor latency condition will not result in fragmentation or discontinuity in graphics, as the GeForce NOW servers can actively reduce the bandwidth demands by tuning the graphic resolution. Instead, with higher latency, clients feel less responsive when generating user inputs until being disconnected by the server.

2) **Xbox Cloud Gaming**: From our experimental results, Xbox Cloud Gaming servers employ different strategies to handle high latency when the gameplay sessions are initiated from different user setups (*i.e.*, Xbox hardware console or PC/mobile browser).

On Xbox’s own **proprietary hardware console**, the platform tends to adjust both graphic resolution and frame rate when extra latencies are introduced (Fig. 6(b)). For instance, in the example we discuss now, with extra latency ranging from 15ms to 50ms, the frame rate drops from 60fps to 30fps, while the graphic resolution remains unchanged at 15ms. Interestingly, when we increase the extra latency to 30ms, the graphic resolution drops from UHD to FHD, but the frame rate returns to 60fps. However, the frame rate drops back to 30fps with 50ms and 100ms extra latency. These adjustments result in different levels of streaming throughput sent by the server. As shown in Fig. 6(b), the maximum throughput for the “0ms” and “30ms” cases is 18Mbps, while for the “15ms”, “50ms” and “100ms” cases, it hovers around 10Mbps. It is worth nothing that the fluctuation in throughput in the “50” ms and “100ms” cases is also caused by error/disordered frames, as discussed in §IV-C.

When the added latency exceeds 200ms, the graphic resolution drops dramatically to an unplayable level (to the point where we can hardly discern the resolution band, but it is worse than SD). As seen in Fig. 6(b), the streaming throughput drops to less than 3Mbps or even zero due to the failure of game initialization.

For the gameplay sessions played via **PC or mobile browsers**, the platform reacts to extra latency in a manner quite similar to GeForce NOW. The graphic resolution and frame rate do not drop until a very high latency (*e.g.*, 1000ms in our example case shown in Fig. 6(d)) is introduced. Again, with an ideal latency of less than 5ms (the leftmost case in Fig. 6(d)), the streaming throughput sent from the cloud server is the most optimized compared to slightly higher latency bands.

3) **Key Takeaways**: Both Xbox Cloud Gaming and GeForce NOW performs well when the latency is 100ms or lower. When the latency reaches 200ms or beyond, the user inputs start becoming unresponsive and eventually unplayable or gets disconnected by the server.

GeForce NOW reacts similarly for all user setups, including PC/mobile devices running native GeForce NOW applications or browsers. However, native applications are more capable than browsers in maintaining good graphic resolution while having smooth frame streaming.

Xbox Cloud Gaming demonstrates different responses to increased latency for cloud games played on the Xbox hardware console and mobile/PC browsers. Both graphic resolution and frame rate are adjusted when the latency is increased at a relatively low level (*e.g.*, from 30ms) on the hardware console, while the browser sessions start to adapt their graphic resolution only at a higher latency level (*e.g.*, 500ms).

D. Packet Drop Rate

In addition to bandwidth limit and latency, packet drop rate is another important network QoS metric that ISP often optimize against. Lost packets can be caused by many reasons related to telecommunications networks, such as network congestion, limited switch/router buffer sizes and hardware/software issues of the routing devices. The impact of packet drop rates on online gaming and video streaming user experience has been extensively studied, while little is known for cloud gaming. Similar to our measurement study on

bandwidth and latency, we use our traffic control proxy shown in Fig. 1 to incrementally introduce higher packet loss rates (from 0% to 50%) till the user experience becomes unplayable (i.e., disconnected from the server and unable to restart a gameplay session).

For given the limitation of pages, we skip our analytical insights for each platform and go straight to key takeaways.

GeForce NOW exhibits better mechanisms to adapt graphic resolutions and maintain decent streaming frame rates under varying packet loss rates compared to Xbox Cloud Gaming. GeForce NOW can provide a playable user experience even with a small level of packet loss (e.g., 2%). In contrast, Xbox Cloud Gaming tends to offer an unplayable user experience under similar conditions. Also, GeForce NOW's native applications on mobile and PC devices adapt their graphic resolutions to different level of packet drop rates more precisely than those via generic browsers, which takes a radical drop of resolution when the drop rate exceeds 1%. In comparison, browser sessions only show a radical drop in resolution when the packet drop rate exceeds 1%.

V. RELATED WORK

The user experience of cloud gaming has been the focus of many prior works – from evaluating QoE demands [12]–[14], measuring user experience [3], [5], [6], [15]–[20], to optimizing cloud gaming system architectures [21]–[23]. K. Chen *et al.* [24] and S. Schmidt *et al.* [25] investigated various factors that can impact cloud gaming user experience, and identified bandwidth, network delay and packet loss as some of the most important ones. H. Iqbal *et al.* [4] measured the user-perceived QoE in browser-based cloud gaming sessions. They highlighted that cloud gaming platforms demonstrated different and often limited capabilities of handling network impairments. S. Bhuyan *et al.* [8] focused on the cloud gaming performance and energy consumption on mobile platforms under wireless (i.e., Wi-Fi and 5G/4G) networks. The works in [2], [7], [9] evaluated the adaptation strategies of cloud gaming platforms under network QoS constraints such as limited bandwidth and increased delays. In this work, we are the first to investigate the adaptability of cloud gaming services under constrained network conditions on different hardware (e.g., PC versus gaming console) and software platforms (e.g., browser versus native application).

VI. CONCLUSION

The increasing popularity of the cloud gaming business model imposes a high demand on the quality of service (QoS) offered by Internet service providers (ISPs), which can play a deterministic role in the gameplay experience (QoE) perceived by users. In this paper, we empirically profile the adaptability of two major cloud gaming platforms, namely Nvidia's GeForce NOW and Microsoft's Xbox Cloud Gaming, to systematically understand how they adapt their network streaming characteristics to client settings and how they optimize gaming streaming experience to network conditions. The insights obtained from this study provide a reference for ISPs

and relevant stakeholders to optimally manage their networks for the expected quality level of cloud gaming user experience on the two major cloud gaming platforms.

REFERENCES

- [1] Statista, "Cloud Gaming - Worldwide," 2023, accessed: 2023-11-23. [Online]. Available: <https://tinyurl.com/yt4vsvdf>
- [2] F. Aumont, F. Humbert, C. Neumann, C. Salmon-Legagneur, and C. Taibi, "Dissecting Cloud Game Streaming Platforms Regarding the Impacts of Video Encoding and Networking Constraints on QoE," in *Proc. ACM GameSys*, Istanbul, Turkey, Sep 2021.
- [3] M. Lyu, S. C. Madanapalli, A. Vishwanath, and V. Sivaraman, "Network Anatomy and Real-Time Measurement of Nvidia GeForce NOW Cloud Gaming," in *Proc. PAM*, Virtual Event, Mar 2024.
- [4] H. Iqbal, A. Khalid, and M. Shahzad, "Dissecting cloud gaming performance with decaf," *Proc. ACM Meas. Anal. Comput. Syst.*, Dec 2021.
- [5] P. Graff *et al.*, "Efficient Identification of Cloud Gaming Traffic at the Edge," in *Proc. IEEE/IFIP NOMS*, Miami, FL, USA, May 2023.
- [6] M. Carrascosa and B. Bellalta, "Cloud-gaming: Analysis of Google Stadia traffic," *Computer Communications*, vol. 188, Apr. 2022.
- [7] M. Suznjevic, I. Slivar, and L. Skorin-Kapov, "Analysis and QoE Evaluation of Cloud Gaming Service Adaptation under Different Network Conditions: The Case of NVIDIA GeForce NOW," in *Proc. QoMEX*, Lisbon, Portugal, Jun 2016.
- [8] S. Bhuyan, S. Zhao, Z. Ying, M. T. Kandemir, and C. R. Das, "End-to-end characterization of game streaming applications on mobile platforms," *Proc. ACM Meas. Anal. Comput. Syst.*, Feb 2022.
- [9] X. Marchal *et al.*, "An Analysis of Cloud Gaming Platforms Behaviour Under Synthetic Network Constraints and Real Cellular Networks Conditions," *J. Netw. Syst. Manage.*, Feb 2023.
- [10] H. Schulzrinne *et al.*, "RTP: A Transport Protocol for Real-Time Applications," RFC 3550, Jul. 2003.
- [11] Vedor, "Sword and Fairy 7 FPS low FPS on equipment page," 07 2023, accessed: 2024-01-03. [Online]. Available: <https://tinyurl.com/56f94eub>
- [12] I. Slivar *et al.*, "The Impact of Video Encoding Parameters and Game Type on QoE for Cloud Gaming: A Case Study Using the Steam Platform," in *Proc. QoMEX*, Pilos, Greece, May 2015.
- [13] S. Zadtootaghaj, S. Schmidt, and S. Moller, "Modeling Gaming QoE: Towards the Impact of Frame Rate and Bit Rate on Cloud Gaming," in *Proc. QoMEX*, Cagliari, Italy, May 2018.
- [14] S. S. Sabet *et al.*, "Towards the Impact of Gamers Strategy and User Inputs on the Delay Sensitivity of Cloud Games," in *Proc. QoMEX*, Athlone, Ireland, May 2020.
- [15] T. Kämäräinen, M. Siekkinen, A. Ylä-Jääski, W. Zhang, and P. Hui, "A Measurement Study on Achieving Imperceptible Latency in Mobile Cloud Gaming," in *Proc. ACM MMSys*, Taipei, Taiwan, Jun 2017.
- [16] C. Baena *et al.*, "Measuring and estimating Key Quality Indicators in Cloud Gaming services," *Computer Networks*, Jul 2023.
- [17] M. Carvalho, D. Soares, and D. F. Macedo, "Transfer Learning-Based QoE Estimation For Different Cloud Gaming Contexts," in *Proc. IEEE NetSoft*, Madrid, Spain, Jun 2023.
- [18] J. R. Ky, P. Graff, B. Mathieu, and T. Cholez, "A Hybrid P4/NFV Architecture for Cloud Gaming Traffic Detection with Unsupervised ML," in *Proc. IEEE ISCC*, Gammarrth, Tunisia, Jul 2023.
- [19] M. Suznjevic *et al.*, "Towards Understanding the Relationship between Game Type and Network Traffic for Cloud Gaming," in *Proc. IEEE ICMEW*, Chengdu, China, Jul 2014.
- [20] A. Di Domenico *et al.*, "A Network Analysis on Cloud Gaming: Stadia, GeForce Now and PSNow," *Network*, Oct 2021.
- [21] H. Chen *et al.*, "T-Gaming: A Cost-Efficient Cloud Gaming System at Scale," *IEEE TPDS*, Dec 2019.
- [22] P. Hamadani, D. Gallatin, M. Alizadeh, and K. Chintalapudi, "Ekho: Synchronizing Cloud Gaming Media across Multiple Endpoints," in *Proc. ACM SIGCOMM*, New York, NY, USA, Sep 2023.
- [23] J. Wu, Y. Guan, Q. Mao, Y. Cui, Z. Guo, and X. Zhang, "ZGaming: Zero-Latency 3D Cloud Gaming by Image Prediction," in *Proc. ACM SIGCOMM*, New York, NY, USA, Sep 2023.
- [24] K.-T. Chen, Y.-C. Chang, H.-J. Hsu, D.-Y. Chen, C.-Y. Huang, and C.-H. Hsu, "On the Quality of Service of Cloud Gaming Systems," *IEEE Transactions on Multimedia*, Feb 2014.
- [25] S. Schmidt, S. Uhrig, and D. Reuschel, "Investigating the Relationship of Mental Immersion and Physiological Measures during Cloud Gaming," in *Proc. QoMEX*, Athlone, Ireland, May 2020.