Streaming Media over LEO Satellite Networking: A Measurement-Based Analysis and Optimization

HAO FANG* and HAOYUAN ZHAO*, Simon Fraser University, Canada
FENG WANG, University of Mississippi, United States
YI CHING CHOU, LONG CHEN, JIANXIN SHI, and JIANGCHUAN LIU†, Simon Fraser University, Canada

Recently, Low Earth orbit Satellite Networks (LSNs) have been suggested as a critical and promising component toward high-bandwidth and low-latency global coverage in the upcoming 6G communication infrastructure. SpaceX's Starlink is arguably the largest and most operational LSN to date. There have been practical uses of Starlink across diverse networked applications, including those with stringent demands, such as multimedia applications. Given the mixed and inconsistent feedback from end users, it remains unclear whether today's LSNs, in particular Starlink, are ready for realtime multimedia. In this paper, we present a systematic measurement study on realtime multimedia services over Starlink, seeking insights into their operations and performance in this new generation of networking. Our findings demonstrate that Starlink can handle most video-on-demand (VoD) and live-streaming services with properly configured buffers but suffers from video pauses or audio cut-offs during interactive videoconferencing. We identify the key factors that impact the performance of LSN, particularly for multimedia services, including satellite switching, routing strategies and weather conditions. Our findings offer valuable hints into future enhancements for multimedia services over LSNs. Specifically, we further propose a Weather Aware Buffer Based Rate Adaption algorithm based on our observations on weather impacts, which is capable of maximizing the quality of experience for VoD applications with seamless integration of dynamic weather conditions.

CCS Concepts: • **Networks** → **Network measurement**; **Wireless access networks**; • **Information systems** → **Multimedia streaming**.

Additional Key Words and Phrases: Low Earth Orbit Satellite Networks, Multimedia Services, Video Streaming, Video Conferencing, Network Performance Measurement & Optimization

1 INTRODUCTION

Recently, Low Earth orbit (LEO) satellites have attracted tremendous attention from both academia and industries. Different from its high orbit counterpart, particularly the Geosynchronous Orbit (GEO) satellite that stays at around 35, 780 km, LEO satellites orbit at distances ranging from 180 km to 2, 000 km from the Earth's surface., which can greatly reduce the space-ground communication delay and increase throughput. As the short orbit distance from the surface of the Earth also reduces

*Both authors contributed equally to this research.
†Corresponding author

This paper extends our previous work published in ACM NOSSDAV'23 [53].
Authors' addresses: Hao Fang, fanghaof@sfu.ca; Haoyuan Zhao, hza127@sfu.ca, Simon Fraser University, 8888 University Dr. W, Burnaby, BC, Canada, V5A 1S6; Feng Wang, fwang@cs.olemiss.edu, University of Mississippi, University, Oxford, United States, 338677; Yi Ching Chou, ycchou@sfu.ca; Long Chen, longchen.cs@ieee.org; Jianxin Shi, jsa394@sfu.ca; Jiangchuan Liu, jcliu@sfu.ca, Simon Fraser University, 8888 University Dr. W, Burnaby, BC, Canada, V5A 1S6.

XXXX-XXXX/2024/10-ART $15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

the coverage of an LEO satellite, the LEO Satellite Network (LSN) constellation, where a large number of LEO satellites work together to achieve full global coverage, has been envisioned to offer anywhere anytime network services to ground users. To this end, many commercial deployments have been started in the past several years. For example, so far *SpaceX's Starlink*, one of the world's leading LSN service providers, already has 5, 700 LEO satellites in operation [30], with an ambitious vision to make the next-generation Starlink constellation eventually harbor up to 30,000 [41]. *OneWeb*, another industrial leader on LSN, is also advancing its LSN deployment with a target of 7,000 broadband satellites [35]. Such rapid growth makes LSN be proposed as a key component to be integrated into the communication infrastructure towards the upcoming 6G and beyond [25].

With improved bandwidth and latency, the emerging LSN brings opportunities to better support various popular application services. For example, prior works have proven that LSN is a viable solution for simple applications like file sharing and Web browsing [21, 26, 31]. Nevertheless, the performance of LSN on other advanced applications, particularly multimedia services, remains uncertain. Given that multimedia services currently dominate more than 60% of internet traffic [42], it is crucial to evaluate their performance over LSN, and take into account the current challenges and potential solutions.

In this paper, we systematically measure and analyze multimedia services over LSN. In detail, we choose three typical multimedia service types: video-on-demand (VoD) service, live streaming service, and videoconferencing service. With over six months of measurement, we have conducted a deep investigation into the general performance and behaviors of these multimedia services over the Starlink network. Our measurements indicate that the Starlink network can provide a decent Quality of Experience (QoE) for VoD and live streaming services with properly configured buffers, but it encounters issues such as video pauses and audio cuts-off in videoconferencing applications. Furthermore, we identify several key factors that can significantly impact the performance of LSN and, thus, its multimedia services. For example, satellite switching, the emerging routing strategy *Inter-Satellite Links* (ISLs), and weather conditions. Those informative discoveries can provide valuable insights for future development in both multimedia services and LSNs. As a demonstration, we further introduce a Weather Aware Buffer Based Rate Adaption (WABB) algorithm aimed at mitigating the detrimental effects during different weather conditions for VoD applications. In contrast to Terrestrial Networks, our measurement reveals that adverse weather conditions, such as thunderstorms, often result in significant performance degradation for multimedia services employing LSN. The outage duration and frequency during thunderstorm days increase by 32.37 and 128.67 times, respectively, compared to the clear days. WABB dynamically adjusts the bitrate and buffer size based on a carefully designed heuristic algorithm to maximize the QoE. Our approach has demonstrated superior performance compared to the majority of state-of-the-art adaptive bitrate (ABR) algorithms, as confirmed through evaluation.

The key contributions of our paper can be summarized as follows:

- We conduct a comprehensive measurement study of Starlink, analyzing its performance across three key multimedia services: VoD, live streaming, and videoconferencing, further identifying major factors that affect the Starlink network performance.
- We examine how satellite switching impacts packet-level behaviors by measuring latency and packet loss rate, identifying two types of packet loss patterns that can have varying impacts on multimedia service performance.
- We examine and compare the performance of existing transport protocols and congestion control algorithms, originally designed for terrestrial networks, on the Starlink network.

- We perform an extensive study on the performance of the Starlink network under varying weather conditions, examining the relationship between network outages and different weather scenarios.
- We introduce an innovative WABB algorithm specifically tailored for LSNs as a demonstration of how our measurement findings can help future enhancements for multimedia services over LSNs. This algorithm is designed to optimize the QoE under a range of weather conditions.

The remainder of this paper is organized as follows. Section 2 reviews the literature of recent works on satellite and multimedia service measurement studies. In Section 3, we provide an overview of the measurement methodology as well as our general observations. The detailed analysis of the performance of three typical multimedia services over the Starlink network is then presented in Section 4. We further discuss some additional impact factors that can significantly influence LSN performance in Section. 5. The proposed method WABB algorithm and its evaluation will be discussed in Section 6. In the end, we conclude our paper in Section 7.

2 RELATED WORKS

While LSNs are still in their nascent stages, some initial explorations into their performance have been conducted. These early studies often employed relatively straightforward setups focusing on single applications [45] or concentrated on the physical layer performance of LSNs [22, 33, 43]. Recent research has expanded to examine various applications and services on LSNs, such as web browsing and QUIC performance on the Starlink network, including potential factors affecting its performance like weather, obstruction, satellite movement, and routing strategy [21, 26, 31]. Unlike these studies, our work adopts a broader approach by examining a range of multimedia services and how their end-to-end performance is influenced by the characteristics of LSNs. This extended version of our conference paper [53] includes additional measurements such as packet loss, various transport layer protocols and congestion control algorithms, and more comprehensive weather analysis. Additionally, we have developed and proposed a novel ABR algorithm, WABB, to enhance the QoE for VoD services under LSNs across different weather conditions.

Multimedia services are some of the most widely utilized services on the Internet, and assessing their performance across different network systems is crucial for advancing both service quality and network technology. There has been a multitude of pioneering research in this area. For instance, several studies have focused on evaluating the performance of live streaming services over 5G networks [34, 44, 48]. Additionally, various research efforts have examined different multimedia services, such as traditional VoD applications [16] and case studies on platforms like *Twitch.TV* [52]. The COVID-19 pandemic has also triggered numerous studies on the network performance of videoconferencing applications [9, 11, 27], with reports indicating a fourfold increase in traffic for popular applications like Zoom since 2020 [13]. In-depth analyses have been conducted on Zoom's network performance at the packet level [32]. Differing from these works that primarily focused on modern terrestrial networks, our research uniquely concentrates on how various multimedia services perform over the newly emerging LSNs. Our study delves into the distinct impacts of LSN characteristics on these services, a relatively unexplored aspect in current literature.

A number of pioneering works have striven to enhance user QoE in multimedia services, particularly via developing various ABR algorithms. These efforts range from buffer-based methods [19] to reinforcement learning techniques [23, 28], and Model Predictive Control (MPC) approaches [50]. Specific algorithms, such as Low-on-Latency (LoL) and LoL$^+$, have been designed with low latency demands of live streaming platforms in mind [8, 20, 24, 29], employing a blend of bitrate adaptation, heuristic-based playback speed control, and thorough throughput calculation. However, despite these advancements, there is still a significant gap in developing ABR solutions tailored for

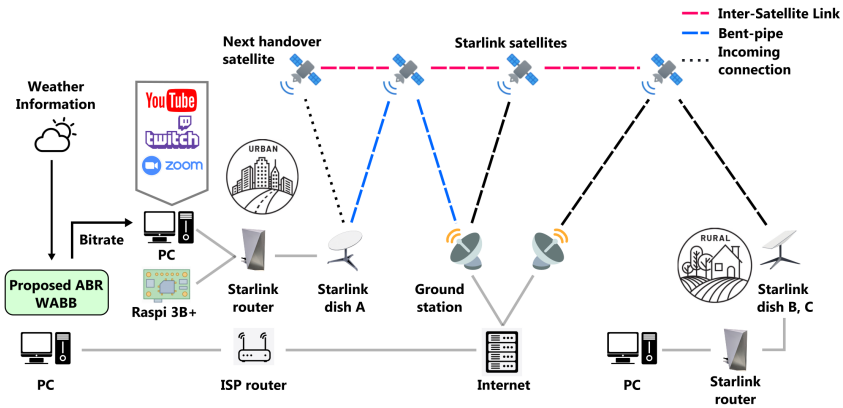

Fig. 1. Measurement setup across various scenarios and tools. Dish B and C utilize the same experimental setup and are positioned in different geographical locations.

LSNs. Although there have been attempts to develop specialized TCP protocols for LEO satellite networks [10], these works did not fully address the unique challenges of LSNs, particularly the impact of weather conditions. In response to this, our proposed ABR algorithm, informed by our comprehensive measurement study, is specifically designed to incorporate the impact of weather conditions on LSNs and their multimedia services. This approach is aimed at significantly enhancing the viewing experience of users under varying weather-related network performance scenarios and serves as a demonstration of how our measurement findings can offer valuable hints into future enhancements for multimedia services and LSNs.

3 MEASUREMENT OVERVIEW

Our measurement, conducted over approximately six months starting from December 2022, involved the deployment of two Starlink dishes in distinct locations, as illustrated in Fig. 1. Dish A, a Gen 1 dish, is situated in an urban area on the West Coast of North America. Dish B, a Gen 2 dish, is located in a rural area in the Mid-South of North America. Dish C, a Gen 2 dish, is located in a rural area on the West Coast of North America. All dishes subscribe to Starlink's standard plan, without network priority, to ensure that the network performance measurements are representative of most users. All our dishes were positioned on vantage points to ensure an unobstructed view of the sky. Since Dish C is located in the middle of a forest, the dish has to be mounted on top of a tree approximately 50 meters high to achieve a clear view of the sky. The obstructed ratio, as reported by the Starlink portal[1], is *0.734%* for dish A, *0.029%* for dish B and *0.753%* for dish C. Additionally, the Starlink mobile application[2] will be used to record the Starlink network outages.

For data collection and analysis, we utilize two types of equipment capable of handling computationally intensive tasks (e.g., high-resolution video playback) and large-scale measurements. The first type consists of typical desktop computers equipped with CPUs of up to 12 cores at 2.1GHz. These computers are directly connected to Starlink routers via WiFi (referred to as Starlink WiFi) and/or Ethernet (referred to as Starlink Ethernet) to access the satellite network. To ensure a good quality WiFi connection, all computers are placed in proximity to the corresponding routers. The second type of device is a Raspberry Pi 3B+ with a 300 MBps Ethernet port, connected to the Starlink router using an Ethernet cable. The Raspberry Pi records long-term *traceroute* and *ping* results, allowing desktop computers to focus on measuring multimedia services. For comparison, we

[1]http://dishy.starlink.com/ (Accessible exclusively via the Starlink Network)
[2]https://play.google.com/store/apps/details?id=com.starlink.mobile

Table 1. Target services and measurement tools.

Service Type	Representative Platform	Streaming Protocol	Software Tools
Video-on-demand	YouTube	DASH	*Stats for nerds* / *Ntopng* / *ping*
Live streaming	Twitch	HLS	*Video Stats* / *Ntopng* / *ping*
Videoconferencing	Zoom	Zoom's own protocol	*Zoom Meeting API*/ *ping*

conduct baseline measurements using a typical terrestrial network service (referred to as Terrestrial Network), providing up to 1000 MBps for download and upload. The stability of this network is verified multiple times, yielding consistent baseline results.

We carefully select three of the most popular multimedia services to systematically measure their end-to-end performance for further analysis. These services include VoD, live streaming, and videoconferencing. For each service type, we opt for the most representative platform as a case study, namely YouTube for VoD service, Twitch for live streaming service, and Zoom for videoconferencing service. This selection ensures coverage across a reasonable variety of prominent multimedia service platforms, each employing different popular streaming protocols. To enhance our understanding and analysis of the results, we utilize various monitoring and analysis tools at both the application and network levels. The details of our target services and measurement tools are summarized in Table 1. At the application level, YouTube and Twitch provide built-in video statistics monitoring tools such as *Stats for nerds* and *Video Stats*. Zoom, on the other hand, utilizes its built-in *Zoom Meeting API* to retrieve video and network statistics for a meeting. Simultaneously, *Ntopng* [15], a web-based network traffic monitoring tool, is employed to track network-level data during the experiments. Latency and routing strategies for the considered LSN were measured and analyzed using *ping* and *traceroute*. To evaluate the packet loss during videoconferencing, we employ *Ringmaster* [49] to record the packet behaviour between two remote hosts and analyze it in conjunction with the satellite handover strategy.

Our findings indicate that present-day LSNs generally possess the capability to support a variety of multimedia services. However, our analysis also brings to light certain limitations, particularly considering that Starlink, despite being an industry-leading pioneer, is still in the process of expanding to its full coverage potential. These limitations can have varying impacts on performance and service quality, contingent on factors such as the type of multimedia service and the streaming protocol employed by the platform. This underscores the need for the development of novel solutions on both the LSN and multimedia service fronts to address these limitations and ensure optimal multimedia performance and service quality. The subsequent sections delve into a presentation and analysis of our measurement results pertaining to diverse multimedia services over the Starlink network. Additionally, we explore further observations and insights that can potentially inform the design of solutions aimed at enhancing multimedia services.

4 MEASUREMENT RESULT AND ANALYSIS

4.1 Video-on-demand Service

To evaluate the overall performance of the VoD service, we select several representative video contents [12, 46] for our testbed, including Sports [2], Gaming [6], Chatting [5], and Outdoor [1]. These videos feature a mix of static backgrounds and rapid scene changes, as well as complex scenes like dense forests and relatively simple backgrounds such as monochrome curtains. This diversity ensures varying levels of video encoding complexity and video chunk sizes, allowing us to evaluate the Starlink network performance under different traffic loads. For each video, we use the 4k resolution based on Starlink's available bandwidth [26, 31] and the actual bitrate of our selected

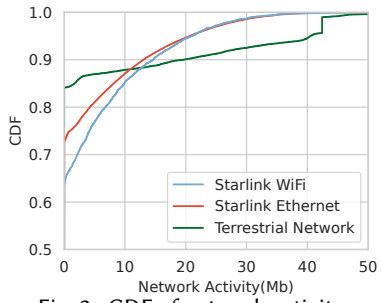

Fig. 2. CDF of network activity.

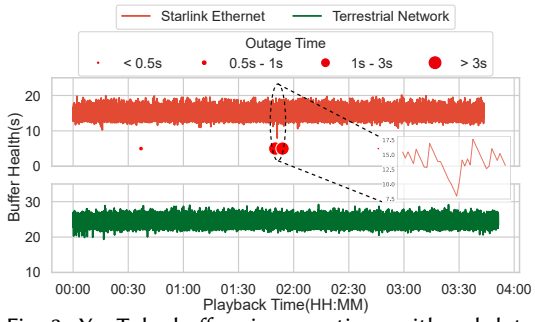

Fig. 3. YouTube buffer size over time, with red dots denoting Starlink network outages.

videos, which varies from 22 to 40 Mbps. During measurement, we used a fixed highest bitrate with ABR streaming disabled. Our purpose is to stress-test the Starlink network performance, allowing us to observe any potential issues caused by LSNs, such as rebuffering events. For each experiment, a Python script is used to automatically loop the video and record video playback status from the YouTube built-in monitoring tool, *Stat for nerds*. The script will automate the browser to select the preset video resolution, open *Stat for nerds* window, start to playback the video, and scrape the video status with an interval of 1 second. Simultaneously, another thread will also start to scrape the network data from *Ntopng's* browser API to monitor low-level network traffic. To prevent advertisements from flushing the video buffer, a premium account is used to ensure a smooth measurement process. The *Stat for nerds* provides video status such as video resolution, frame drops, connection speed, network activity, and buffer health for current playback. In detail, i) Connection speed reflects the actual download speed, updating when the client actively downloads video chunks. ii) Network activity denotes the per-second downloaded size in bytes, which can reflect the actual video chunk size downloaded. iii) Buffer health represents the length of the downloaded video content in seconds and iv) Frame drops stand for the cumulative lost or corrupted frames. To ensure consistency, the script will clear the browser cache after each video loop, compelling the client to download each video segment afresh from the server. Eventually, we collected approximately 34 hours of video data from two dishes with various video types.

We first compare the performance between different networks. We observe that the average connection speeds for Starlink Ethernet, Starlink WiFi, and the Terrestrial Network are 67 Mbps, 62 Mbps, and 381 Mbps, respectively during the video playback. The cumulative distribution function (CDF) of overall network activity is depicted in Fig. 2. Notably, the Terrestrial Network exhibits 88% idle time, whereas Starlink Ethernet and WiFi show 73% and 65%, respectively (indicated by network activity at 0 KB). This implies that the Starlink network tends to require more time to download the subsequent video segment and experiences more frequent network activities than the Terrestrial Network. This result suggests that the Terrestrial Network surpasses Starlink Ethernet/WiFi in terms of both connection speed and network activity. Moreover, throughout the entire measurement process, we also observe that Starlink WiFi closely mirrors the performance of Starlink Ethernet, indicating that the majority of network impacts stem from satellite links. As a result, our subsequent analysis will focus on evaluating the performance of Starlink Ethernet, as it serves as an upper bound for the overall performance of the Starlink network.

Fig. 3 depicts the evolution of YouTube's buffer size over time, where red dots correspond to network outages reported by the Starlink mobile application, and their size is proportionate to the duration of the outage. The average buffer size is approximately 15 seconds for Starlink Ethernet, while it extends to around 25 seconds for the Terrestrial Network. This disparity arises from the Terrestrial Network's greater bandwidth, allowing it to prefetch more segments for each

Table 2. Starlink Ethernet performance across common video contents.

	Sport	Gaming	Chatting	Outdoor
Connection Speed (Mbps)	92.82 ± 33.11	101.76 ± 29.51	59.13 ± 8.98	53.97 ± 11.28
Network Activity (Mb)	23.40 ± 16.82	21.59 ± 15.02	13.53 ± 10.51	10.06 ± 7.43
Buffer Size (s)	8.37 ± 1.68	12.77 ± 1.48	19.90 ± 1.86	16.73 ± 12.43

The values follow the *Mean* ± *Std* format.

downloading event. For Starlink Ethernet, we observe an average outage time of about 6.46 seconds, with the longest outage lasting 18 seconds, which will lead to notable rebuffering events. Two illustrative instances of such rebuffering can be identified at 00:15 and 01:50 in Fig. 3. Despite these outages, the substantial buffer size typically employed in VoD services proves generally adequate to compensate for Starlink outages, ensuring a seamless viewing experience.

Table 2 illustrates the video status using Starlink Ethernet across various types of content. Sports and gaming videos, which feature rapid scene changes and complex visuals, exhibit higher network activity and variance compared to other types. Despite these challenges, Starlink Ethernet effectively manages these demanding playback tasks. Notably, sports and gaming videos tend to have shorter buffer sizes, potentially due to larger video chunk sizes that might be constrained by YouTube's memory allocation limits. Additionally, it's important to mention that during the playback of outdoor videos, the Starlink network experienced several outages, leading to a significant variance in buffer size. Overall, the Starlink network demonstrates the capability to support popular video content at 4k resolution.

In conclusion, our findings indicate that the end-to-end performance of the VoD service over the Starlink network is comparable to that of the conventional Terrestrial Network. This aligns with the results of previous research studies [26, 31]. Generally, users should not perceive significant differences in their experiences while using the VoD service on either network. However, it is essential to note that while Starlink currently effectively supports 4K video at 24 FPS, the performance may experience a notable decline when dealing with more bandwidth-intensive content. A prime example is high-definition 360 videos, as they pose a greater demand on Starlink's bandwidth.

4.2 Live streaming Service

As the video content and the duration of live streaming are under the control of the streamer, we select a consistently active channel [4] with relatively fixed streaming content for our measurement target. The live streaming channel maintains a resolution of 1,920 x 1,080 at 35 FPS, with an average bitrate of around 6 Mbps. Similar to Section 4.1, ABR streaming is disabled, and we manually select the highest bitrate. In addition, a Python script is used to automate the browser and record the data from *Video Stats* and *Ntopng* simultaneously. The collected dataset encompasses approximately 77 hours of live streaming service, sampled at a rate of one data point per second. The built-in tool *Video Stats* offers metrics such as video resolution, FPS, skipped frames, buffer size, latency to broadcaster (LtB), and playback bitrate. LtB signifies the duration from the moment the video content is captured at the broadcaster's end to when it is displayed on the viewer's monitor. Streamers can adjust LtB based on the channel's focus. For instance, opting for low-latency mode when realtime interactions with viewers via text chat are crucial. In platforms like Twitch, LtB can dynamically adapt using the LoL method [8, 24]. This method involves assigning larger LtB values to viewers with poor network conditions, enabling a longer pre-loaded video content duration but simultaneously diminishing the realtime interaction experience for viewers.

The overall performance metrics for Twitch are detailed in Table 3. The Terrestrial Network exhibits a significant advantage over Starlink Ethernet/WiFi, with the exception of the playback bitrate. This observation suggests that both the Starlink and Terrestrial Networks have sufficient

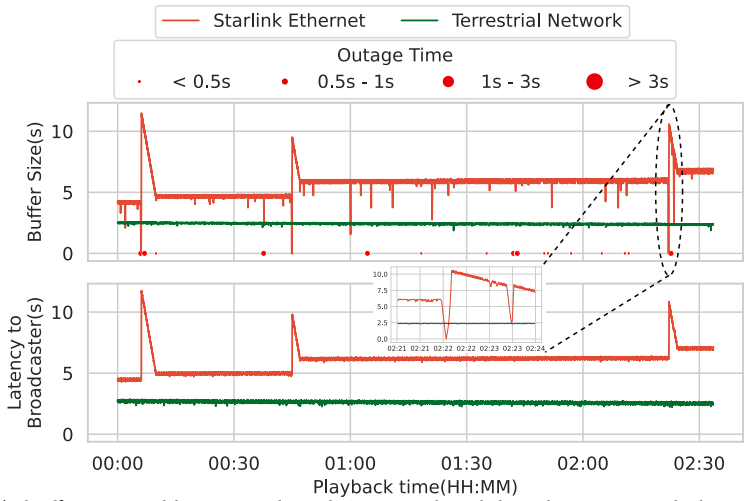

Fig. 4. Twitch's buffer size and latency to broadcaster, with red dots denoting Starlink network outages.

bandwidth for playback. In Fig. 4, we analyze how the buffer size and LtB in Twitch evolve over a measurement round lasting approximately 2.5 hours. The service via Starlink Ethernet encounters a higher frequency of buffer drops compared to the service via the Terrestrial Network. Furthermore, some of these drops are substantial enough to lead to rebuffering. For instance, at 00:06, 00:45, and 02:22, there are severe buffer drops that deplete all preloaded video contents. Following each outage, the LoL mechanism elevates the LtB to around 10 seconds, gradually reducing until reaching a stable state at a higher latency level compared to the previous state. The higher LtB facilitates clients to prefetch more video content and is an affordable measure for potential future outages. For example, at 02:23, another extended outage happens just after the buffer drops, yet the streaming continues without interruption due to the extended buffer size. It is noteworthy that the LtB decrement is achieved by accelerating video playback at a rate of 1.025, and this speed change is generally imperceptible to humans. In live streaming services, the length of available buffered video is notably shorter compared to VoD services. Therefore, outages in the Starlink network are more likely to deplete the current buffer, resulting in streaming freezes and directly impacting the user's experience.

Examining Fig. 4 yields another notable observation: certain instances of buffer drops do not align with the outages reported by the Starlink mobile application. A deeper investigation uncovers that these drops arise due to substantial network jitters stemming from satellite switching—a topic that will be elaborated upon in Section 5. In short, the failure of satellite handover or a considerable delay difference between the two satellites involved can lead to notable network jitter over a brief period, posing a detrimental impact on the quality of multimedia services.

In summary, compared to VoD service, live streaming service over the Starlink network is more likely to experience severe buffer drops and rebuffering.

Table 3. Average statistics of Twitch service performance.

	Starlink Ethernet	Starlink WiFi	Terrestrial Network
Latency to broadcaster(s)	5.55	7.91	2.54
Playback bitrate(kbps)	6000.93	5986.71	6001.97
Frame drops per hour	3.20	10.43	0.25

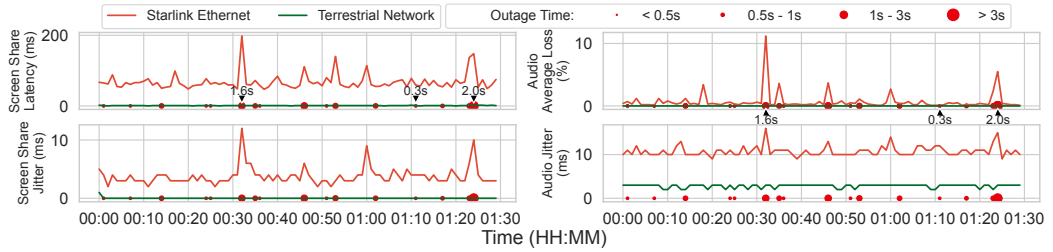

Fig. 5. Screen share and audio statistics in a Zoom meeting, with red dots denoting Starlink network outages.

4.3 Videoconferencing Service

To measure the network statistics of the Zoom platform, we utilized Zoom's *Zoom Meeting API* [55]. The API enables us to retrieve the per-minute statistical data on network metrics such as average latency, jitter, and average loss for both audio and video. We collected totally over 60 hours of Zoom meeting data during our measurements, where we observed that the bitrate and frames per second (FPS) of both video and audio fluctuate over time for both the Starlink and terrestrial networks without any apparent pattern, with the average frame rate and bitrate for Starlink Ethernet as 11.92 FPS and 1437.97 kbps, respectively, while the corresponding values for Terrestrial Ethernet were 11.84 FPS and 1415.06 kbps. Additionally, we observed that screen share latency and video frame rate showed a similar network pattern. As a result, our analysis will primarily focus on screen share latency/jitter and audio average loss. Another observation is that Zoom optimizes its performance based on the participants' network configuration. For instance, if both participants are on the same local area network (LAN), Zoom will use peer-to-peer (P2P) by default. To avoid this, we make sure that all participants are not on the same LAN. Furthermore, our measurements on the cases of two and three participants show no significant differences in network performance. This is because, according to Zoom's documentation [54], all participants in a Zoom meeting will communicate with a Zoom server during the meeting time, so the network statistics for each user should be independent of the number of users in the meeting. To ensure consistency across our measurements, we used Zoom's screen share function rather than a video camera. This allows us to display the same content during each Zoom meeting session, facilitating better analysis of the performance statistics. The shared content was a pre-recorded 2k video with a frame rate of 60 FPS, a video data rate of 9001 kbps, and an audio bit rate of 126 kbps. As videoconferencing applications have high demands for internet connection stability, we ensure that all participants are connected to the network via Ethernet cables during Zoom meetings.

4.3.1 Screen Share from One User. To gain a better understanding of the Starlink network's performance during typical Zoom meetings, we conducted experiments focusing on two key scenarios: first, when a Starlink user shares their screen, and second, when the Starlink user views screen share content from another participant. Fig. 5 shows the measurement results of a typical Zoom meeting with two participants, with one participant video sharing using Terrestrial Ethernet and the other participant watching via the Starlink Ethernet (dish A). It is evident that Terrestrial Ethernet is stable and low for both audio and video latency/jitter. In contrast, the Starlink network exhibits significant variations, particularly during network outages lasting longer than 1 second. However, Starlink network outages shorter than 0.5 seconds have marginal impacts on the Zoom meeting performance. Furthermore, we can see that unlike screen share latency, the audio loss rate is highly correlated with Starlink network outages. The average audio loss rate shown on the top right of Fig. 5 remains consistent most of the time, but whenever a Starlink network outage happens,

there is an increase in the average audio loss rate, with longer network outages yielding higher average audio loss rates. In addition to the above measurement, we also conduct an experiment in which a participant using a Starlink Ethernet shares the screen with the participant using Terrestrial Ethernet watching the screen share. We observe that the overall pattern of performance remains consistent with our previous findings, indicating that both downlink and uplink are sensitive to Starlink outages. Furthermore, we find that the average latency for Starlink Ethernet's uplink is 61 ms, while the average latency for its downlink is 45 ms. This phenomenon may be attributed to limitations in the user's dish antenna size and power. Overall, Starlink network is capable of handling realtime videoconferencing traffic, but users may experience interruptions in screen sharing and audio cut-offs during long network outages.

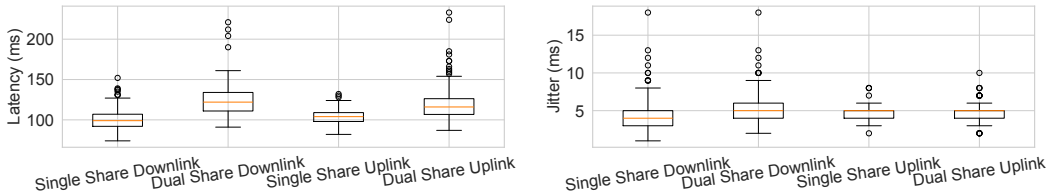

Fig. 6. Latency and jitter statistics of Starlink users in a 2-participant meeting. Single share means only one user shares a screen, and dual share means both users share screens.

4.3.2 Screen Share from Multiple Users. In a typical Zoom meeting, participants can only view one person's screen share content at a time. If a participant is sharing their screen, they cannot simultaneously view someone else's screen share content. However, Zoom offers a feature for users with multiple monitors that allows them to share their screen while also watching another participant's screen share content. We conducted an experiment to test this feature, as we believe it can provide insight into how the Starlink network performs when both the uplink and downlink are being utilized simultaneously and how this will affect users' realtime videoconferencing experience. During the experiment, we had two Starlink users share their screens while simultaneously viewing each other's screen share content. Our results show that the Starlink network can still provide decent network quality when both the downlink and uplink are utilized at the same time during a Zoom meeting. However, as shown in Fig. 6, there is an increase in network latency for users in a dual screen share session for both downlink and uplink. We conjecture that this is also limited by the User's dish power supply or antenna size. So the dish can only send or receive a constrained amount of data simultaneously.

Table 4. Statistics of user interactions.

	Terrestrial Ethernet	Starlink Ethernet
# of interactions	1136	869
Average of RTT (s)	0.57	0.69
Variance of RTT (σ^2)	0.0073	0.0095

4.3.3 User Interactions. Since videoconferencing services often involve many realtime user interactions, we also conducted an experiment to investigate the impacts of the Starlink network on such interactions. To better automate the experiment and make it more objective, each participant used two monitors, enabling them to share their screens while viewing each other's screen share content. A Python script is run on each participant to display a pure black or white image on one monitor while detecting the colour displayed on the other monitor. At the beginning, both participants

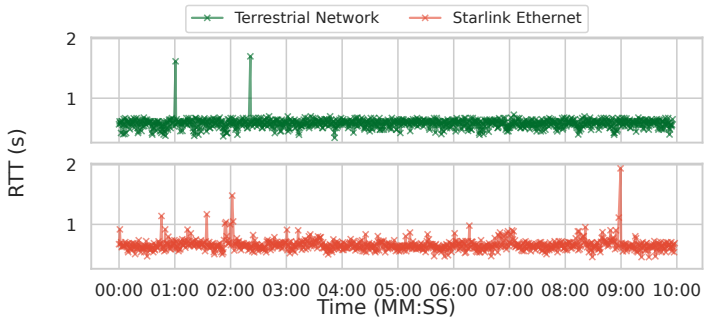

Fig. 7. RTT during a 10-minute interactive Zoom session

display white screens. Then participant A first changes the displayed color (e.g., from white to black). If participant B detects the colour change on participant A's screen share, participant B's changes its display color, and then participant A will do the same upon detecting color change from participant B. Both participants' actions are simulated and controlled by the Python script to mimic actual user interaction and provide accurate measurement results. Every time a participant detects a color change, the Python script records a timestamp, allowing us to analyze the latency on both sides. Before each experiment, we use Windows' sync time feature to ensure that both systems' system clocks are synced to the same server. The script also records the code execution overhead, such as the time used to switch the display or detect a color change, ensuring that the latency is as accurate as possible.

Fig. 7 compares the results of the user interaction experiment we conducted on both Terrestrial Network and Starlink Ethernet during a 10-minute session, where the RTT is defined as the duration between the moment when participant A's screen changes and the moment when participant A detects the screen change of participant B. It is evident that the Terrestrial Network has more stable interactions compared to the Starlink network. Table 4 further provides the statistics of interactions made via each network, with Terrestrial Network completing about 23.5% more interactions than the Starlink network. All these indicate that Zoom users with the Starlink network may have less fluent experiences on interactions.

In summary, compared with Terrestrial Network, although videoconferencing service over the Starlink network can still achieve reasonably good performance, it may still experience higher average latency and jitter as well as larger network variance. Moreover, compared to VoD and live streaming services, where the corresponding platforms can utilize buffers to mitigate the impact of network instabilities in Starlink, videoconferencing service is generally more sensitive to network variations. Therefore, users using Starlink for videoconferencing services may experience more disruptions, such as frame losses and audio cut-offs.

5 SOME KEY IMPACT FACTORS ON LSN AND ITS MULTIMEDIA SERVICES

5.1 Satellite Switching

A distinctive feature of LSN, setting it apart from other network systems, is the periodic satellite switching, which can result in significant fluctuations in latency or bursts of packet loss [10, 36, 40]. The fundamental cause is that the LEO satellites are not geo-synchronized with the Earth, which causes the relative locations of the LEO satellites to change over time even in a relatively short period. As illustrated in Fig. 1, where a dish (e.g., dish A in the figure) will switch its connection from one satellite (e.g., the right satellite currently connecting to dish A) to another (e.g., the left satellite where dish A will handover next) because the former may move too far away to communicate with.

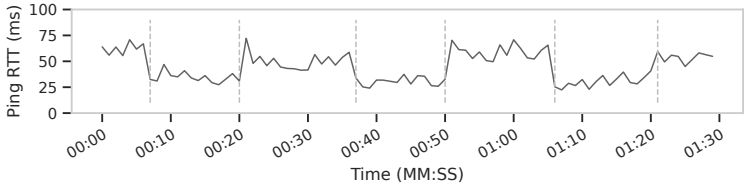

Fig. 8. A typical pattern for *ping* RTTs, where the interval for each steady state is indicated by vertical lines.

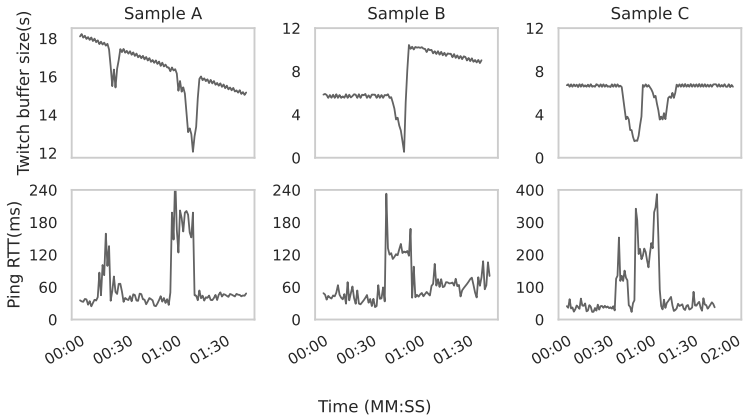

Fig. 9. Impact of large jitters in observed patterns on Twitch's buffer size.

Impact on Latency. To investigate the impact of satellite switching on multimedia services, we utilize the *ping* command to collect the RTT time to the ground station (GS), and the GS's private IP address is 100.64.0.1 from the User's dish viewpoint [36]. The *ping* packet is sent with a one-second interval, and we eventually collect 887,628 data points from both dish A and dish B.

Our observation reveals a pattern where the RTT often changes from a steady state to another steady state with only small fluctuations. Fig. 8 shows a typical example of this pattern, where starting from time *00:07*, the RTT remains steady with only small fluctuations for approximately 15 seconds and then switches to another steady state, and so on and so forth. We also notice that sometimes the average RTT is nearly the same for every other state, which may indicate that the connection is oscillating between two satellites. The switching frequency is consistent with the 15-second satellite reallocation interval observed in [40]. In addition, the RTT difference between each steady state in Fig. 8 is approximately 20 ms, which can increase further if a handover fails or when the UE switches to a satellite that is far from the UE. For instance, the average network jitter we observed is around 10.65 ± 0.08 ms. However, during satellite reallocation, the average jitter doubles to 21.01 ± 0.05 ms, with a 20% probability of exceeding 30.70 ± 0.20 ms. Such large jitters can significantly impact multimedia services, causing performance degradation. Fig. 9 shows three Twitch buffer health examples with corresponding synchronized *ping* RTTs, where buffer drops can be easily observed when the RTT changes to the next state. For example, in sample A, the RTT increases to 180ms at *01:00* and leads to a 5-second buffer drop. Similar consequences can also be observed from samples B and C. Interestingly, during these periods of time, the Starlink mobile application does not report any network issues or outages, indicating that this could also be one reason for the non-outage-caused buffer drops in Fig. 4.

Impact on Packet Loss. As live broadcasting and videoconferencing platforms like Twitch, TikTok, and Zoom continue to grow in popularity, there is an increasing demand from users for high-bandwidth and low-latency uplink capabilities to enable smooth, realtime video uploading. While the current uplink bandwidth of LSN is generally adequate to support realtime video uploading of

1080p video at 60 FPS, it still encounters frequent frame losses, as discussed in the Zoom section. Furthermore, frame loss events often result in significant degradation of QoE, and we notice that most frame losses are attributed to periodic satellite handover. This observation underscores the need to closely examine packet loss patterns during satellite handover periods, specifically in the context of video uploading, and to explore potential solutions to mitigate these losses.

Tanveer et al. [40] conducted an investigation into satellite handover scheduling by extensively sending packets (50 packets per second) using *iRTT* [17]. Their findings reveal that a global network controller will allocate a suitable satellite to each UE every 15 seconds, with allocations occurring at the 12th, 27th, 42nd, and 57th seconds of every minute. It's important to note that handovers do not occur with every allocation. The controller decides to re-allocate if the current link is deemed suboptimal, taking into account factors such as satellite load, position, and user priority. Through this exploration of the Starlink handover schedule, we can precisely assess how satellite handovers impact realtime multimedia services at the packet level. To examine packet level behaviors, we employ the *Ringmaster* platform designed for benchmarking videoconferencing applications. *Ringmaster* facilitates the emulation of videoconferencing scenarios between a server and clients by establishing a one-to-one video call over UDP. To assess the real impact of packet loss, we deactivate Forward Error Correction, allowing unacknowledged packets to be retransmitted up to three times. In the videoconferencing scenario, if any video segments fail to arrive within the suggested playout latency of 1 second according to the guidelines outlined in *RFC 9317* [18], we mark the corresponding frame as lost and skip it.

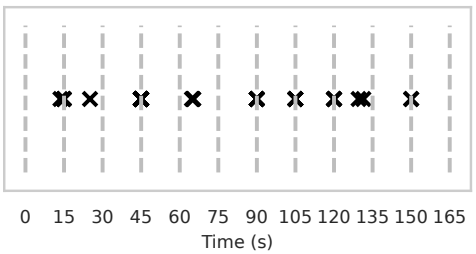

Fig. 10. Type A packet loss: The packet loss burst during the handover period.

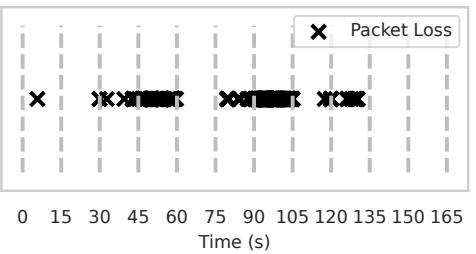

Fig. 11. Type B packet loss: The packet loss burst between two handover periods.

Using the same experiment setup, the computer connected to Starlink Ethernet serves as the sender, while the computer connected to the Terrestrial Network functions as the receiver. We observe noticeable packet losses at the satellite reallocation period (refer to Fig. 10), where the dotted vertical line represents the timestamp of reallocation. Some scattered packet losses occur between these periods, likely due to network congestion or unstable wireless communication, as there is no significant correlation between these scattered losses and the satellite reallocation periods. Furthermore, we notice intensive bursts of packet loss occurring between two satellite reallocation periods (see Fig. 11). We infer that this phenomenon is likely due to the UE being linked to an overloaded satellite or being obstructed, resulting in subpar network conditions for the 15-second window until the next handover. During our 3-hour experiment, we recorded a total of 1547 packet loss events[3]. Notably, 284 of these events occurred precisely during the handover period. The probability of observing a packet loss event for each second is $14.32\%(\frac{1547}{3\times60\times60})$, and the probability of observing a packet loss event during the handover period is $39.44\%(\frac{284}{3\times60\times4})$[4]. Such a marked deviation suggests that packet loss is more likely to occur during the handover period.

[3]We define a packet loss event as any instance where packet loss is observed within a one-second interval.
[4]There are four satellite reallocation events per minute.

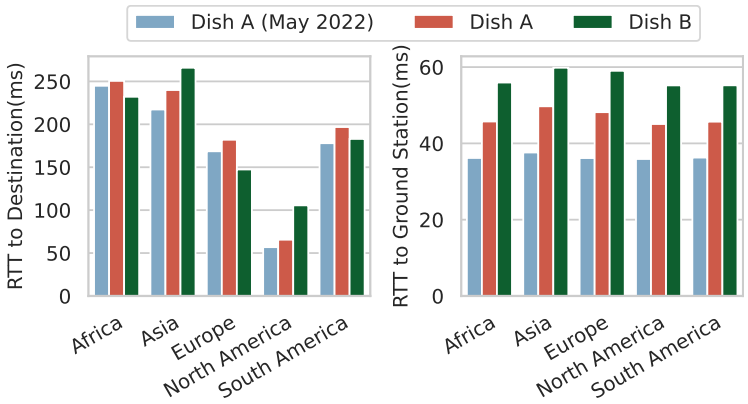

Fig. 12. *traceroute* RTTs between different destinations.

Concerning frame losses, we notice that only 9.18% of frame losses are attributed to Type A packet loss (Fig. 10), while 41.84% of frame losses are associated with Type B packet loss (Fig. 11). It is essential to highlight that we implement a stringent threshold during the identification of Type B packet loss to minimize the risk of false positives. Consequently, the actual real-world impact may be even more substantial than the reported 41.84%. Type A packet loss is more frequent, involving the loss of a few sequential packets, but most lost packets can be retransmitted and acknowledged promptly. On the other hand, Type B packet loss persists for a longer duration, leading to persistent delays in the delivery of some packets, as all retransmitted packets may be lost.

With the capability to anticipate upcoming satellite handovers, it becomes viable to mitigate the impact of Type A and B packet loss in multimedia services. For example, if the server encounters a series of packet losses immediately after the satellite handover period, it may opt to reduce the video bitrate and enhance the number of retransmitted packets to address the potential occurrence of Type B packet loss. Moreover, the client can periodically identify and mark satellites as defective or overloaded based on historical packet loss data. Additionally, it can forecast impending Type B packet losses by integrating satellite trajectory data obtained from Two-Line Element sets [47].

5.2 Starlink Routing Strategy

During early last year before starting this measurement work on multimedia services, we also did some preliminary measurements on the Starlink network. At that time, we found that the Starlink network will connect to the GS geographically closest to the dish, which also matches what has been reported in [26].

However, we notice that our measurement results in this work show a different routing behavior, where all the packets are directed to a GS that may be geographically far away from our dish. This might indicate that *SpaceX* has changed their routing strategy or started to use *Inter-Satellite Links* (ISLs) since later 2022, especially considering that some Starlink users have reported that they received the ISL service enabling announcement by the end of 2022 [51]. ISLs enable direct communication between satellites using laser beams without relying on GS as intermediaries. Consequently, satellites can relay user traffic to distant GS, facilitating long-range communication (see Fig. 1). This approach is particularly promising for users who are geographically distant from GS, such as those on islands or cargo ships. Furthermore, the speed-of-light advantage of ISL potentially enables faster intercontinental communication compared to submarine optical cables, as light travels faster through the vacuum of space than through optical cables. Yet, given the ISL is still not fully supported by all satellites and communication efficiency may degrade after traversing multiple satellites, evaluating the performance of this emerging routing strategy is essential. We

Table 5. Average bandwidth of different transport protocols and congestion control algorithms.

	TCP Reno	TCP Cubic	UDP
Average Bandwidth (Mbps)	52.11	60.16	205.38
Standard Deviation (Mbps)	28.50	36.43	77.10
Max (Mbps)	275.00	340	316

compare our current *traceroute* data with the previous *traceroute* data collected in May 2022 on dish A, and plot the RTTs to different continents in Fig. 12 with outlier data[5] removed. It is apparent that the current RTT to each destination is larger than the RTT observed early last year. Furthermore, we can see a similar to-ground-station RTT increment for all the destinations, which indicates that Starlink may have used a different routing strategy from early last year. Currently, the new routing strategy adds more latency on the path to GS, which may be caused by ISL, as additional overhead such as the processing delay and queuing delay introduced by multiple satellites.

5.3 Starlink Network Performance During Peak Traffic Hours

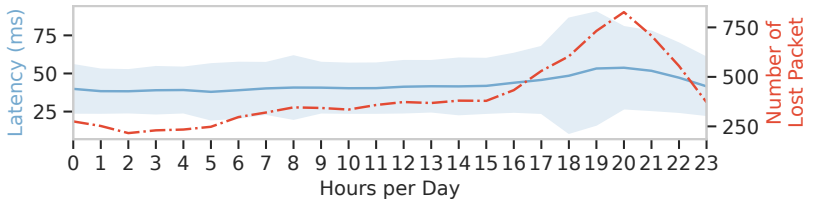

Fig. 13. Correlation between Starlink network performance and internet rush hours, with the blue area indicating the standard deviation of latency.

By extending our data analysis over a longer period, we confirmed a clear correlation between Starlink's network performance and Internet rush hours. Using the same dataset from Section 5.1, we segment all ping data into 24-hour segments, analyzing average latency and packet loss as illustrated in Fig. 13. Our findings show substantial increases in latency, jitter, and packet loss, with 24.30%, 72.33%, and 106.08% respectively, during the hours of 15:00 to 22:00, coinciding with peak Internet usage times [14]. This pattern is also observed with Dish B and Dish C, which are deployed in rural areas, though the performance degradation is less severe, likely due to lower Starlink user density in these regions. Hence, the observed performance degradation is likely tied to traffic congestion caused by increased data flow. Unlike Terrestrial Networks, where urban areas can alleviate rush hour traffic by building more data centers, the LEO satellites are constantly moving, making it challenging to concentrate resources over high-traffic areas. Despite ongoing efforts by Starlink to launch additional satellites, the exponentially increasing number of subscribers and the limited number of Ku-band antennas that LEO satellites can equip suggest that the oversaturated traffic load is likely to persist in the near future.

5.4 Transport Protocols and Congestion Control Algorithms

To better understand the LSN's network performance, we also conduct lower-level measurements, including transport protocols and congestion control algorithms originally designed for Terrestrial Networks. We test both TCP and UDP protocols on Starlink using *iPerf3*. For TCP, we also evaluate the TCP Reno and TCP Cubic congestion control algorithms. The results are shown in Table 5. For

[5]The data falling above 75th- and below 25th-percentile considered as a outlier.

TCP, the TCP Reno congestion control algorithm exhibits lower performance compared to TCP Cubic. We anticipate this is due to TCP Reno's more drastic decrease in congestion window size when packet loss is detected, which is quite frequent in LSNs. On the other hand, UDP has the highest average bandwidth since it does not have a congestion control mechanism. Our measurement results suggest that current transport protocols cannot fully utilize LSNs' performance due to their unique network characteristics. This observation aligns with existing works, such as [10], which were motivated by these findings and have made efforts to modify lower-level transport protocols to improve network performance on LSNs.

5.5 Weather Impact on Starlink

Table 6. Correlation between weather conditions and each Starlink network outage duration, with *Average Outage Duration* representing the mean length of each outage and *Outage Probability* representing the likelihood of an outage occurring every minute.

Weather	Average Outage Duration (s)	Standard Deviation (s)	Outage Probability (%)
Clear/Cloudy	2.07	17.49	0.09
Rain	12.91	48.69	1.27
Drizzle	20.36	51.16	1.33
Thunderstorm	67.00	300.42	11.67

We conduct an extensive study to evaluate the impact of weather on Starlink network outages. This involves tracking outages via the Starlink application while concurrently gathering corresponding weather statistics through the *Weather API* from OpenWeather [3]. The comprehensive statistics linking Starlink outages with various weather conditions are detailed in Table 6. In meteorological terms, *Drizzle* refers to rain with drop diameters between 0.2 and 0.5 mm, and *Rain* is defined as raindrops exceeding 0.5 mm in diameter [7]. Our research findings distinctly underscore the significant impact of weather on the performance of the Starlink network. The data clearly indicates that network outages are notably more frequent and last longer during periods of rainy weather, in contrast to clear or cloudy weather conditions. It's important to note that due to the insufficient occurrence of snow days during our measurement period, snow-related data was not extensively summarized in our study.

One particularly striking observation is the pronounced impact of a thunderstorm on network performance. This is most evident with dish B, which experienced frequent and extended outages during the thunderstorm period. For example, Fig. 14 showcases a 12-hour outage comparison between a day with a thunderstorm and a clear day. This comparison distinctly highlights a significant network disruption, especially between 5 pm and 6 pm during the day of the measurement, aligning with reports of intense thunderstorm activity. Furthermore, Fig. 15 provides a contrast in the

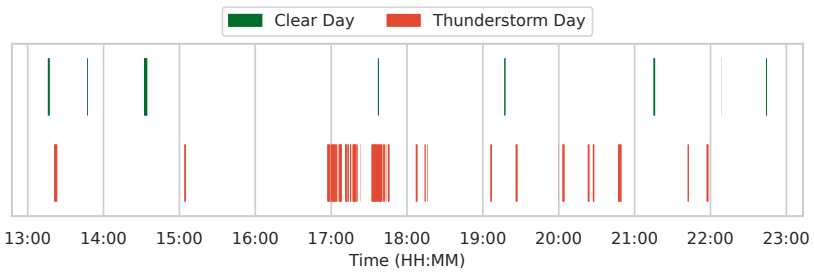

Fig. 14. A 12-hour outage history, where outages lasting less than 2 seconds are multiplied by a scaler of 80 for better visualization.

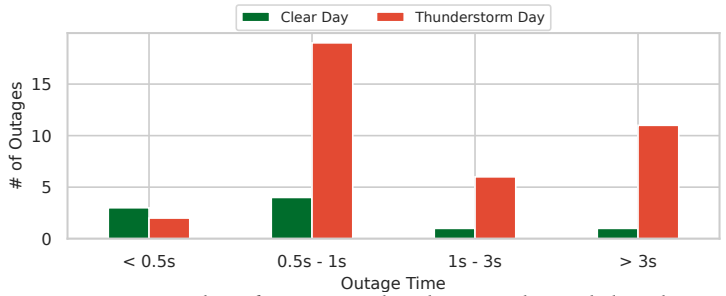

Fig. 15. Number of outages in thunderstorm day and clear day.

number and duration of outages between a thunderstorm and a clear day. Such prolonged outages have a substantial impact on multimedia services. Zoom meetings may experience disruptions, and streaming on platforms like YouTube or Twitch could be paused or become unavailable. This underscores the challenges faced by users in regions prone to frequent thunderstorms, such as West Central Florida in the US, which experiences over 100 thunderstorms annually, surpassing other areas in the country[6]. These users may encounter notably worse network performance when utilizing LSNs.

6 WEATHER AWARE BUFFER BASED RATE ADAPTION

As discussed earlier, our comprehensive measurements have systematically investigated the performance of typical multimedia services over LSNs, as well as revealed some key impact factors therein. These informative discoveries can provide valuable insights for future development in both multimedia services and LSNs. As a demonstration along this direction, in this section, we further develop and evaluate a new ABR algorithm for VoD services that are motivated by our insights, especially on the impact of weather conditions. This algorithm can intelligently adjust the bitrate and maximum buffer size based on the current weather conditions to improve user QoE. We first analyze the key factors influencing the QoE of VoD applications and define the Buffer-Protection based Bitrate Selection (BPBS) problem. We then introduce our WABB algorithm, which is designed to address the BPBS problem effectively.

6.1 Problem Formulation

In the VoD scenario, the client side downloads video segments at a bitrate determined by the ABR algorithm. These segments are transmitted via LSNs and stored in the local buffer, awaiting playback by the video player. A crucial aspect of maintaining a high QoE is avoiding an empty buffer, as this would cause pauses in the video stream, significantly degrading the user's viewing experience. Therefore, a key strategy in enhancing QoE is to ensure that the buffer consistently has a sufficient number of segments for uninterrupted playback. This approach is particularly effective in mitigating the impact of short-term network outages, which can be caused by adverse weather conditions and satellite handovers, thereby preventing users from experiencing any noticeable disruptions in service.

We use the discrete-time model to split the time into a set of intervals, denoted as $T = \{t\}$. Given an interval $t \in T$, we define the starting and end time as t_s and t_e, respectively, and the length of each interval equals the video segment length. Based on the above assumptions, we define the outage events during t as a set of tuples as $O(t) = \{(\xi_o, \zeta_o)\}$, where ξ_o is the time when the outage happens and ζ_o is the corresponding duration. $O(t)$ is a function that takes the current weather

[6]https://www.weather.gov/key/tstmhazards

condition $h_{weather}(t_s)$ and handover behavior $g_{handover}(t_s)$ as inputs. For each time interval t, we use the weather condition and handover behavior at the starting point t_s to represent their status, as they are expected to exhibit consistency within such a brief period (the typical segment length is less than 10 seconds [37]). Formally, we have

$$O(t) : h_{weather}(t_s) \times g_{handover}(t_s) \rightarrow \mathbb{R}^* \times \mathbb{R}^*, \tag{1}$$

where $\mathbb{R}^*$ is the set of non-negative real numbers.

We assume that the video has m available bitrates $\mathcal{K} = \{k_i | i = 1, 2, \ldots, m\}$ sorted with an ascending order, and the selected bitrate for each interval t is donated as $r(t)$. Hence, for each time interval t, we define the buffer occupancy ratio at time τ as $R(\tau)$, where $\tau \in [t_s, t_e]$. The buffer occupancy ratio represents the amount of occupied buffer to the maximum buffer size, and $R(\tau) = 0$ represents that after the playback completion of the current segment, the buffer runs dry, and the next segment is not yet prepared, leading to rebuffering. $R(\tau)$ is a function that takes the initial buffer length $I(t_s)$, the outage events $O(t)$, the consumed buffer $C(\tau - t_s)$ and selected bitrate $r(t)$ as inputs. We have:

$$R(\tau) : I(t_s) \times O(t) \times C(\tau - t_s) \times r(t) \rightarrow [0, 1], \tag{2}$$

where $R(\tau)$ returns a non-negative real number ranging from 0 to 1. In the first video segment, we assume that the buffer occupancy is not zero since applications will warm up the buffer by downloading some segments in advance.

Based on the above definitions, we formulate the BPBS problem as an optimization problem. The aim is to maximize the utility, which is a weighted sum of the occurrences of the buffer becoming empty and bitrate $r(t + 1)$ within each interval. We have

$$\textbf{maximize} \sum\nolimits_{t \in T} w_1 \int_{t_s}^{t_e} \mathbb{1}_{R(\tau)!=0} d\tau + (1 - w_1) r(t + 1), \tag{3}$$

$$s.t. \quad (1) - (2),$$

$$w_1 \in (0, 1), \tag{4}$$

where $\mathbb{1}_X$ is a function that returns 1 if X holds true and 0 otherwise. We use w_1 to represent the weight, which is a non-negative real number ranging from 0 to 1. This utility function ensures that the WABB will minimize the occurrences of rebuffering and, meanwhile, maximize the total playback bitrate.

6.2 Proposed Algorithm

Algorithm 1 WABB Algorithm

1: **Input:** $\{h_{weather}(t_s) | \forall t \in T\}, \{g_{handover}(t_s) | \forall t \in T\}, \{R(\tau) | \forall \tau \in t\}$
2: **Output:** $\{r(t + 1) | \forall t + 1 \in T\}$
3: Determine the maximum buffer size B_m based on the available resources.
4: Initialize the set of target buffer size $\mathcal{L} = \{L_i | i = 1, 2, \ldots, n\}$.
5: **for** $t \in T$ **do**
6: Initialize t_s and t_e as the starting and end time of t.
7: Update $h_{weather}(t_s), g_{handover}(t_s) \leftarrow t_s$
8: $L_i \leftarrow$ Get the type by the current weather condition $h_{weather}(t_s)$.
9: $r(t + 1) = \mathcal{M}(\frac{B_m R(t_e)}{L_i})$.
10: **end for**

The WABB algorithm operates by considering three key factors: current weather conditions, handover behavior and the buffer occupancy ratio $R(\tau)$. The current realtime weather conditions can be obtained using the OpenWeather API, the detailed handover behaviour for Starlink, as discussed in Section 5.1, is already known, and the buffer occupancy ratio can also be directly obtained. As shown in Algorithm 1, we set different maximum buffer sizes and target buffer sizes based on available memory usage and current weather so that our method can be adaptive to different types of devices and weather conditions. The algorithm then initializes n types of weather conditions aligned with a set of predefined target buffer size $\mathcal{L} = \{L_i | i = 1, 2, \ldots, n\}$, where the L_i is a real number ranging from 0 to B_m. For each interval t, the L_i is selected from the predefined n types of target buffer size based on current weather condition $h_{weather}(t_s)$. The optimal bitrate for the current interval t is determined by the function $\mathcal{M}$, as shown below:

$$\mathcal{M}(x) = \underset{k \in \mathcal{K}}{\operatorname{argmin}} |xk_m - k|, \tag{5}$$

In general, the next bitrate $r(t+1)$ will be determined by x, and x is a ratio of the current buffer size ($B_m R(t_e)$) to the target buffer size (L_i). This ratio will then be mapped to a bitrate that can approach the target buffer size while also maintaining the highest possible bitrate. In Equation 5, we incorporate the highest bitrate k_m as the multiplication factor to promote the selection of the highest bitrate when the current buffer size is close to the target buffer size, which is aligned with the optimization goal in Eq.3. In our experiment, we observe that using the highest bitrate multiplication factor also avoids frequent bitrate changes when the buffer size is saturated. For example, without this approach, bitrates might fluctuate between k_m and k_{m-1}, particularly if a lower multiplication factor is selected.

6.3 Experimental Setup and Performance Evaluation

Due to the practical challenges of replicating diverse weather conditions for real-world experiments, we opt to evaluate the performance of our algorithm in a simulated environment. For this purpose, we used the *Sabre* video streaming simulator, renowned for its ability to accurately replicate adaptive streaming environments [38]. This approach allows us to conduct comprehensive experiments under controlled conditions, effectively simulating a range of weather scenarios that would be challenging to reproduce in field tests. We generate 100 network traces for each weather type, drawing on data from our measurement datasets. The video used in the simulation is a 3-hour movie from *Sabre's* dataset, encoded in standard definition as per the guidelines of [39]. It features four different bitrates: 1000, 2500, 5000, 8000 Kbps, with each segment lasting 2000 ms. To evaluate the effectiveness of our proposed algorithm, we focus on three key metrics: average rebuffering time, average bitrate, and the number of bitrate changes. The objective is to minimize rebuffering time and the frequency of bitrate changes while maintaining a higher average bitrate to ensure the best possible viewing experience.

To comprehensively evaluate the effectiveness of our proposed algorithm, we conducted a comparison with the following baseline algorithms:

- RobustMPC [50]: RobustMPC incorporates the maximum prediction error from previous chunks as boundaries. This enhances its resilience to uncertainties and variability, optimizing the selection of bitrates for improved QoE.
- Buffer Occupancy based Lyapunov Algorithm (BOLA) [39]: Utilizing Lyapunov optimization, BOLA aims to minimize rebuffering while maximizing video quality. It operates independently of network bandwidth predictions, providing near-optimal performance.

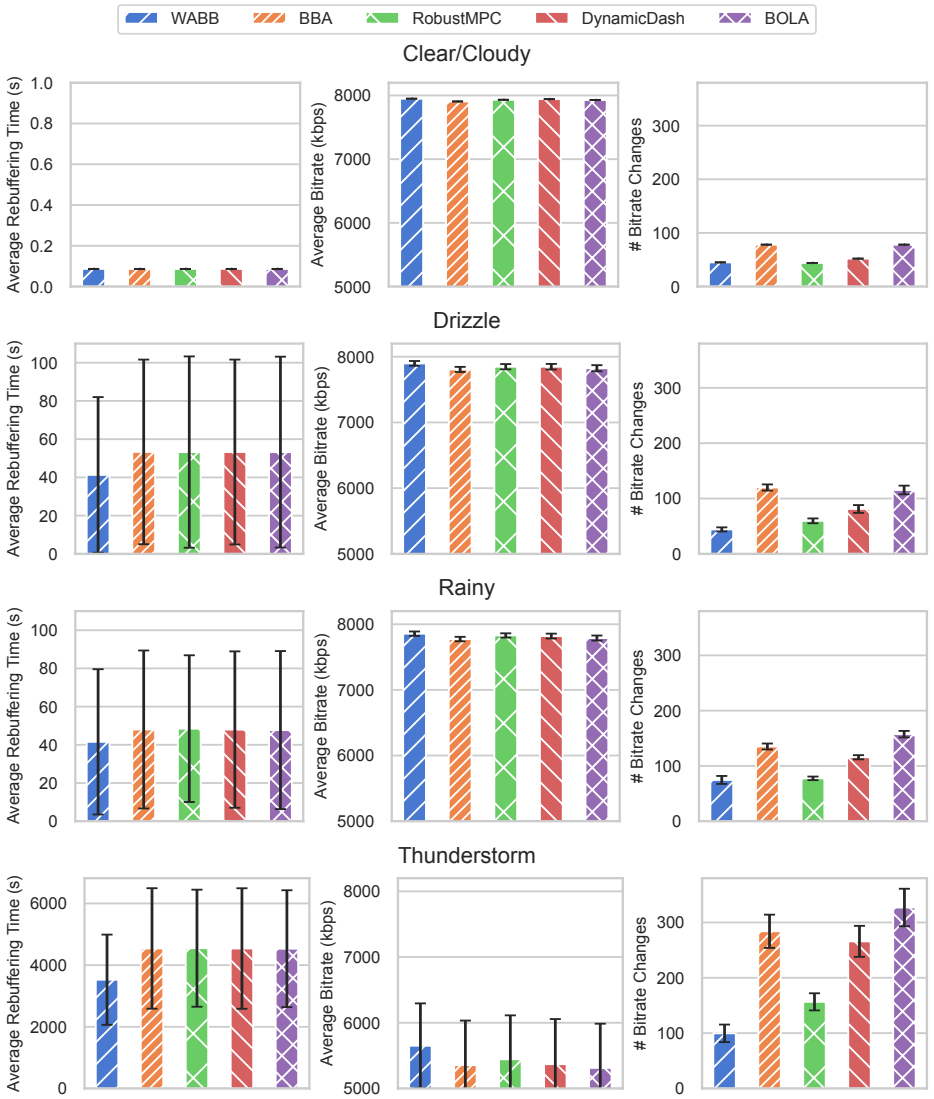

Fig. 16. Performance evaluation of different ABR algorithms under different weather conditions.

- DynamicDash[7]: This ABR strategy alternates between throughput-based ABR and BOLA depending on buffer levels, combining the strengths of both approaches. It aims to optimize the balance between video quality and buffering time.
- Buffer Based Rate Adaption (BBA) [19]: This algorithm uses a 15-second buffer (5 seconds of reservoir and 10 seconds of cushion) to maintain buffer occupancy and manage bitrate selections, providing a cushion against network disruptions.

Fig. 16 presents the results of different ABR algorithms under various weather conditions. In clear and cloudy weather, the Starlink network effectively meets the demands of typical VoD services. It maintains an average rebuffering event duration of approximately 0.08 seconds alongside a high average bitrate. The superiority of our proposed algorithm is particularly noticeable during drizzle

[7]https://reference.dashif.org/dash.js/

and rainy weather. In these conditions, our algorithm outperforms others by reducing the average rebuffering time by around 20% while still allowing for a high bitrate and fewer bitrate changes. During thunderstorm conditions, which pose a challenge to all ABR algorithms due to prolonged and frequent Starlink network outages, WABB demonstrates exceptional performance and offers the highest overall QoE. It achieves a 22.5% reduction in rebuffering time and records the fewest bitrate changes, significantly enhancing user QoE. It is worth noting that there is a significant standard deviation observed in the average rebuffering time. This is primarily due to the wide range of outage durations observed in the Starlink network, as detailed in Section 5.3. These outages vary from as brief as 0.2 seconds to as long as 23.57 seconds. Such a broad spectrum of outage durations inevitably leads to considerable fluctuations in the average rebuffering times. In addition, the rebuffering times of other baseline algorithms were similar to each other. This similarity is attributed to the fact that during longer Starlink outages, all selected ABR algorithms struggle to cope with such extensive disruptions, leading to an empty buffer and consequent playback pauses.

7 CONCLUSION

This paper presents a systematical measurement and analysis to understand the performance of multimedia services over the LEO satellite networks, with the Starlink network as a case study. Our six-month measurement covers VoD, live streaming and videoconferencing services, with data collected via three Starlink dishes located in totally different geographical areas. Our findings show that the Starlink network can generally provide reasonable accommodations to support multimedia services, although the performance may degrade due to such factors as extreme weather, satellite switching and routing path changes. Moreover, such performance degradation may have different impacts depending on the multimedia service type, with VoD impacted the least, followed by live streaming and videoconferencing mostly impacted due to its realtime two-way user interactions. These insights are invaluable for guiding future developments in both multimedia services and LSNs. To demonstrate this, we introduce the WABB algorithm that is able to maximize the QoE by solving the BPBS problem. We also evaluate our ABR's effectiveness, demonstrating its superior performance compared to other baseline ABR algorithms. In the future, we aim to improve our ABR algorithms for live streaming and videoconferencing services, which are more challenging due to the need to account for user interactions and latency. As we better understand LSN's unique network characteristics, we could utilize satellite handover patterns and weather data to forecast network outages and take proactive measures to mitigate them.

ACKNOWLEDGMENTS

We appreciate the constructive comments from the reviewers. This research is supported by an NSERC Discovery Grant, a British Columbia Salmon Recovery and Innovation Fund (BC-SRIF_2022_401), a MITACS Accelerate Cluster Grant, and partly supported by an NSF I/UCRC Grant (1822104).

REFERENCES

[1] YouTube Video Trace (Outdoor), 2023. Retrieved June 6th, 2024 from https://youtu.be/Ac07Qt84WDw.
[2] YouTube Video Trace (Sport), 2023. Retrieved June 6th, 2024 from https://www.youtube.com/watch?v=vIZo_ndXc-A&ab_channel.
[3] Open Weather Map, 2024. Retrieved June 6th, 2024 from https://openweathermap.org/current.
[4] Twitch Streaming Channel, 2024. Retrieved June 6th, 2024 from https://www.twitch.tv/pcentinela.
[5] YouTube Video Trace (Chatting), 2024. Retrieved June 6th, 2024 from https://www.youtube.com/watch?v=TXtnFw9eDmM&ab_channel=ConnerArdman.
[6] YouTube Video Trace (Gaming), 2024. Retrieved June 6th, 2024 from https://www.youtube.com/watch?v=UPtoIy-oPWQ&ab_channel=Throneful.

[7] American Meteorological Society. Raindrop, 2024. Retrieved June 6th, 2024 from https://glossary.ametsoc.org/wiki/Raindrop.

[8] Bentaleb, A., Akcay, M. N., Lim, M., Begen, A. C., and Zimmermann, R. Catching the moment with lol+ in twitch-like low-latency live streaming platforms. *IEEE Transactions on Multimedia (TMM'21) 24* (2021), 2300–2314.

[9] Bieringa, R., Radhakrishnan, A., Singh, T., Vos, S., Donkervliet, J., and Iosup, A. An empirical evaluation of the performance of video conferencing systems. In *Companion of the ACM/SPEC International Conference on Performance Engineering (ICPE'23)* (2021), pp. 65–71.

[10] Cao, X., and Zhang, X. SaTCP: Link-layer informed TCP adaptation for highly dynamic LEO satellite networks. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM'23)* (2023).

[11] Chang, H., Varvello, M., Hao, F., and Mukherjee, S. Can you see me now? a measurement study of zoom, webex, and meet. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC'21)* (2021), pp. 216–228.

[12] Cheng, Y., Zhang, Z., Li, H., Arapin, A., Zhang, Y., Zhang, Q., Liu, Y., Du, K., Zhang, X., Yan, F. Y., et al. GRACE:Loss-Resilient Real-Time Video through Neural Codecs. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI'24)* (2024).

[13] Choi, A., Karamollahi, M., Williamson, C., and Arlitt, M. Zoom session quality: A network-level view. In *International Conference on Passive and Active Network Measurement (PAM'22)* (2022), Springer, pp. 555–572.

[14] Cisco. New cisco study reveals peak internet traffic increases due to social networking and broadband video usage, 2009. Retrieved June 6th, 2024 from https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2009/m10/new-cisco-study-reveals-peak-internet-traffic-increases-due-to-social-networking-and-broadband-video-usage.html.

[15] Deri, L. Ntopng, 2024. Retrieved June 6th, 2024 from https://github.com/ntop/ntopng.

[16] Dobrian, F., Sekar, V., Awan, A., Stoica, I., Joseph, D., Ganjam, A., Zhan, J., and Zhang, H. Understanding the impact of video quality on user engagement. In *In Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM'11)* (2011), pp. 362–373.

[17] Heist, P. Isochronous Round-Trip Tester, 2024. Retrieved June 6th, 2024 from https://github.com/heistp/irtt.

[18] Holland, J., Begen, A. C., and Dawkins, S. Operational Considerations for Streaming Media, 2022. Retrieved June 6th, 2024 from https://www.rfc-editor.org/info/rfc9317.

[19] Huang, T.-Y., Johari, R., McKeown, N., Trunnell, M., and Watson, M. A buffer-based approach to rate adaptation: Evidence from a large video streaming service. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM'14)* (2014), pp. 187–198.

[20] Karagkioules, T., Mekuria, R., Griffioen, D., and Wagenaar, A. Online learning for low-latency adaptive streaming. In *Proceedings of the ACM Multimedia Systems Conference (MMSys'20)* (2020), pp. 315–320.

[21] Kassem, M. M., Raman, A., Perino, D., and Sastry, N. A browser-side view of starlink connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC'22)* (2022), pp. 151–158.

[22] Khalife, J., Neinavaie, M., and Kassas, Z. Z. The First Carrier Phase Tracking and Positioning Results With Starlink LEO Satellite Signals. *IEEE Transactions on Aerospace and Electronic Systems (AESS'22) 58*, 2 (Apr. 2022), 1487–1491.

[23] Li, Y., Zheng, Q., Zhang, Z., Chen, H., and Ma, Z. Improving ABR performance for short video streaming using multi-agent reinforcement learning with expert guidance. In *Proceedings of the ACM Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'23)* (2023), pp. 58–64.

[24] Lim, M., Akcay, M. N., Bentaleb, A., Begen, A. C., and Zimmermann, R. When they go high, we go low: low-latency live streaming in dash.js with LoL. In *Proceedings of the ACM Multimedia Systems Conference (MMSys'20)* (2020), pp. 321–326.

[25] Lin, X., Cioni, S., Charbit, G., Chuberre, N., Hellsten, S., and Boutillon, J.-F. On the Path to 6G: Embracing the Next Wave of Low Earth Orbit Satellite Access. *IEEE Communications Magazine (MCOM'21) 59*, 12 (2021), 36–42.

[26] Ma, S., Chou, Y. C., Zhao, H., Chen, L., Ma, X., and Liu, J. Network characteristics of leo satellite constellations: A starlink-based measurement from end users. In *In Proceedings of the IEEE Conference on Computer Communications (INFOCOM'23)* (2022), IEEE, pp. 1–10.

[27] MacMillan, K., Mangla, T., Saxon, J., and Feamster, N. Measuring the performance and network utilization of popular video conferencing applications. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC'21)* (2021), pp. 229–244.

[28] Mao, H., Netravali, R., and Alizadeh, M. Neural adaptive video streaming with pensieve. In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM'17)* (2017), pp. 197–210.

[29] Mao, Y., Wang, J., and Sheng, B. DAB: Dynamic and agile buffer-control for streaming videos on mobile devices. *Elsevier Procedia Computer Science 34* (2014), 384–391.

[30] McDowell, J. Jonathan's Space Pages, 2024. Retrieved June 6th, 2024 from https://planet4589.org/space/con/star/stats.html.

[31] Michel, F., Trevisan, M., Giordano, D., and Bonaventure, O. A first look at starlink performance. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC'22)* (2022), pp. 130–136.

[32] Michel, O., Sengupta, S., Kim, H., Netravali, R., and Rexford, J. Enabling passive measurement of zoom performance in production networks. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC'22)* (2022), pp. 244–260.
[33] Neinavaie, M., Khalife, J., and Kassas, Z. M. Acquisition, Doppler Tracking, and Positioning With Starlink LEO Satellites: First Results. *IEEE Transactions on Aerospace and Electronic Systems (AESS'22) 58*, 3 (2022), 2606–2610.
[34] Nightingale, J., Salva-Garcia, P., Calero, J. M. A., and Wang, Q. 5g-qoe: Qoe modelling for ultra-hd video streaming in 5g networks. *IEEE Transactions on Broadcasting (TBC'18) 64*, 2 (2018), 621–634.
[35] OneWeb. OneWeb Streamlines Constellation, 2021. Retrieved June 6th, 2024 from https://oneweb.net/resources/oneweb-streamlines-constellation.
[36] Pan, J., Zhao, J., and Cai, L. Measuring a low-earth-orbit satellite network. In *Proceedings of the IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'23)* (2023), pp. 1–6.
[37] Pantos, R., and May, W. HTTP Live Streaming, 2017. Retrieved June 6th, 2024 from https://www.rfc-editor.org/info/rfc8216.
[38] Spiteri, K., Sitaraman, R., and Sparacio, D. From theory to practice: Improving bitrate adaptation in the dash reference player. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM'19) 15*, 2s (2019), 1–29.
[39] Spiteri, K., Urgaonkar, R., and Sitaraman, R. K. Bola: Near-optimal bitrate adaptation for online videos. *IEEE/ACM transactions on networking (TON'20) 28*, 4 (2020), 1698–1711.
[40] Tanveer, H. B., Puchol, M., Singh, R., Bianchi, A., and Nithyanand, R. Making sense of constellations: Methodologies for understanding starlink's scheduling algorithms. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT'23)* (2023), pp. 37–43.
[41] Tereza Pultarova, E. H. Starlink satellites: Facts, tracking and impact on astronomy, 2024. Retrieved June 6th, 2024 from https://www.space.com/spacex-starlink-satellites.html.
[42] The Internet and Television Association. Report: Where Does the Majority of Internet Traffic Come From?, 2019. Retrieved June 6th, 2024 from https://www.ncta.com/whats-new/report-where-does-the-majority-of-internet-traffic-come.
[43] Tregloan-Reed, J., Otarola, A., Unda-Sanzana, E., Haeussler, B., Gaete, F., Colque, J., González-Fernández, C., Anais, J., Molina, V., González, R., et al. Optical-to-NIR magnitude measurements of the Starlink LEO Darksat satellite and effectiveness of the darkening treatment. *Astronomy & Astrophysics (A&A'21) 647* (2021), A54.
[44] Uitto, M., and Heikkinen, A. Evaluation of live video streaming performance for low latency use cases in 5g. In *2021 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit'21)* (2021), IEEE, pp. 431–436.
[45] Uran, C., Horvath, K., and Wöllik, H. Analysis of a starlink-based internet connection–roadmap-5g, 2021.
[46] Wang, F., Li, Q., Shi, W., Tyson, G., Jiang, Y., Ma, L., Zhang, P., Lan, Y., and Li, Z. Reparo: Qoe-aware live video streaming in low-rate networks by intelligent frame recovery. In *Proceedings of the 31st ACM International Conference on Multimedia (MM'23)* (2023).
[47] Wikipedia. Two-line element set, 2024. Retrieved June 6th, 2024 from https://en.wikipedia.org/wiki/Two-line_element_set.
[48] Xu, D., Zhou, A., Zhang, X., Wang, G., Liu, X., An, C., Shi, Y., Liu, L., and Ma, H. Understanding operational 5g: A first measurement study on its coverage, performance and energy consumption. In *Proceedings of the ACM Conference on Special Interest Group on Data Communication (SIGCOMM'20)* (2020), pp. 479–494.
[49] Yan, F. Y. Ringmaster, 2023. Retrieved June 6th, 2024 from https://github.com/microsoft/ringmaster.
[50] Yin, X., Jindal, A., Sekar, V., and Sinopoli, B. A control-theoretic approach for dynamic adaptive video streaming over HTTP. In *Proceedings of the ACM Conference on Special Interest Group on Data Communication (SIGCOMM'15)* (2015), pp. 325–338.
[51] Zafar, R. Starlink Turns On Laser Satellites For Region With Four Months Long Night, 2022. Retrieved June 6th, 2024 from https://wccftech.com/starlink-turns-on-laser-satellites-for-region-with-fourth-month-long-night/.
[52] Zhang, C., and Liu, J. On crowdsourced interactive live streaming: a twitch. tv-based measurement study. In *In Proceedings of the 25th ACM workshop on network and operating systems support for digital audio and video (NOSSDAV'15)* (2015), pp. 55–60.
[53] Zhao, H., Fang, H., Wang, F., and Liu, J. Realtime multimedia services over starlink: A reality check. In *Proceedings of the 33rd Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV'23)* (2023), pp. 43–49.
[54] Zoom Video Communications Inc. Zoom Whitepaper, 2020. Retrieved June 6th, 2024 from https://explore.zoom.us/docs/doc/Zoom%20Connection%20Process%20Whitepaper.pdf.
[55] Zoom Video Communications Inc. Zoom Developers, 2024. Retrieved June 6th, 2024 from https://developers.zoom.us/docs/api/rest/reference/zoom-api/methods/#overview.