

Systematic assessment of cloud game adaptability for network conditions and user experience

Minzhao Lyu ^{ID}*, Yifan Wang ^{ID}, Vijay Sivaraman ^{ID}

School of Electrical Engineering and Telecommunications, University of New South Wales, 2052, Sydney, NSW, Australia

ARTICLE INFO

Keywords:

Cloud gaming
Network measurement
Quality-of-service
Quality-of-experience

ABSTRACT

Cloud gaming platforms lower the access barriers to graphics-intensive games by rendering computationally heavy game scenes on cloud GPU servers and streaming them back to players as real-time video, which in turn places significant demands on carrier networks to deliver these video streams with high throughput, low latency and minimal packet loss. To achieve decent user experience, cloud gaming platforms adapt streaming behaviors based on network conditions and allow users to adjust their graphics settings. Knowing the level of game streaming adaptability offered by various cloud gaming providers is helpful for network operators to effectively provision network resources for subscriber satisfaction, and for game development community to incentivize cloud gaming providers to better optimize their streaming techniques. Toward this objective, we develop a systematic framework to assess the adaptability of a cloud gaming platform in reducing network demand for lower client graphics settings; and in adjusting streaming quality under constrained network conditions for smooth gaming experience. Focusing on four popular platforms (NVIDIA GFN, Xbox, PlayStation and Amazon Luna), we begin by empirically profiling and comparing how they adapt game streaming characteristics to various levels of client graphics settings and network conditions. Building on the insights, we develop our systematic assessment framework, which provides quantitative scores for both fine-grained metrics by processing labeled traffic traces, as well as aggregated scores tailored to an assessor's preference. We showcase our quantitative assessments of the four platforms.

1. Introduction

Cloud gaming aims to reduce the hardware barriers for casual and regular players to access graphics intensive games that often require highly configured graphics 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 all their game storage, logical processing, and graphics rendering performed on cloud platforms. User actions and game 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. For example, a generally good user experience of cloud gaming often requires a minimum of tens of Mbps bandwidth, less than 100ms latency, and less than 5% packet drop rate [1].

Given the important role Internet Service Providers (ISP) can play in cloud gaming user experience, understanding the right levels of network QoS (e.g., bandwidth, latency and packet drop rate) they should provide

to support a certain level of cloud gaming user QoE across (a diverse range of) browsers or dedicated software applications on either mobile or PC devices is the first step ISPs should take to guarantee superlative cloud gaming user experience. Such understandings of demands can be summarized as two broad questions: (i) How do cloud games adapt their network streaming characteristics to different options of **active client settings** such as frame rate and graphic resolution? (ii) How does the user-perceived experience change as cloud games adapt their streaming quality to various levels of **passive network conditions** like bandwidth, latency and packet drop? In addition to help network operators better understand the diverse demands across cloud gaming contexts for effective network resource provisioning, being able to answer the two questions can also incentivize cloud gaming providers to better optimize their streaming techniques for the resilience in user-perceived streaming quality under (unavoidable) degraded network conditions.

Prior works have studied cloud gaming from various aspects including detection of cloud games and measurement of user experience metrics [5,6,8], analysis of video encoding and decoding on client and server hardware [2], and network anatomy of specific cloud gaming platforms

* Corresponding author.

E-mail addresses: minzhao.lyu@unsw.edu.au (M. Lyu), wangyifan.frank@student.unsw.edu.au (Y. Wang), vijay@unsw.edu.au (V. Sivaraman).

<https://doi.org/10.1016/j.comnet.2026.112214>

Received 18 August 2025; Received in revised form 6 January 2026; Accepted 10 March 2026

Available online 12 March 2026

1389-1286/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Table 1
Summary of recent research works that analyze cloud gaming network traffic.

Work	Platforms	Game streaming QoE metrics	Investigated impact factors	Developed methodology
[1]	GFN, Stadia	Image quality scores; playability score	Bandwidth; latency; packet loss; video codec	Only analysis study
[2]	GFN, Stadia; Luna	Resolution; frame rate; responsiveness	Bandwidth; latency; packet loss; game genre	CG client performance monitoring
[3]	Stadia	Resolution; responsiveness	Bandwidth; latency; video codec	Only analysis study
[4]	GFN, xCloud, Stadia, PS5	Resolution; frame rate	Bandwidth; latency; packet loss	Only analysis study
[5,6]	GFN, xCloud	Resolution; frame rate	User setup	Session detection and QoE measurement
[7]	GFN	Resolution; frame rate	Game title/genre; player activity	Streaming context classification
Ours	GFN, xCloud, Luna, PS5	Resolution; frame rate; responsiveness	Bandwidth; latency; packet loss	Platform adaptability assessment framework

like Google's Stadia [9] and NVIDIA's GeForce NOW [10]. While prior works [4,11] analyzed the changes of network behavior (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 streaming experience, have not been fully studied.

In this paper, we systematically assess how cloud games adapt their streaming characteristics to various client settings and constrained network conditions, with a particular focus on four major cloud gaming platforms, namely NVIDIA GeForce NOW (GFN), Microsoft Xbox Cloud Gaming (xCloud), Sony PS5 Cloud Streaming and Amazon Luna, that together take over 60% of the global market [12]. Using our controlled lab setup (in Section 3) that allows us to enforce traffic control policies (e.g., limiting bandwidth) and measure network traffic before and after the shaping point, we make four specific contributions.

For the **first contribution** (in Section 4), we empirically investigate how cloud gaming platforms adapt their network streaming characteristics to client graphics settings including frame rate and graphic resolution, by adjusting the available graphics settings across various user setups on all four platforms and measuring their network characteristics. We observe interesting insights such as that GeForce NOW reacts to both changes in client resolution and frame rate settings, whereas Xbox Cloud Gaming does not adapt its network characteristics at all.

For the **second contribution** (in Section 5), we empirically study how user experience metrics (i.e., frame rate and graphic resolution) change as the cloud gaming platforms adapt their streaming quality to network Quality-of-Service (QoS) conditions 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 four platforms. For example, GeForce NOW, aligned with previous findings [10], prioritizes frame rate over graphic resolution under limited bandwidth conditions and hence can keep relatively good smoothness even with low available bandwidth, while Xbox does not adapt to the bandwidth condition change at all, leading to poor experience even with moderately limited bandwidth.

Driven by our analytical insights from representative case studies, our **third contribution** (in Section 6) develops a systematic schema that quantitatively assesses network adaptability of cloud gaming platforms, which can be applied on large-scale dataset for converged quantitative conclusions and easily replicated for other emerging platforms or for regular assessments. The schema is structured with four hierarchies from fine-grained adaptability metrics of cloud game streaming on a certain user setup under a specific active or passive condition (e.g., user actively changing frame rate or varying available bandwidth), to aggregated scores of each metric under active or passive conditions and the two categories, and overall adaptability of an assessed cloud gaming platform.

The **fourth contribution** (in Section 7) implements an automatic process to quantitatively assess the network adaptability of a given cloud gaming platform. The process takes labeled packet captures of cloud gameplay sessions and generates scores at both fine-grained and aggre-

gated levels with default or customized weighting for each fine-grained category of user setups and active/passive conditions. We then report our assessment scores for the four studied cloud gaming platforms using labeled packet captures under three example assessment priorities including equal weights and prioritization on casual or mobile gamers. The results from our lab dataset covering 439 cloud gaming sessions showcase that GeForce NOW and Amazon Luna outperform Xbox and PS5 Cloud Streaming in terms of overall rating. PS5 Cloud Streaming has the highest adaptability under constrained network QoS conditions but performs poorly for active client configurations.

2. Related work

2.1. Analyzing cloud gaming quality-of-experience (QoE)

The user experience of cloud gaming has been the focus of many prior works – from evaluating QoE demands [13–16], measuring user experience [5,6,8,9,17–22], to optimizing cloud gaming system architectures [23–25]. As examples, K. Chen et al. [26] and S. Schmidt et al. [27] 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, among other indirect factors such as jitter and mobile handover. H. Iqbal et al. [2] 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. The works in [1,4,10] evaluated the adaptation strategies of cloud gaming platforms under network QoS constraints such as limited bandwidth and increased delays. For recent research works that analyze cloud gaming network traffic as provided in Table 1, they cover one or few popular cloud gaming platforms and analyze certain set of game streaming quality-of-experience (QoE) metrics as impacted/determined by different factors including network QoS conditions [1,3,4] and game contexts [2,5–7]. They either conducted pure analysis study of cloud gaming traffic characteristics under various network QoS conditions [1,4], or developed real-time traffic analysis technologies [5–7] to detect cloud gaming sessions and measuring user experience metrics for operational networks. In this work, we develop a measurement process to assess the streaming adaptability of cloud gaming platforms under constrained network conditions on different hardware (e.g., PC vs gaming console) and software platforms (e.g., browser vs native application) by analyzing cloud gaming traffic traces labeled by their respective user setups, streaming settings, network conditions, and subjective user experience levels.

2.2. QoE-driven network analysis of multimedia applications

Given their different demands and traffic behaviors, various methods have been proposed in the past to measure and enhance the user experience of multimedia applications (e.g., video streaming and conferencing) from the network perspective. *iTelescope* [28] developed an SDN-based traffic processing architecture for the classification and telemetry of video streams. *ReLive* [29] compared the traffic characteristics of on-demand and live video streaming. O. Michel et al. [30] analyzed Zoom's network behaviors, and developed a methodology for

Table 2
Controlled parameters in our cloud gaming measurement.

	GeForce NOW	Xbox Cloud Gaming	Amazon Luna	PS5 Cloud Streaming
User Setup	mobile/PC; browser/APP	mobile/PC/Xbox console; browser/APP	mobile/PC; browser/APP	mobile/PC/PlayStation console; APP
Client Setting	graphic resolution; streaming frame rate bandwidth; latency;	graphic resolution bandwidth; latency;	streaming frame rate bandwidth; latency;	graphic resolution bandwidth; latency;
Network Condition	packet drop rate CS2; Cyberpunk 2077;	packet drop rate GTA V; FIFA 23;	packet drop rate	packet drop rate Slime Rancher; Baldurs Gate 3;
Game Title	Monster Hunter; Flight Simulator	Sword & Fairy 7	Fortnite; Kunai	Riders Republic




Fig. 1. Our lab setup.

detecting Zoom meetings and identifying meeting types (e.g., single vs multiple users) by inspecting its packet headers. The authors in [31] proposed a signature-based online gaming detection method and studied the relationship between server locations and latencies of online games across three different Internet Service Providers (ISP). Cloud gaming is characterized by its unique attributes of combining computation-intensive cloud computing with bandwidth-demanding video streaming and latency-sensitive online gaming. In this work, we specifically focus on understanding the network demands of cloud gaming and how it reacts to different network conditions, enabling ISPs to better plan for the growth of cloud gaming traffic in their infrastructure and for cloud gaming platforms to optimize their adaptability to network constraints.

3. Measurement setup and methodology

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

3.1. 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 laptop and Windows PC) and Xbox gaming console. They are wired/wireless connected via 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 [32] 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 communicate 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 before additional constraints by our traffic control proxy are enforced. The setup and data collection process can be easily replicated by professional data engineering teams in a Digital Twin of an industry assessor (e.g., ISP) to enrich the dataset as per their specific operational demands, such as to incorporate more devices popular among user or to cover different network environments.

3.2. Dataset

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.

User Setup: As shown in Table 2, we have collected network traffic traces¹ (i.e., PCAP files) of cloud gaming sessions served by GeForce NOW, Xbox Cloud Gaming, Amazon Luna and PS5 Cloud Streaming via their supported user setups. Specifically, 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 third-party browsers and native applications on mobile devices, browsers on PC devices, and native applications on the Xbox proprietary hardware gaming console. Amazon Luna supports both browsers and native applications on mobile devices but only supports browsers on PC devices, which are all covered in our dataset. PS5 Cloud Streaming offers cloud gaming services via its native applications on mobile, PC and gaming consoles that are all included in our data collection.

Game Selection: To capture possible variations introduced by game genres and titles, we played popular titles representing diverse genres, such as CS2 (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, Sword & Fairy 7 and Baldur's Gate 3), as well as games that reportedly have low configuration requirements such as GTA V [34] and CS2.

Active Client Settings: GeForce NOW and Amazon Luna allow users to select their gaming graphic resolution and streaming frame rate to better suit their display hardware, while Xbox Cloud Gaming and PS5 Cloud Streaming only provide the flexibility in graphic resolution. Therefore, as indicated by the second row of Table 2, 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, covering the most popular graphics settings [35]. The results will be discussed in Section 4.

Passive Network Conditions: All 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 the player experience (QoE) is affected by the cloud gaming platforms' adaptations of their streaming graphics configurations. The results will be discussed in Section 5.

¹ Our cloud gaming PCAP dataset is publicly available at [33].

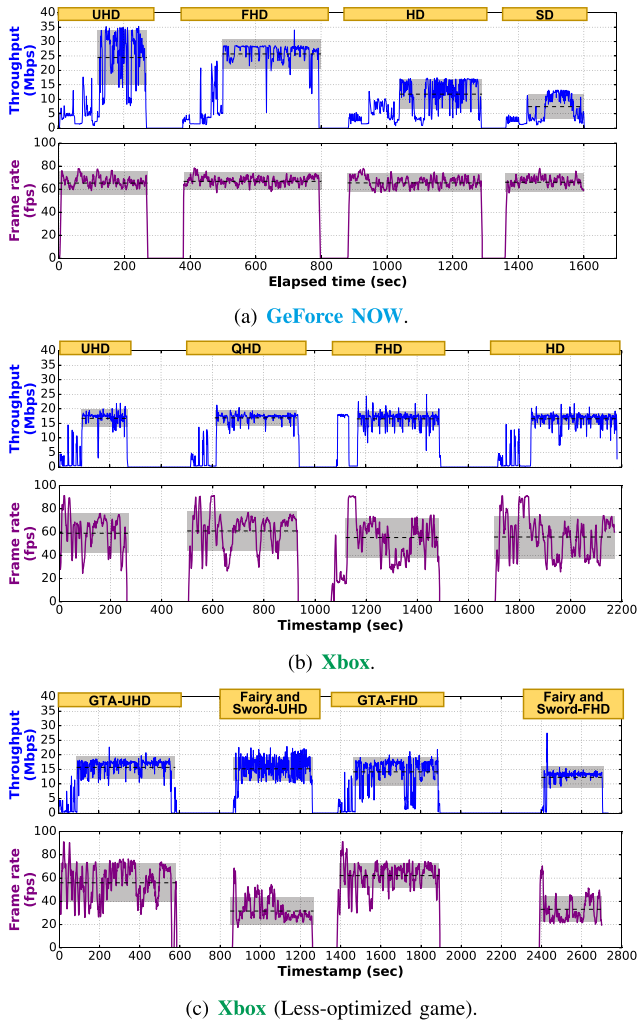


Fig. 2. Time series plots for frame rates and streaming throughput of cloud games with different client resolution settings. The range of timestamps in each plot represents the total duration of each continuously conducted experiment, which can be different across our collected sessions.

4. Analyzing cloud gaming network characteristics across client settings

In this section, through case studies, we profile how cloud games adapt their network streaming characteristics to support different streaming settings (*i.e.*, graphic resolution and frame rate) exactly as configured by the client, under unconstrained network conditions. For conciseness, insights into Sony PS5 Cloud Streaming and Amazon Luna are discussed in our quantitative assessment (Section 6). Key observations are summarized in Section 4.3.

4.1. NVIDIA's GeForce NOW

Fig. 2(a) shows the downstream throughput and gameplay streaming frame rate of First Person Shooting (CS2) gameplay sessions on Windows PC native 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 to either 30 or 60 fps. Therefore, we fix the maximum frame rate option to 60fps in Fig. 2(a). Our other experiments with frame rate set to 30fps reached consistent conclusions, thus, are not shown here for brevity. To measure the actual frame rate received by the client, we use

the number of frame markers per second in the streaming RTP flows [36] as the ground-truth indicator, as introduced in Lyu et al. [5].

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. The dashed lines indicate the average throughput and frame rate during active gameplay, with the gray bands representing the dispersion of the per-second statistics accumulated over at least 200 data points for each experiment.

In addition to the CS2 cloud game sessions with a 60fps frame rate, we also investigated into 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 the distribution of streaming throughput among various game titles and genres. Similar insights are obtained for other supported user setups, including PC browsers and native mobile applications.

4.2. Microsoft's Xbox cloud gaming

Xbox does not allow users to choose the game streaming frame rate. In Fig. 2(b), we 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 (at 18Mbps) for the four resolution bands on the Xbox hardware console. This suggests that Xbox Cloud Gaming always attempts to stream at the highest resolution, and only performs downscaling at the client side to suit the user setting and/or fit the display device.

The second observation we made for Xbox Cloud Gaming platform is that, the streaming throughput reduces for game titles that are inherently offered with low frame rate by the cloud servers, possibly due to high graphics 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 [37] for graphics 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.

4.3. 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 (ISP) 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



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

Cloud Gaming maintains a fixed frame rate for various game titles based on their well- or less-optimized graphics rendering costs. Frequent instances of low frame rate (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 (mistakenly) attributing the network QoS provided by their ISPs as the root cause.

Having analyzed how cloud gaming platforms actively adapt their network streaming characteristics (especially streaming bitrate) to the client graphics settings, we now move on to investigate their adaptability to passive network conditions, namely how they reactively adapt the streamed user experience under constrained network conditions.

5. Analyzing cloud gaming adaptability to network conditions

In this section, through representative case studies, we profile how the perceived cloud gaming user experience (QoE) changes as GeForce NOW and Xbox Cloud Gaming adapt their streaming quality to constrained network conditions (Section 5.1), including limited bandwidth (Section 5.2), increased latency (Section 5.3) and increased packet drop rate (Section 5.4).

5.1. 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 quality 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 game scenes, is named **unplayable resolution**. As shown in Fig. 3(b), the game scene is fragmented into multiple blocks that are inconsistently synchronized with each other, as opposed to just being blurry in the case of "reduced resolution". Users find it challenging to proceed with gameplay due to such a distorted graphic experience.

In addition to the picture quality (i.e., graphic resolution and bitrate), the cloud gaming user experience is also directly dependent on the streaming frame rate. A high value (e.g., near 60fps or above) 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 Section 4, the streaming frame

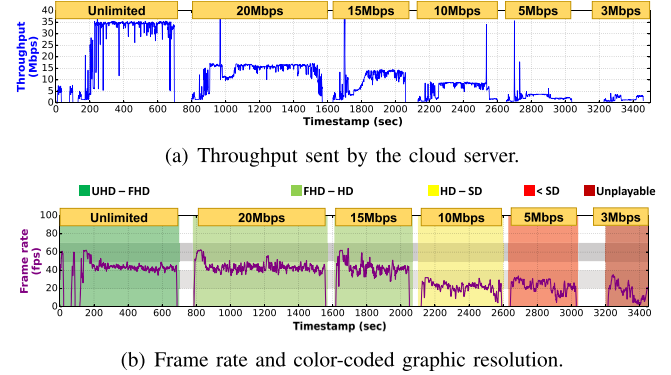


Fig. 4. Frame rates and streaming throughput of GeForce NOW cloud games under limited bandwidths.

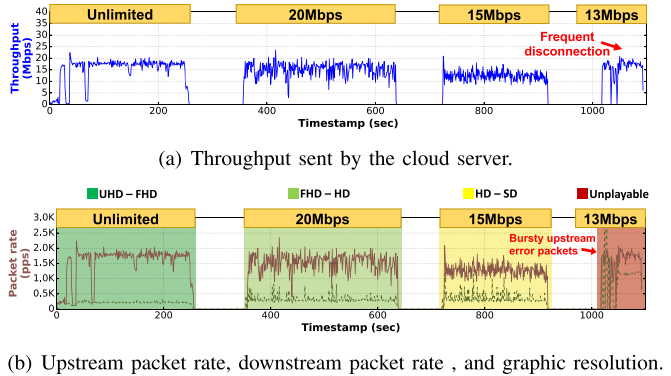


Fig. 5. Streaming throughput, upstream/downstream packet rate of Xbox cloud games under limited bandwidths.

rate can be directly and quantitatively measured by counting the frame markers in the respective RTP flows for our analysis.

5.2. Bandwidth

Clients who subscribe to broadband network services often have an expected bandwidth as part of their service level agreements. However, in practice, the actual 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 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 Section 3. 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

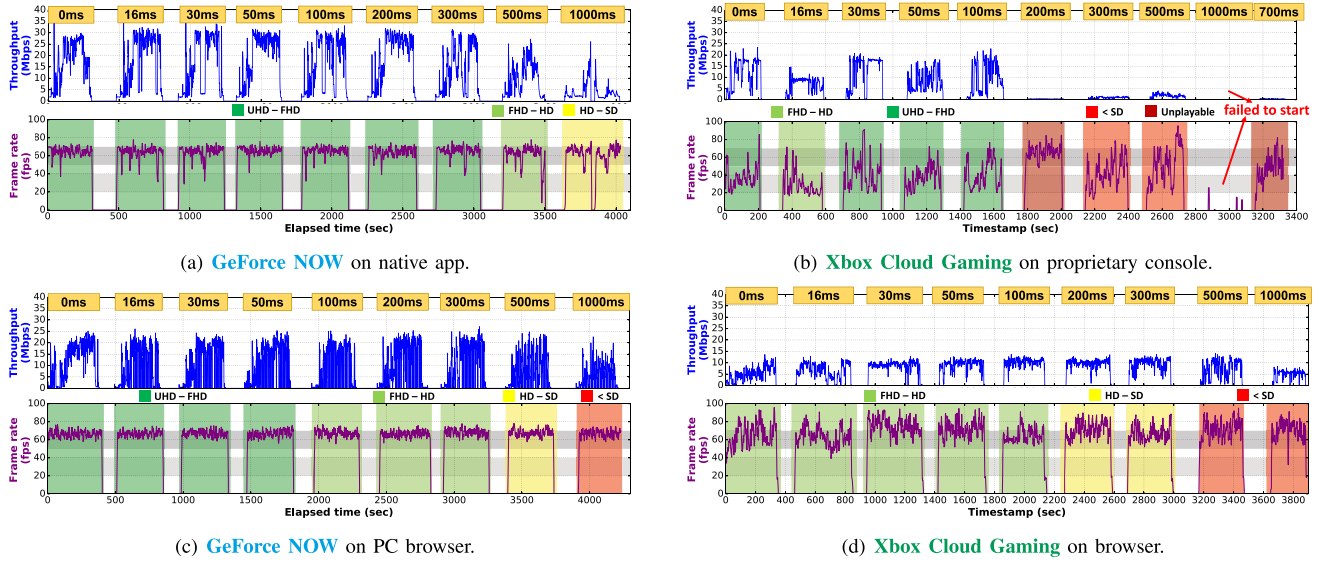


Fig. 6. Streaming throughput, frame rate and graphic quality of cloud games with increased latency.

fragmented graphics on each cloud gaming platform. 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 post traffic shaping (indicated by the green arrow in Fig. 1), we verify that the downstream traffic throughput is consistently capped at the configured levels.

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 pre-shaping vantage point (indicated by the red arrow in Fig. 1). Fig. 2(b) shows the streaming frame rate received by the client device, as measured at the post-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 available network bandwidth. With the default settings on our Windows PC connected to a UHD display, the servers sent a maximum of 36Mbps traffic to the client under unlimited available bandwidth. The bitrate from cloud servers dropped to 18Mbps when a “20Mbps” bandwidth limit is imposed on the client device. Similar observations can be made from Fig. 2(a) for the other bandwidth limits.

As shown in Fig. 2(b), to adapt the video bitrate 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 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 rate for such unplayable scenarios are also unstable and consistently below 30fps.

Xbox Cloud Gaming: The Xbox Cloud Gaming platform exhibits a significantly different behavior when the available bandwidth of the client is insufficient to support the default client QoE settings. We reduce the bandwidth limit from unlimited (1Gbps) to 20Mbps, 15Mbps,

and 13Mbps. The downstream throughput sent by the cloud gaming servers is shown in Fig. 5(a), and the upstream/downstream packet rate measured at server/client sides is shown in Fig. 5(b).

Unlike GeForce NOW cloud servers, which adapt their streaming bitrate 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 bitrate from the server drops when bandwidth limits are introduced, its pattern becomes quite unstable. Notably, the streaming throughput from the server before the traffic shaping 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 bitrate 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 the 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 fragmented.

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, which is aligned with previous findings [4]. 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 support dynamic adjustment of the streaming bitrate from the server side based on client-side bandwidth limitation, sending the same amount and quality of gaming video regardless. Consequently, users perceive unplayable graphic quality and/or frequent disconnection when the expected client bandwidth cannot be achieved.

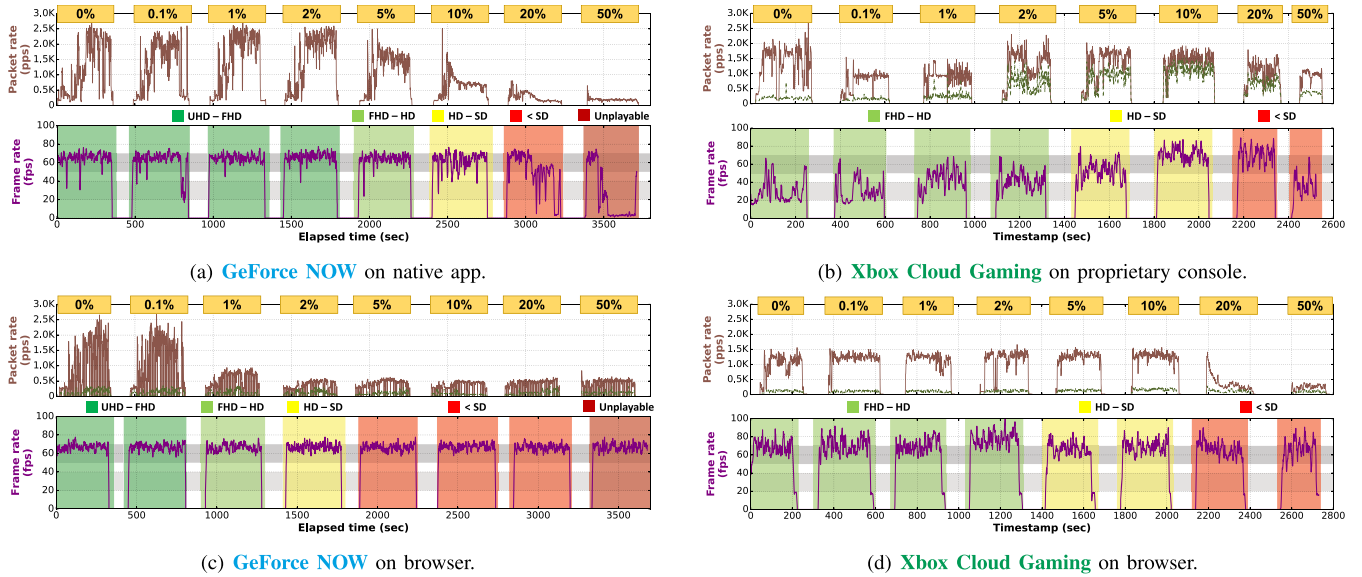


Fig. 7. Packet rate, frame rate and graphic quality of cloud games with increased packet drop rates.

5.3. Latency

Latency is another critical network QoS metric that can directly impact the user experience of networked applications. We now investigate how the streamed user experience are affected as cloud games served by GeForce NOW and Xbox adapt 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.

CeruleanGeForce 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 16ms to 30ms, 50ms, 100ms, 200ms, 300ms, 500ms until 1000ms, when we experience frequent disconnections.

The streaming bitrate sent by the sever is visually presented in Fig. 5(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 bitrate (less than 30Mbps) as it can stream more efficiently (*i.e.*, demanding fewer bytes on the wire), possibly due to less aggressive Forward Error Correction (FEC) mechanisms. 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 adaptation also results in lower streaming bitrate, *i.e.*, 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 latency in terms of maintaining graphic resolution while having smooth frame delivery. In the shown example, the resolution starts dropping right after 100ms, much earlier than the 500ms threshold via native applications.

Compared to limited bandwidth, as discussed in Section 5.2, 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 experience less responsive in-game motions and actions, until being disconnected by the server.

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 devices).

On Xbox’s own proprietary hardware console, the platform tends to adjust both graphic resolution and frame rate when extra latency is introduced (Fig. 5(b)). For instance, with an extra latency of 16ms, we observe both resolution drop from UHD to FHD and frame rate drop from 60fps to 30fps. Interestingly, when we increase the extra latency to 30ms, the frame rate returns to 60fps while the resolution also returns back to UHD. However, the frame rate drops back again 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 “16ms”, “50ms” and “100ms” cases, it hovers around 10Mbps. It is worth noting that the fluctuation in throughput in the “50ms” and “100ms” cases is also caused by error/disordered frames (*e.g.*, due to packets dropped upon queue overflow), as discussed in Section 5.2.

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 failure of game initialization.

For the gameplay sessions played via PC browsers or mobile browsers/apps, 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 is introduced (*e.g.*, no resolution change until 200ms and no frame rate change even with 1000ms latency in our example case shown in Fig. 5(d)). With an ideal latency of less than 5ms (the leftmost case in Fig. 5(d)), the streaming bitrate sent from the cloud server is the most optimized compared to slightly higher latency bands, as it requires the lowest amount of bandwidth while still maintaining the same high resolution and frame rate.

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 become very unresponsive and eventually the cloud gaming session becomes 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, the reaction, *i.e.*, dropping graphic resolution for smooth frame streaming, to increased latency is more optimized for cloud games on native applications.

Xbox Cloud Gaming demonstrates different responses to increased latency for cloud games played on the Xbox hardware console and PC/mobile devices. 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 PC/mobile sessions start to adapt their graphic resolution only at a higher latency level (*e.g.*, 500ms).

5.4. Packet drop rate

In addition to bandwidth limit and latency, packet drop rate is another important network QoS metric that ISP often optimize against. The impact of packet drop rate on online gaming [38] and video streaming [39] 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).

GeForce NOW: GeForce NOW tends to actively reduce the graphic resolution of gameplay sessions that are experiencing observable (*i.e.*, medium or high) packet losses, while the streaming frame rate is kept at a relatively optimal level for the smoothness of user motions and dynamic game scenes. We now discuss a representative experimental result for CS2 gameplay sessions via native GeForce NOW PC applications, which is visually depicted in Fig. 7(a) as time-series plots.

Gameplay sessions via GeForce NOW native application behave nearly optimally when experiencing packet loss. From Fig. 7(a), we can see that the graphic resolution gradually drops from FHD to HD and finally SD as the packet drop rate increases from 1% to 50% in five steps. The frame rate always stays at a relatively high level (around 60fps) to guarantee smoothness in user motions and dynamic game scenes, even when the loss rate reaches 50%. Considering that prior works, such as Wu et al. [25], have proposed advanced techniques for image and user action prediction in cloud gaming, it is plausible that the native application might incorporate similar techniques for a highly optimized gameplay experience even in degraded network QoS conditions.

Unlike native GeForce NOW application that maintains good smoothness by optimally reducing graphic resolution even with significant packet drops, gameplay sessions via mobile/PC browsers seem to follow a more brute-force approach. This involves a sudden drop in graphic resolution even with a slight level of packet loss, as clearly observed from Fig. 7(c) from streaming throughput, packet rate and user experience metrics of representative sessions via PC browsers.

For the first two very low levels of packet drop rate (*i.e.*, 0% and 0.1%) shown as the two leftmost blocks in Fig. 7(c), we do not observe any difference in graphic resolution or frame rate. The streaming bitrate stays at nearly 20Mbps with a packet rate of 2Kpps. As the packet drop rate increases to even a small amount of 1% and 2%, the graphic resolution is radically adjusted to a very low quality (*e.g.*, HD and SD on a UHD display), as visually shown in Fig. 7(c). Meanwhile, the frame rate remains at a decent level (*e.g.*, around 60fps) for a smooth streaming experience. Surprisingly, the perceived smoothness of the streaming experience does not change significantly from 1% to 10% packet drop rates. Therefore, we believe that a significant resolution drop at low packet drop levels (*e.g.*, 1%) is an excessive approach more suitable for higher drop rates. When the packet drop rate reaches 20%, we begin to experience frequent stutters in game scenes and noticeable “teleportation” of our in-game characters. Eventually, the gameplay becomes entirely unplayable at a 50% packet drop rate, requiring a long wait for even a single change in the video frame.

Xbox Cloud Gaming: From our experiments, Xbox Cloud Gaming gameplay sessions accessed via all supported PC, mobile and hardware console do not react optimally to packet drops. The game scenes via

PC and mobile browsers become fragmented and discontinued (such as shown in Fig. 3(c)) even if the correctly displayed segments are at good resolution when the packet drop rate reaches 5% or higher, as shown in the representative sessions in Fig. 7(d). In addition, Xbox Cloud Gaming does not adapt its streaming bitrate or packet rate until a very high packet drop rate is applied. Therefore, the streaming bitrate and packet rate from the cloud server remains at the highest level (*i.e.*, when no packet loss is introduced) until a 20% packet drop rate is introduced when the user experience becomes almost unplayable. As visually shown in Fig. 7(d), the packet rate drops from 1400pps to 400pps for the two high drop rates.

The gameplay sessions via Xbox hardware console similarly underperform compared to those via PC or mobile browsers, as illustrated in the example in Fig. 7(b). However, given that the hardware console sessions do actively acknowledge error and lost packets to the cloud server, as shown by the relatively higher upstream throughput for 2%, 5%, and 10% packet loss rates in Fig. 7(b), the streaming bitrate and packet rate show arbitrary changes. Similarly, the frame rate shows random, unexpected increase under high packet drop rate (*i.e.*, 5% to 20%) due to retransmitted video frames that failed to be delivered to the client. We could not find a clear pattern from our repeated experiments.

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 with a small level of packet loss (*e.g.*, 2%), whereas Xbox sessions can quickly become unplayable (*e.g.*, having fragmented graphics) even with a small drop rate. Therefore, ISPs will need to guarantee minimized packet drop rates for Xbox Cloud Gaming players to achieve an equivalent experience from GeForce NOW.

As for GeForce NOW, its native applications on mobile and PC devices adapt their graphic resolutions to different level of packet drop rates more precisely than those via generic third-party browsers, which takes a radical drop of resolution when the drop rate exceeds 1%. Therefore, by identifying the software agent a GeForce NOW cloud gaming user is using for a particular session that may experience packet drops, ISPs can precisely manage the packet drop rate to achieve desirable user satisfaction.

6. Systematic schema for assessing cloud gaming network adaptability

Through representative case studies, we have discussed our qualitative understanding of how cloud gaming platforms actively adapt network demands (*i.e.*, throughput) for different levels of active streaming configurations, and passively adapt their streaming behaviors (*e.g.*, frame rate or graphic resolution that impact subjective user experience) under degraded network conditions. Such heuristic conclusions drawn from manual analysis require extensive labor efforts and cannot be easily expanded for all the traffic traces we have collected in Table 2, a necessary step for converged quantitative conclusions. They also cannot be replicated for other emerging platforms or used to provide periodic assessment results for the networking and game development communities. Therefore, in this section, we design a systematic schema that quantitatively assess the network adaptability of a cloud gaming platform. Our schema (Section 6.1) uses performance scores (Section 6.2) to indicate the adaptability for client settings impacting network throughput; and network conditions impacting subjective user experience, which are further aggregated into summarized scores using customized weights defined by the assessors (*e.g.*, network operators) (Section 6.3).

6.1. Assessment schema and categories of scores

We define our schema that contains assessment scores covering all variables as discussed in Section 4 that could be considered as adaptability for both client configurations and network conditions.

Table 3
Five quantitative levels of cloud gaming user experience with ACR scale.

	Excellent	Good	Fair	Poor	Bad
Graphic Quality	High graphic quality	Reduced graphic quality with minor artifacts	Noticeably reduced visual quality with occasional artifacts, blurry during fast motion	Significant reduction in visual quality, persistent blur/pixelation, distorted textures	Unrecognizable (pixelated) visuals, persistent freezing
Responsiveness	No noticeable input lag	Slightly input lag, minimal impact on gameplay	Noticeable input lag with impact on reaction time	Laggy controls and require anticipation to compensate	Unresponsive controls, significant input lag
Streaming Smoothness	No noticeable stuttering or frame drops	Slight frame drops or stutters, minimal impact on gameplay	Rare frame drops and stutters, noticeable impact on gameplay	Occasional stuttering and frame drops, still playable	Frequent game freezes and significant frame drops, not playable

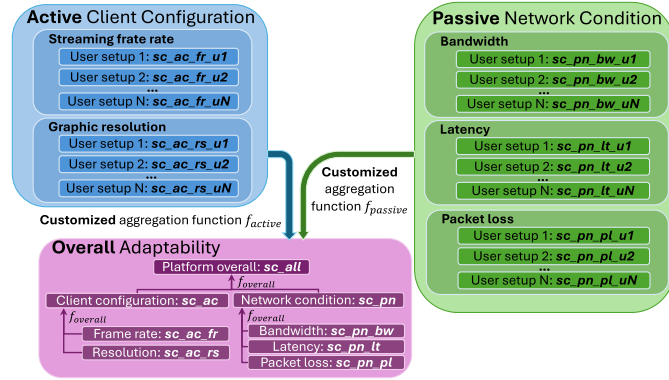


Fig. 8. Assessment schema of cloud gaming platform network adaptability.

As visually shown in Fig. 8, the scores are grouped into three categories (*i.e.*, boxes), namely active client configuration, passive network condition and overall adaptability.

Two basic categories: Within the category of active client configuration, each score indicates the network adaptability with respect to either streaming frame rate or graphic resolution of the assessed cloud gaming platform accessed through a certain user setup such as mobile browser or PC APP. Similarly, within the green box in Fig. 8 for passive network condition, each score (*e.g.*, $sc_{pn_bw_u1}$ and $sc_{pn_lt_u1}$) is for the adaptability of cloud games in terms of either available network bandwidth, latency or packet loss through one type of user setup.

One category for the overall adaptability: Given that the scores can be large in quantity, thus, become difficult for a user of the assessment schema (*e.g.*, network operators who want to estimate the gaming experience of their clients) to evaluate the performance of a cloud gaming platform in an abstract view. We therefore design the third category, shown as the pink block in Fig. 8, that hosts eight scores aggregated from those in both client configuration and network condition categories. The eight scores in this abstract category are organized in three levels of hierarchy, starting from the basic variables (*e.g.*, streaming frame rate sc_{ac_fr} and bandwidth sc_{pn_bw}), to their upper categories (*i.e.*, either client configuration sc_{ac} or network condition sc_{pn}), and the one overall score sc_{all} for the assessed cloud gaming platform.

6.2. Generating and aggregating scores

After introducing the three blocks in our assessment schema, we now explain how the scores are generated in each block.

Client configuration: For the fine-grained block of client configuration, each score represents the ability of the assessed cloud gaming platform to adapt its network volumetric metrics “ net_vol ” (*i.e.*, throughput or packet rate) to various levels of a certain considered variable “ var ” (*e.g.*, frame rate). Assuming that we have a sequence of considered variable values $[var_1, var_2, \dots, var_n]$ and their corresponding network volumetric metric values $[net_vol_1, net_vol_2, \dots, net_vol_n]$, there are many statistical methods, from linear regression to nonlinear models (*e.g.*, loga-

rithmic), that take the two input sequences (*i.e.*, net_vol and \vec{var}) to come up with their correlation. The correlation, which indicates the adaptability of a cloud gaming platform for a certain variable, can be formally expressed as:

$$corr = f_{adapt}(net_vol, \vec{var}) \quad (1)$$

As will be discussed in Section 7, the sequences of throughput and other five variables including streaming frame rate, graphic resolution, available bandwidth, latency and packet loss are produced from traffic traces of labeled cloud gaming sessions. While we use linear regression as the correlation function in the following lab assessment as a demonstrative example, a systematic evaluation may be performed by an operator to select the most suitable correlation function for its specific use cases, which is beyond the scope of this study.

Network condition: Similar methods are used to compute adaptability scores for the fine-grained block of network condition, whereas the correlations are computed between the **subjective user-perceived streaming quality** as rated according to Table 3 and the corresponding network conditions (*e.g.*, latency). Also, a smaller change in video quality indicate a better platform adaptability. Therefore, to make the value of score consistent with human intuition that larger scores mean better results, we take the inversion of normalized correlations as the final adaptability score.

Following the common practice in user experience research [40], we choose Absolute Category Rating (ACR) to map subjective user experience to five quantitative levels by collectively considering the graphic quality, responsiveness and streaming smoothness, as detailed in Table 3. As found in [7], subjective user experience is not solely determined by the objective QoE metrics like streaming frame rate, as certain game titles and genres may inherently demand higher or lower frame rate and graphic resolution to support a good user experience than others. For example, a good experience for card games like Hearthstone may only require a streaming frame rate of 30fps, whereas the same objective frame rate metric leads to poor or even bad experience for competitive shooting games like CS2. Therefore, we directly capture user-perceived experience via subjectively rated ACR scores while playing the respective games, rather than inferring from objective, game-agnostic QoE metrics. This provides a uniform representation of the subjective experience across game titles and genres for a consistent assessment of a platform’s network adaptability, while inherently accounting for their unequal demands.

In gameplay, poorer performance of one metric (*e.g.*, graphic resolution) always lead to a worse overall user experience. Therefore, we use the lowest ACR score across all three aspects to represent the overall experience of an entire cloud gaming session. In our assessment results as will be discussed in Section 7.2, we statistically model the changes of user experience scores as a certain network condition metric degrades to various levels. For a more fine-grained and holistic representation of the user experience, a composite or weighted score covering all three aspects may be used, which is beyond the scope of this work.

Overall adaptability: Presenting only individual scores can make executive-level comparison difficult. Therefore, as designed in our

schema (Section 6.1), we aggregate the assessment scores per variable via each user setup into three levels of overall scores through their respective aggregation functions. The first aggregation function f_{active} takes all scores in the active client configuration block to produce two overall scores for frame rate sc_{ac_fr} and resolution sc_{ac_rs} , respectively. The second function $f_{passive}$ produces three scores for bandwidth sc_{pn_bw} , latency sc_{pn_lt} and packet loss sc_{pn_pl} from the passive network condition block. As shown in Fig. 8, the third function $f_{overall}$ in the overall adaptability block further aggregates the scores to higher hierarchies including client configuration sc_{ac} , network condition sc_{pn} and ultimately one overall adaptability score sc_{all} . The mechanisms of our aggregation functions are discussed next.

6.3. Customized aggregation functions

We now discuss how our three aggregation functions (i.e., f_{active} , $f_{passive}$ and $f_{overall}$) progressively concise the number of adaptability scores with customized assessment priorities defined by the users.

Aggregation function for client configuration: First, the aggregation function f_{active} summarizes all scores of streaming frame rate each for a unique user setup type into one score sc_{ac_fr} . The original scores can be represented as an array S_{ac_fr} defined as:

$$S_{ac_fr} = [sc_{ac_fr_u1} \quad sc_{ac_fr_u2} \quad \dots \quad sc_{ac_fr_uN}] \quad (2)$$

Users can specify their priorities within the range between 0 (lowest priority) and 1 (highest priority) on each user setup type for their own assessment needs. For example, a mobile telecommunications operator may give higher priority on mobile-related user setups like mobile APP and mobile browser so that the results are not biased by user setups that are not popular among its served users (e.g., PC or gaming console). The priorities assigned for all user setups of a given cloud gaming platform can be represented as an array W_{ac_fr} defined below:

$$W_{ac_fr} = [wt_{ac_fr_u1} \quad wt_{ac_fr_u2} \quad \dots \quad wt_{ac_fr_uN}] \quad (3)$$

Therefore, the weighted average score sc_{ac_fr} by combining S_{ac_fr} and W_{ac_fr} can be expressed as:

$$sc_{ac_fr} = \frac{S_{ac_fr} W_{ac_fr}^T}{\sum_{i=1}^N W_{ac_fr_ui}} \quad (4)$$

We use the same approach in f_{active} to compute the summarized score sc_{ac_rs} for graphic resolution.

Aggregation function for network condition: The adaptability scores aggregated by $f_{passive}$ for the three variables in passive network condition, i.e., bandwidth sc_{pn_bw} , latency sc_{pn_lt} and packet loss sc_{pn_pl} are computed by the same method as used in client configuration f_{active} . Users can provide their customized priorities for each user setup type, or use default uniform values for all original scores.

Aggregation function for overall adaptability: Unlike the previous two aggregation functions that summarize scores from two detailed blocks into the overall adaptability block in Fig. 8, the third aggregation function $f_{overall}$ is applied within the overall adaptability block from the bottom to top hierarchy.

It first aggregates the two scores for frame rate sc_{ac_fr} and resolution sc_{ac_rs} into one score for active client configuration sc_{ac} . Similar to other aggregation functions, each original score is assigned a customized priority (in the array W_{sc_ac}) by the users, as shown in the formula below:

$$sc_{ac} = \frac{S_{sc_ac} W_{sc_ac}^T}{\sum W_{sc_ac}}, \text{ where :} \quad (5)$$

$$S_{sc_ac} = [sc_{ac_fr} \quad sc_{ac_rs}]$$

$$W_{sc_ac} = [wt_{ac_fr} \quad wt_{ac_rs}]$$




Fig. 9. Processing workflow for assessing the adaptability of a certain cloud gaming platform.

The function also aggregate the three scores for passive network conditions into one score sc_{pn} , as shown below:

$$sc_{pn} = \frac{S_{sc_{pn}} W_{sc_{pn}}^T}{\sum W_{sc_{pn}}}, \text{ where :} \quad (6)$$

$$S_{sc_{pn}} = [sc_{pn_bw} \quad sc_{pn_lt} \quad sc_{pn_pl}]$$

$$W_{sc_{pn}} = [wt_{pn_bw} \quad wt_{pn_lt} \quad wt_{pn_pl}]$$

As the last step shown in the formula below, $f_{overall}$ further aggregates the two scores for client setting and network condition in the middle hierarchy into one score for overall adaptability sc_{all} with user-specified priorities.

$$sc_{all} = \frac{S_{sc_{all}} W_{sc_{all}}^T}{\sum W_{sc_{all}}}, \text{ where :} \quad (7)$$

$$S_{sc_{all}} = [sc_{ac} \quad sc_{pn}]$$

$$W_{sc_{all}} = [wt_{ac} \quad wt_{pn}]$$

6.4. Discussion on schema extensibility

The schema is independent of the quantity and combination of user setups in the assessed dataset, as the five fine-grained adaptability scores (e.g., streaming frame rate) calculated for a new user setup (e.g., Linux browser) will be added to their respective pillar along with all other user setups (e.g., as “ $sc_{ac_fr_linux_browser}$ ”), and then aggregated to various levels of categories. Similarly, the schema can be easily extended to include more assessment aspects in both active client configuration and passive network condition pillars. Three changes will be made in the schema for a new assessment aspect (e.g., “jitter” representing the variation of latency). First, a sub-block containing detailed scores of the newly added assessment aspect per user setup will be added into the respective pillar (e.g., network condition pillar for the network jitter). Second, a summarized score for this aspect will be added into the bottom hierarchy in the overall adaptability block, such as “jitter: sc_{pn_jit} ”. Third, score priorities associated with this assessment aspect will be assigned by the user accordingly. In the “jitter” example, the priorities of jitter scores per user setup and the priority of the overall jitter score will be assigned. Other cloud gaming platforms can also be assessed the same way as those four discussed in this paper, given their labeled PCAP files for all the fine-grained assessment metrics across various user setups.

7. Assessing network adaptability of four popular cloud gaming platforms

In this section, we systematically assess the network adaptability of four popular cloud gaming platforms using our schema designed in Section 6. Our assessment workflow that takes labeled ground-truth traffic capture (PCAP) files and user-specified priorities to generate scores is explained in Section 7.1. The assessment results are discussed in Section 7.2.

Table 4

Breakdown assessment scores of active client configuration and passive network condition for GeForce NOW and Xbox Cloud Gaming platforms. The cells under PC categories are for Windows/macOS and those under mobile categories are for Android/iOS. Adaptability scores higher than **85**/100 or lower than **50**/100 are highlighted.

		GeForce NOW				Xbox Cloud Gaming		
		PC APP	PC browser	Android APP	mobile browser	Xbox console	PC browser	mobile browser
Client configuration	Frame rate	90.8 /83.4	78.0/82.9	52.9	99.7 / 93.8	0	0	0
	Resolution	83.3/83.7	81.7/ 99.0	3.86	89.1 /60.9	0	0	0
Network condition	Bandwidth	88.9 / 86.3	89.9 / 97.9	89.9	89.9 /69.2	95.5	81.4/73.3	95.5 /77.6
	Latency	80.7/77.9	83.4/68.6	83.5	83.5/73.8	76.4	87.0 /65.9	77.9/ 92.0
	Packet loss	98.1 / 93.9	92.9 /83.2	98.1	93.6 /75.1	94.5	83.1/75.1	80.7/75.1

Table 5

Breakdown assessment scores of active client configuration and passive network condition for Amazon Luna and PS5 Cloud Streaming platforms. Adaptability scores higher than **blue85**/100 or lower than **50**/100 are highlighted.

		Amazon Luna		PS5 Cloud Streaming	
		PC browser	mobile browser	PlayStation console	PC APP
Client configuration	Frame rate	55.4	0	0	0
	Resolution	98.7	93.3	25.6	49.2
Network condition	Bandwidth	92.5	85.4	93.6	91.9
	Latency	76.4	91.5	86.2	97.9
	Packet loss	76.5	72.8	68.9	81.2

7.1. Workflow

We implement an assessment pipeline that automatically generates scores in our schema by taking cloud gaming packet capture files as abstracted in Fig. 9. The pipeline requires three types of input files, including packet capture (PCAP) files of cloud game sessions served by the assessed platform, ground-truth labels of the provided PCAP files in YAML metadata files, and user-specified score priorities for aggregation purposes in a YAML configuration files.

Starting from the left side of Fig. 9, the **packet processing module** receives PCAP files and the corresponding metadata files in YAML format. Noting that the YAML files contain the values of enforced variables such as frame rate or available bandwidth, timestamp ranges of different values of an enforced variable, and perceived user experience ratings that cannot be directly derived from the PCAP files. Table statistics in CSV format is generated per predefined time interval (e.g., one second in our implementation) for three metrics that are directly computed from the PCAP files (i.e., median throughput, median packet rate, median frame rate) and the enforced variable value. As an offline process after data collection, this module can be sufficiently implemented with common traffic processing tools like tshark, without stringent performance constraints.

The following **assessment module** consumes all table statistics for each adaptability variable (e.g., resolution) per user setup to generate corresponding fine-grained scores in our adaptability assessment schema (e.g., *sc_ac_rs_ul* in Fig. 8). For comparison purposes, we use the linearly normalized correlation as the adaptability score for a given variable with a certain type of user setup, e.g., available bandwidth with mobile browser *sc_pn_bw_ul*.

To quantify the adaptability to active client configuration options, namely resolution (e.g., UHD, FHD) and frame rate (e.g., 30fps, 60fps), which are discrete categorical values, we represent resolution values by the numbers of pixels per frame (e.g., the product of 3840×2160 for UHD), represent frame rate by the number of frames per second (e.g., 60), and use throughput as the adaptability metric *net_vol*. Previous findings show that the bitrate of video streams linearly scales with the pixel dimension [41] and frame rate [42] within a practical range (e.g., 30–120Hz), with small variations introduced by compression in

the encoding process. We then apply linear regression (as f_{adapt}) to capture their correlations. Specifically, we measure how well the changing ratio of the adaptability metrics (i.e., throughput) match that of the variable values (i.e., resolution and frame rate) using the Pearson correlation coefficient. For example, the active adaptability to client resolution settings for GFN on Mac native app is calculated by performing a linear regression between the resolution variables from UHD (3840×2160) to HD (1280×720) and the measured network throughput metrics, and taking the resulting Pearson correlation coefficient as the final score.

Assessment for network conditions variables (i.e., bandwidth, latency and packet loss) follow the same procedures, with the adaptability metric being the subjective ACR score rather than throughput. Additionally, for latency and packet loss whose values are negatively correlated with the QoE, we take the absolute value of the Pearson correlation coefficient so that the score reflects the level of adaptability, regardless of the direction of the correlation.

With the user-specified score priorities, the fine-grained scores under the two categories, i.e., active client configuration (blue block in Fig. 8) and passive network condition (green block in Fig. 8), are further processed by the **assessment score aggregation module** that applies aggregation functions (introduced in Section 6.3) to generate overall adaptability scores for a certain cloud gaming platform at three hierarchical levels.

7.2. Assessment results

We have collected packet capture (PCAP) files for a total of 439 cloud gameplay sessions on the four popular platforms. The labeled PCAP files cover all available variations of supported active client settings (at the time when we conduct our experiment) and network QoS constraints. For NVIDIA GFN and Xbox Cloud Gaming platforms that are available in our geography (i.e., Australia), the traffic traces are collected within our lab using the setup discussed in Section 3. For Amazon Luna and PS5 Cloud Streaming which are not offered in Australia, the labeled data are collected by a volunteer (acknowledged in §VIII) recruited from the U.S. with a Linux-based home router pre-configured by the authors for adjusting network QoS conditions and capturing traffic traces. To eliminate the potential biases on network demands and user experience

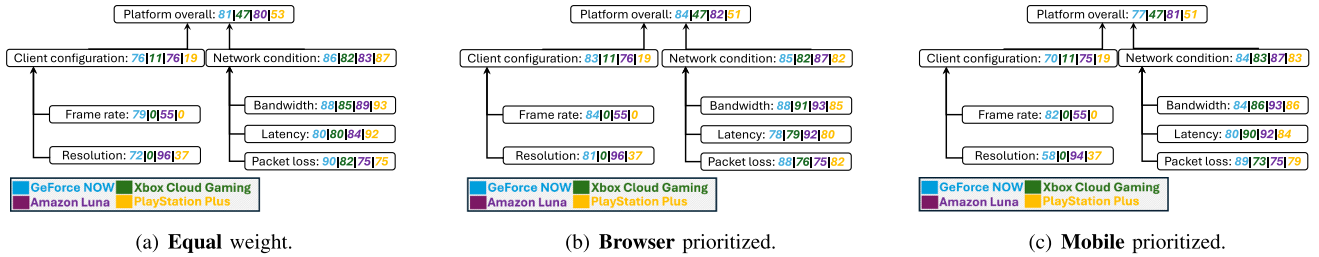


Fig. 10. Assessors can customize weights when assessing the adaptability of cloud gaming platforms to network conditions and client configurations, such as (a) **equally** consider all user setup types and metrics, (b) **prioritize browser** metrics, or (c) **prioritized sessions on mobile** devices.

expectations imposed by different games, we select the same game title (*i.e.*, Fortnite) throughout all our cloud gameplay sessions, which is available on all four studied services and user platforms. All trace files along with their ground-truth labels are fed into our automatic process (Fig. 9) to generate assessment results. We acknowledge that the results presented here are only representative of our region and our specific network environments. For network operators planning to periodically reassess updated cloud gaming platform behaviors and/or to make assessment under their specific network environments such as different geographical locations or certain access types (*e.g.*, cellular mobile), they may recollect a bespoke set of labeled data through dedicated data engineering teams [43,44], or extend our dataset through traffic augmentation [45–47], which can then be readily fed into our assessment pipeline to generate results.

Fine-grained adaptability scores: Assessors can directly refer to the detailed assessment score for a specific metric in client configuration or network condition category for a certain user setup type. We show the detailed scores for GeForce NOW and Xbox Cloud Gaming in Table 4; and Amazon Luna and PS5 Cloud Streaming in Table 5. Aligned with our empirical findings discussed in Sections 4 and 5, platforms showing a perceptibly good adaptability to certain variables (*e.g.*, GFN native PC app to bandwidth changes) are assigned high scores (*e.g.*, 86.3/100) by our assessment module, whereas platforms demonstrating poor adaptability such as GFN Android app to client resolution settings are assigned very low scores (*e.g.*, 3.86/100). Therefore, for illustrative purposes, high scores over 85/100 and low scores below 50/100 are highlighted by blue and red color in the tables, respectively.

From Table 4 we can see that **GeForce NOW** exhibits good adaptability under constrained network conditions with all assessment scores above 75/100 on all user setups. The game streaming experience of GFN is exceptionally good when the bandwidth is limited for PC browser sessions and the packet loss is relatively high for PC APP and Android APP sessions, with their respective assessment scores above 90/100, possibly due to more effective FEC mechanisms used to compensate such poor network conditions. The GFN’s network adaptability on active client configurations including frame rate and resolution settings are good on PC APP and PC browser as all scores are above 75/100. However, the scores are quite low for both frame rate and resolution on Android APP, indicating that game streaming sessions even with relatively low client setting levels are with bandwidth demands higher than necessary or expected.

Xbox Cloud Gaming has very good adaptability on its hardware console and Android mobile browsers for ensuring cloud streaming user experience when bandwidth is limited, resulting high scores over 95/100. However, its network adaptability on all client configuration does not exist for any available user setup type as indicated by the 0/100 scores, hinting the need for the platform operator to efficiently minimize the network demands to the actual client setting levels for an optimal bandwidth usage.

From Table 5, it can be concluded that **Amazon Luna** achieves good adaptability for network conditions especially on mobile browsers, with

the scores for bandwidth and packet loss above 85/100. Its network adaptability on user requested graphic resolution is good as well with all scores above 85/100. However, different user requested frame rates have very minor or even no impact on the high bandwidth demands on both available user setup types, resulting in very low scores of 55.4/100 and 0/100.

As for **PS5 Cloud Streaming**, it has good adaptation to ensure user experience under constrained network conditions. However, its network adaptability for both user requested frame rate and graphic resolution is very poor (with all scores less than 50/100) on both available user setup types (PlayStation console and PC APP), indicating a constantly high network demand by cloud gaming sessions regardless of the requested frame rate and graphic resolution.

Adaptability by aggregating weighted scores: In addition to fine-grained adaptability scores for each assessment metric of a certain user setup type, as discussed in Section 6.3, our assessment framework also produces aggregated views of adaptability for specific metrics, client configuration or network condition categories, and the entire platform with customized weights assigned by an assessor for effortless comprehension. As shown in Fig. 10(a) where all metrics are assigned with **equal weights**, we can conclude that GeForce NOW has the best overall platform adaptability with a score of 81, followed by Amazon Luna (80), PS5 Cloud Streaming (53), and Xbox Cloud Gaming (47). By looking at the aggregated scores one level below, we can see that the four platform have similar adaptability for “network condition” with their scores all within a small range between 82 and 87. However, the four platforms are significantly different in “client configuration”. Notably, we see very low adaptability scores of 11 and 19 for Xbox Cloud Gaming and PS5 Cloud Streaming, indicating that their network streaming behaviors have almost no change for varying client configurations. An assessor can further look at scores aggregated for each client configuration or network condition metric for more specialized conclusions.

An assessor may **prioritize** certain assessment aspects compared to others. For example, one may suggest that cloud game streaming may be more favored by casual gamers who use **browsers** for their cloud game sessions than installing native applications, therefore, the platform adaptability for browser sessions is more critical in overall user satisfaction. This prioritization can be directly captured in our assessment pipeline by setting high weights for browser metrics (*e.g.*, those for “PC browser” and “mobile browser” in Tables 4 and 5). We show the results after setting the weights for browser metrics to 0.8 and native APP metrics to 0.2 in Fig. 10(b). The overall adaptability scores of GeForce NOW and Amazon Luna slightly increase from 81 to 84 and from 80 to 82, whereas the overall scores for Xbox and PlayStation remain unchanged or slightly drop compared to the results with equal metric weights in Fig. 10(a).

Another set of example results is shown in Fig. 10(c) where we **prioritize** all metrics for **mobile** user setups by setting their weights to 0.8 and otherwise 0.2. Such assessment scores are useful for 5G network operators whose clients are mostly connected via mobile devices. It can be observed that the adaptability of GeForce NOW and

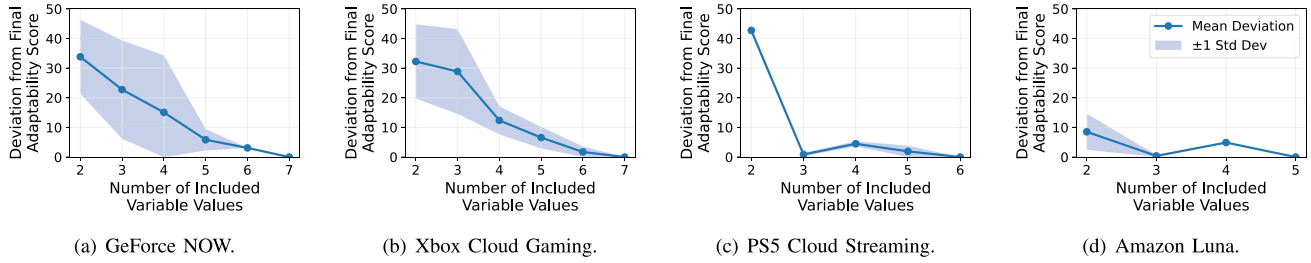


Fig. 11. Deviation from final adaptability scores using a subset of bandwidth variable values.

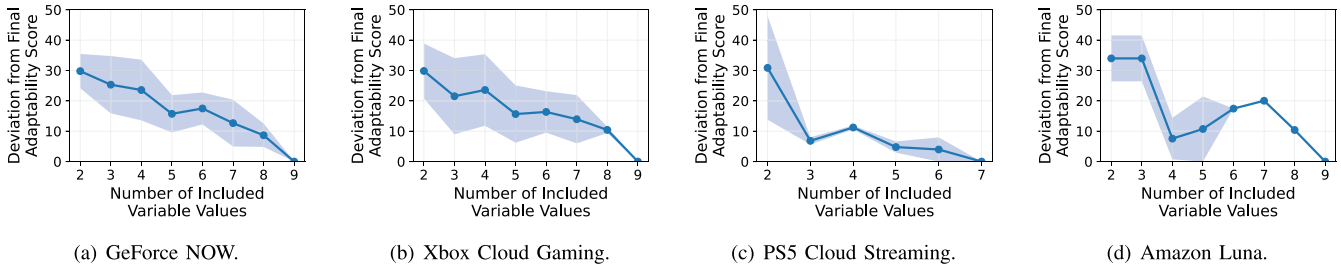


Fig. 12. Deviation from final adaptability scores using a subset of latency variable values.

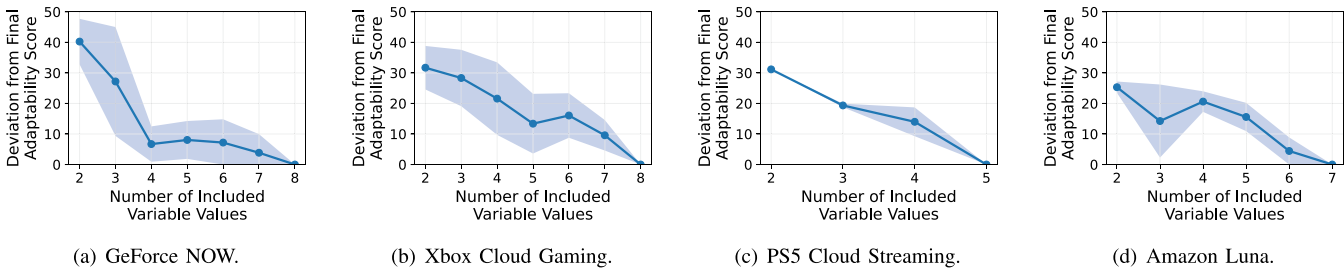


Fig. 13. Deviation from final adaptability scores using a subset of packet drop rate variable values.

PlayStation are slightly weaker on mobile devices compared to others, as their aggregated scores have small drops compared to Fig. 10(a), especially for GeForce Now in client configuration and its graphic resolution metric, hinting more static network bandwidth demands by mobile-based streaming sessions compared to the PC sessions, even when the gamers choose lower streaming graphic settings.

Sensitivity analysis: While the variable values for active client configurations (*i.e.*, resolution and frame rate) are directly determined by the cloud gaming platform and user setup, those for passive network conditions (*i.e.*, bandwidth, latency, packet drop rate) are manually selected. Therefore, we conduct a sensitivity analysis in order to verify that the adaptability results generated by our assessment framework are consistent even if an operator adopts a different set of values for a certain variable (*e.g.*, bandwidth) in their evaluation. Specifically, for each of the network condition variables, we compute the adaptability scores using subsets of our dataset with progressively increased numbers of variable values (from the minimum values required for correlation to using all values) across all available user setups and compare them against the final scores as reported in Tables 4 and 5.

Figs. 11, 12 and 13 depict the mean deviation between the subset score and the final score (as indicated by the dot markers) and the dispersion across user setups (as indicated by the shaded areas). Apparently, with a small subset of variable values (*e.g.*, 2 or 3), the deviations are typically very large, up to 45 points different from the final results. However, the scores quickly converge to the final results as more variable values from each experiment dataset are included in the assessment. We also observe that GFN, xCloud and PS5 almost monotonically converge to the final score as the number of variable values increases, often

bringing the deviation below 15 with only 5 or more variable values. In contrast, Amazon Luna shows an oscillating behavior before eventually converging to the final score, therefore, network operators may need to carefully select a large number of unique variable values when assessing Amazon Luna to reach a stable result. Additionally, our results show that using the currently selected variable values for passive network conditions in our lab dataset, we cover experiments with both excellent/good ACR scores and poor/bad ACR scores for every included cloud gaming platform and user setup, indicating that the selected values provide a representative spread across the entire user experience spectrum.

7.3. Discussion on practical applicability

Our automatic assessment process is designed for seamless integration into the operational workflows of ISPs and game developers, given that standard Data Operations (DataOps) and Machine Learning Operations (MLOps) practices are performed for their specific assessment settings. To satisfy their diverse operational requirements, such as testing across various network environments or different game titles, collecting an extended set of labeled ground-truth data is needed as the **preliminary step**. Additionally, periodic data cleansing and replenishment are necessary to mitigate biases from cloud gaming platforms updating their adaptation behaviors, a phenomenon known as concept drift [48,49]. These tasks are typically achieved through dedicated data engineering teams utilizing a Digital Twin to mimic the operational network, which are standard DataOps and MLOps practices widely used in the industry [50–52].

The modular architecture of our assessment process ensures scalability and customizability, with additional engineering adjustments required according to the changed or extended assessment objectives. Our **packet processing module** operates offline on the labeled dataset to extract statistical volumetric information, which can be achieved via network traffic analysis tools (e.g., tshark) and libraries (e.g., gopacket), as demonstrated in our reference implementation [33]. The **assessment module** implements the score aggregation framework detailed in Section 6, which intrinsically allows for customized choices of the adaptability correlation function and schema extensions (see Section 6.4). As it is decoupled from the other stages of the process, assessors (e.g., ISPs) can adjust the assessment logic to cater to their specific demands with minimal engineering overhead.

8. 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 crucial role in the gameplay experience (i.e., QoE) perceived by users. In this paper, we empirically profile the streaming adaptability of four major cloud gaming platforms to systematically understand how they adapt their network streaming characteristics to client settings and how they optimize game streaming experiences to network conditions. A systematic framework is developed and implemented to quantitatively assess the cloud game streaming adaptability of a platform with ground-truth traffic traces. The insights obtained from this study using our assessment framework provide a reference for ISPs and relevant stakeholders to benchmark and optimally manage their networks for supporting expected quality level of cloud gaming user experience.

Funding

This project is funded by the [Australian Government](#)'s Cooperative Research Centres Projects (CRC-P) Grant CRCPXIV000099. This article is an extended and improved version of our paper presented at the IFIP Networking 2024 conference [53].

CRediT authorship contribution statement

Minzhao Lyu: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Yifan Wang:** Writing – review & editing, Visualization, Validation, Software, Investigation, Data curation; **Vijay Sivaraman:** Writing – review & editing, Supervision, Resources, Project administration, Investigation, Funding acquisition.

Data availability

I have shared the link to my data in the paper.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.


Acknowledgment

The authors would like to thank our volunteer Mr. Aaron King from North Carolina State University who collected packet capture (PCAP) files of cloud gaming sessions on the Amazon Luna and PS5 Cloud Streaming platforms.


References

- [1] F. Aumont, F. Humbert, C. Neumann, C. Salmon-Legagneur, C. Taibi, Dissecting cloud game streaming platforms regarding the impacts of video encoding and networking constraints on QoE, in: Proc. ACM GameSys, Istanbul, Turkey, 2021.
- [2] H. Iqbal, A. Khalid, M. Shahzad, Dissecting cloud gaming performance with DECAF, Proc. ACM Meas. Anal. Comput. Syst. 5 (3) (2021) 1–27.
- [3] M. Carrascosa, B. Bellalta, Cloud-gaming: analysis of google stadia traffic, Comput. Commun. (2022).
- [4] X. Marchal, P. Graff, J.R. Ky, T. Cholez, S. Tuffin, B. Mathieu, O. Fester, An analysis of cloud gaming platforms behaviour under synthetic network constraints and real cellular networks conditions, J. Netw. Syst. Manage. 31 (2023) 1–31.
- [5] M. Lyu, S.C. Madanapalli, A. Vishwanath, V. Sivaraman, Network anatomy and real-time measurement of Nvidia GeForce NOW cloud gaming, in: Proc. PAM, Virtual Event, 2024.
- [6] M. Lyu, V. Sivaraman, A large-scale network measurement study of NVIDIA GeForce NOW cloud gaming in the wild, IEEE Trans. Netw. (2025) 1–16.
- [7] Y. Wang, M. Lyu, V. Sivaraman, Games are not equal: classifying cloud gaming contexts for effective user experience measurement, in: Proc. ACM IMC, Madison, WI, USA, 2025.
- [8] P. Graff, X. Marchal, T. Cholez, B. Mathieu, O. Fester, Efficient identification of cloud gaming traffic at the edge, in: Proc. IEEE/IFIP NOMS, Miami, FL, USA, 2023.
- [9] M. Carrascosa, B. Bellalta, Cloud-gaming: analysis of google stadia traffic, Comput. Commun. 188 (2022).
- [10] M. Suznjivic, I. Slivar, 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, 2016.
- [11] S. Bhuyan, S. Zhao, Z. Ying, M.T. Kandemir, C.R. Das, End-to-end characterization of game streaming applications on mobile platforms, Proc. ACM Meas. Anal. Comput. Syst. 6 (1) (2022) 1–25.
- [12] Market.us, Cloud Gaming Market to Soar to USD 143.4 Bn by 2032, 2024, Accessed: 2025-10-30, <https://finance.yahoo.com/news/cloud-gaming-market-soar-usd-121200504.html>.
- [13] I. Slivar, M. Suznjivic, L. Skorin-Kapov, 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, 2015.
- [14] S. Zadtotooghaj, S. Schmidt, S. Moller, Modeling gaming QoE: towards the impact of frame rate and bit rate on cloud gaming, in: Proc. QoMEX, Cagliari, Italy, 2018.
- [15] S.S. Sabet, S. Schmidt, S. Zadtotooghaj, C. Griwodz, S. Moller, Towards the impact of gamers strategy and user inputs on the delay sensitivity of cloud games, in: Proc. International Conference on Quality of Multimedia Experience, Athlone, Ireland, 2020, pp. 1–3.
- [16] Z. He, Y. Yang, S. Li, L. Qiu, D. Dai, Y. Yang, VIGOR: reviving cloud gaming sessions, Proc. ACM Netw. (2024).
- [17] T. Kämäräinen, M. Siekkinen, A. Ylä-Jääski, W. Zhang, P. Hui, A measurement study on achieving imperceptible latency in mobile cloud gaming, in: Proc. ACM MMSys, Taipei, Taiwan, 2017.
- [18] C. Baena, O.S. Peñaherrera-Pulla, R. Barco, S. Fortes, Measuring and estimating key quality indicators in cloud gaming services, Comput. Networks (2023).
- [19] M. Carvalho, D. Soares, D.F. Macedo, Transfer learning-based QoE estimation for different cloud gaming contexts, in: Proc. IEEE NetSoft, Madrid, Spain, 2023.
- [20] J.R. Ky, P. Graff, B. Mathieu, T. Cholez, A hybrid P4/NFV architecture for cloud gaming traffic detection with unsupervised ML, in: Proc. IEEE ISCC, Gammarth, Tunisia, 2023.
- [21] M. Suznjivic, J. Beyer, L. Skorin-Kapov, S. Moller, N. Sorsa, Towards understanding the relationship between game type and network traffic for cloud gaming, in: Proc. IEEE ICMEW, Chengdu, China, 2014.
- [22] A. Di Domenico, G. Perna, M. Trevisan, L. Vassio, D. Giordano, A network analysis on cloud gaming: stadia, GeForce now and PSNow, Network 1 (3) (2021) 247–260.
- [23] H. Chen, X. Zhang, Y. Xu, J. Ren, J. Fan, Z. Ma, W. Zhang, T-gaming: a cost-efficient cloud gaming system at scale, IEEE Trans. Parallel Distrib. Syst. 30 (12) (2019) 2849–2865.
- [24] P. Hamadianian, D. Gallatin, M. Alizadeh, K. Chintalapudi, Ekho: synchronizing cloud gaming media across multiple endpoints, in: Proc. ACM SIGCOMM, New York, NY, USA, 2023.
- [25] J. Wu, Y. Guan, Q. Mao, Y. Cui, Z. Guo, X. Zhang, ZGaming: zero-latency 3D cloud gaming by image prediction, in: Proc. ACM SIGCOMM, New York, NY, USA, 2023.
- [26] K.-T. Chen, Y.-C. Chang, H.-J. Hsu, D.-Y. Chen, C.-Y. Huang, C.-H. Hsu, On the quality of service of cloud gaming systems, IEEE Trans. Multimedia 16 (2) (2014) 480–495.
- [27] S. Schmidt, S. Uhrig, D. Reuschel, Investigating the relationship of mental immersion and physiological measures during cloud gaming, in: Proc. QoMEX, Athlone, Ireland, 2020.
- [28] H. Habibi Gharakheili, M. Lyu, Y. Wang, H. Kumar, V. Sivaraman, iTelescope: software network middle-box for real-time video telemetry and classification, IEEE Trans. Netw. Serv. Manage. 16 (3) (2019) 1071–1085.
- [29] S.C. Madanapalli, A. Mathai, H. Habibi Gharakheili, V. Sivaraman, ReCLive: real-time classification and QoE inference of live video streaming services, in: Proc. IEEE/ACM IWQoS, Tokyo, Japan, 2021.
- [30] O. Michel, S. Sengupta, H. Kim, R. Netravali, J. Rexford, Enabling passive measurement of zoom performance in production networks, in: Proc. ACM IMC, Nice, France, 2022.
- [31] S.C. Madanapalli, H. Habibi Gharakheili, V. Sivaraman, Know thy lag: in-network game detection and latency measurement, in: Proc. PAM, Virtual Event, 2022.
- [32] A.N. Kuznetsov, tc(8) - Linux manual page, 2001, <https://man7.org/linux/man-pages/man8/tc.8.html>.

- [33] M. Lyu, Multimedia Network Traffic Dataset, 2025, Accessed: 2025-10-20, <https://minzhaolyu.github.io/dataset/MultimediaNetworkTrafficDataset>.
- [34] Zorbak1, GTA V is the most optimized PC game I've ever played, 2015, Accessed: 2024-01-03, <https://steamcommunity.com/app/271590/discussions/0/617335934127794199/?ctp=6>.
- [35] Steam, Steam Hardware & Software Survey: December 2025, 2025, Accessed: 2026-01-03, <https://store.steampowered.com/hwsurvey/Steam-Hardware-Software-Survey-Welcome-to-Steam>.
- [36] H. Schulzrinne, S.L. Casner, R. Frederick, V. Jacobson, RTP: A Transport Protocol for Real-Time Applications, RFC 3550, 2003. <https://doi.org/10.17487/RFC3550>.
- [37] Vedor, Sword and Fairy 7 FPS low FPS on equipment page, 2023, Accessed: 2024-01-03, https://www.reddit.com/r/XboxGamePass/comments/14twotm/sword_and_fairy_7_fps_low_fps_on_equipment_page/.
- [38] K.-T. Chen, P. Huang, C.-L. Lei, Effect of network quality on player departure behavior in online games, *IEEE Trans. Parallel Distrib. Syst.* (2009) 281–293.
- [39] D. Loguinov, H. Radha, Measurement study of low-bitrate internet video streaming, in: *Proc. ACM SIGCOMM Workshop on Internet Measurement*, San Francisco, CA, USA, 2001.
- [40] R.D. Tripathi, M. Lyu, V. Sivaraman, Assessing the impact of network quality-of-service on metaverse virtual reality user experience, in: *Proc. IEEE MetaCom*, 2024, pp. 206–213.
- [41] Y.O. Sharrab, N.J. Sarhan, Modeling and analysis of power consumption in live video streaming systems, *ACM Trans. Multimedia Comput. Commun. Appl.* 13 (4) (2017).
- [42] D. Danier, F. Zhang, D.R. Bull, BVI-VFI: a video quality database for video frame interpolation, *IEEE Trans. Image Process.* 32 (2023) 6004–6019.
- [43] A.I. Scale, Accelerating the Development of AI Applications, 2025, Accessed: 2025-10-30, <https://scale.com/>.
- [44] Labelbox, The Data Factory of AI Teams, 2025, Accessed: 2025-10-30, <https://labelbox.com/>.
- [45] A. Shirmarz, G. de Castro Ariel, V.H.S. Lopes, F. Verdi, M. Luizelli, C. Rothenberg, CGSynth: cloud gaming synthesizer, in: *Proc. ACM SIGCOMM Posters and Demos*, Coimbra, Portugal, 2025, p. 175–177.
- [46] A. Shirmarz, A.G. de Castro, F.L. Verdi, C.E. Rothenberg, CGReplay: capture and replay of cloud gaming traffic for QoE/QoS assessment, in: *Proc. IEEE NetSoft*, Budapest, Hungary, 2025, pp. 312–314.
- [47] X. Jiang, S. Liu, A. Gember-Jacobson, A.N. Bhagoji, P. Schmitt, F. Bronzino, N. Feamster, NetDiffusion: network data augmentation through protocol-constrained traffic generation, *Proc. ACM Meas. Anal. Comput. Syst.* 8 (1) (2024).
- [48] L. Yang, W. Guo, Q. Hao, A. Ciptadi, A. Ahmadzadeh, X. Xing, G. Wang, CADE: detecting and explaining concept drift samples for security applications, in: *Proc. USENIX Security*, 2021, pp. 2327–2344.
- [49] S. Liu, F. Bronzino, P. Schmitt, A.N. Bhagoji, N. Feamster, H.G. Crespo, T. Coyle, B. Ward, LEAF: navigating concept drift in cellular networks, *Proc. ACM Netw.* 1 (CoNEXT2) (2023).
- [50] B.P.R. Rella, MLOps and DataOps integration for scalable machine learning deployment, *Int. J. Multidiscipl. Res.* 4 (2022) 20.
- [51] R. Poorzare, D.N. Kanellopoulos, V.K. Sharma, P. Dalapati, O.P. Waldhorst, Network digital twin toward networking, telecommunications, and traffic engineering: a survey, *IEEE Access* 13 (2025) 16489–16538.
- [52] Databricks, Big Book of Data Engineering, 2025, Accessed: 2026-01-03, <https://www.databricks.com/sites/default/files/2025-11/big-book-data-engineering.pdf>.
- [53] M. Lyu, Y. Wang, V. Sivaraman, Do cloud games adapt to client settings and network conditions?, in: *Proc. IFIP Networking*, Thessaloniki, Greece, 2024.



Minzhao Lyu received his B Eng (First Class Hons.) and Ph.D. degree from the University of New South Wales, Sydney, Australia in 2017 and 2022 respectively. He is a Lecturer at the University of New South Wales, Sydney, Australia, where he has worked as a Research Associate. His research interests include network measurement, network security, multimedia Quality-of-Experience (QoE), and AI for networking.



Yifan Wang received his B Eng (First Class Hons.) in Computer Engineering from the University of New South Wales, Sydney, Australia in 2022. He is now a PhD student at the School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, Australia. His research topics focus on AI-empowered network measurement for multimedia streaming applications.



Vijay Sivaraman received his B Tech from the Indian Institute of Technology in Delhi, India, in 1994, his M S from North Carolina State University in 1996, and his PhD from the University of California at Los Angeles in 2000. He has worked at Bell-Labs as a student Fellow, in a silicon valley start-up manufacturing optical switch-routers, and as a Senior Research Engineer at the CSIRO in Australia. He is now a Professor at the University of New South Wales in Sydney, Australia. His research interests include Software Defined Networking, network architectures, and cyber-security particularly for IoT networks.