Realtime Multimedia Services over Starlink: A Reality Check

Haoyuan Zhao[*]
Simon Fraser University
Burnaby, Canada
hza127@sfu.ca

Hao Fang[*]
Simon Fraser University
Burnaby, Canada
fanghaof@sfu.ca

Feng Wang
University of Mississippi
Oxford, United States
fwang@cs.olemiss.edu

Jiangchuan Liu
Simon Fraser University
Burnaby, Canada
jcliu@sfu.ca

ABSTRACT

Recently, Low Earth orbit Satellite Networking (LSN) has 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 operable LSN to date. There have been practical uses of Starlink with diverse networked applications, including multimedia applications of stringent demands. Given the mixed and inconsistent feedbacks from end users, it remains unclear whether today's LSNs, in particular, Starlink, have been 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 networking. Our findings demonstrate that Starlink can effectively handle most video-on-demand (VoD) and live-streaming services with properly configured buffers, but suffer from video pauses or audio cut-offs during interactive video conferencing, especially in extreme weather. We also examine the impact of satellite switching and evolution of satellite routing strategies, offering hints into the future enhancements for multimedia services and for LSNs.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Wireless access networks**; *Network reliability*.

KEYWORDS

Low Earth orbit, Starlink, Multimedia Services, Video Streaming, Video Conferencing, Network Performance Measurement

ACM Reference Format:
Haoyuan Zhao, Hao Fang, Feng Wang, and Jiangchuan Liu. 2023. Realtime Multimedia Services over Starlink: A Reality Check. In *The 33rd edition of the Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV '23), June 7–10, 2023, Vancouver, BC, Canada.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3592473.3592562

1 INTRODUCTION

Recently, Low Earth orbit (LEO) satellites have attracted tremendous attentions from both academia and industries. Different from their high orbit counterpart, particularly, Geosynchronous orbit (GEO) satellites that stay at around 35,780 km, LEO satellites operate at the distance of around 180 to 2,000 km from the surface of the Earth, which can greatly reduce the space-ground communication delay and increase the throughput. As the short orbit distance from the surface of the Earth also reduces the coverage of an LEO satellite, LEO Satellite Network (LSN) constellation, where a large number of LEO satellites work together to achieve the full global coverage, has been envisioned to offer anywhere anytime network connections to ground users. The world's leading LSN service provider, *SpaceX's Starlink*, already has 3754[1] LEO satellites in operation, with an ambitious vision to make the next-generation Starlink constellation eventually harbor up to 30,000.[2] Such a rapid growth also makes LSN a key component to be integrated into the communication infrastructure towards 6G and beyond [10].

There have been efforts on examining and optimizing LSNs in different layers [4, 8, 15–17]. The much improved bandwidth and latency of LSNs potentially enable a wider range of services and applications that were never possible with traditional space networking.[3] Recent works have evaluated the end-to-end performance over Starlink for different applications, e.g., file sharing and Web browsing [7, 11, 13], indicating that Starlink is a viable option for such simple applications. Yet advanced applications remain to be examined. In particular, it is known that multimedia traffic dominates the current Internet (>60%);[4] understanding their performance over Starlink is critical for both service and network enhancements of future LSNs.

In this paper, we present a systematical measurement and analysis on the performance of representative realtime multimedia services over Starlink, including video-on-demand (VoD), live-streaming and video-conferencing services. We confirm that the current Starlink network can support these typical multimedia services with

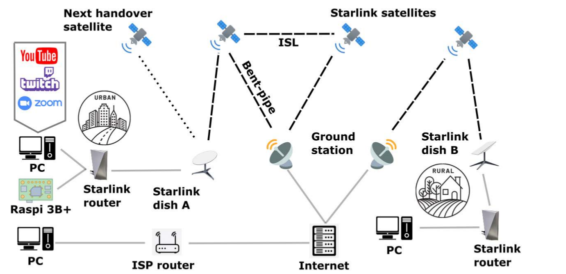

Figure 1: Measurement setup with diverse scenarios and tools.

[*]Both authors contributed equally to this research.

[1]https://planet4589.org/space/con/star/stats.html
[2]https://www.space.com/spacex-starlink-satellites.html
[3]The round-trip propagation time for GEO satellite communication is at least 240 ms, and the end-to-end delay for network transactions is indeed much longer (averaged around 500 ms) [18], making it hardly practical for realtime communications.
[4]https://www.ncta.com/whats-new/report-where-does-the-majority-of-internet-traffic-come

NOSSDAV '23, June 7–10, 2023, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0184-9/23/06…$15.00
https://doi.org/10.1145/3592473.3592562

reasonable QoS under normal conditions. The performance, however, can degrade due to certain factors such as satellite switches, extreme weather, satellite routing strategy, and when both uplink and downlink are utilized at the same time. In particular, we have found that Starlink performs poorly in thunderstorm weather, which can result in severe disruptions for users of certain multimedia services, if not all. We observed significant network jitters caused by satellite switching, which can disrupt users with video pauses during live-streaming. Realtime interactions are an integral aspect of video conferencing services, making them especially susceptible to disruptions on the network. We have devised an experiment to gauge the interactivity of video conferencing services on the Starlink network and found users with Starlink network may have a less fluent video conferencing experience. Our measurement provides useful insights for multimedia service and LEO network providers to improve the performance and stability of their services and networks, respectively. It also contributes to the existing knowledge of LEO networks' ability to support multimedia applications as well as methodologies for their performance measurements.

2 MEASUREMENT OVERVIEW

Our measurement started from December 2022 and lasted 4 months. We have examined a wide range of realtime multimedia services, including VoD, live-streaming, and video conferencing (see Table. 1). Using YouTube, Twitch, and Zoom as representatives, we have analyzed their data from both the application and network level.

Figure 1 shows our measurement setup. We deployed two Starlink dishes in different locations to collect the experiment data: Dish A in an urban area of Pacific West Coast, whereas Dish B in a rural area of the US Mid-South. Both dishes were carefully placed to ensure an unobstructed view of the sky, with an obstructed ratio of *0.734%* for A and *0.029%* for B as reported by the Starlink portal[5]. Since multimedia applications are bandwidth- and computation-intensive, we prepared two desktop PCs with an i7-12700 CPU at 2.1GHz running the applications and connected to each dish's router via Ethernet (denoted as Starlink Ethernet) or WiFi[6] (denoted as Starlink WiFi). We then connected a Raspberry Pi 3B+ through the 300 Mbps Ethernet port of Dish A's router, so as to record the long-term `traceroute` and `ping` results, yet without interfering with the desktop PC running multimedia applications. For comparison, we used a typical terrestrial network service via Ethernet as the baseline (denoted as terrestrial network), which has up to 1 Gbps for download and upload. The stability of this network has been verified multiple times, so as to ensure we obtain a reliable baseline.

To capture application-level data, YouTube and Twitch provide built-in video statistics monitoring tools `Stats for nerds` and `Video Stats`, and Zoom has its own `Zoom Meeting API`[7] to retrieve video and network statistics of a meeting. We used the `Ntopng`[8], a Web-based network traffic monitor to track the network-level data simultaneously during the experiments. `ping` and `traceroute` were also used to measure and analyze the latency pattern and routing strategy for the considered LSN. The Starlink mobile application[9] has been utilized to track the history of Starlink network outages.

[5]http://dishy.starlink.com/
[6]All WiFi adaptors support WiFi 5 protocol
[7]https://marketplace.zoom.us/docs/api-reference/zoom-api/methods/#overview/
[8]https://github.com/ntop/ntopng

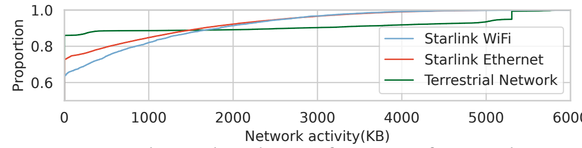

Figure 2: Cumulative distribution function of network activity.

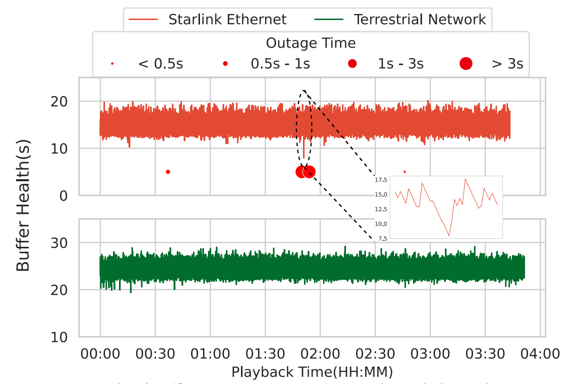

Figure 3: YouTube buffer size over time, with red dots denoting Starlink network outages.

3 MEASUREMENT RESULT AND ANALYSIS

3.1 Video-on-demand Service (YouTube)

In order to assess the end-to-end performance of VoD service [6], we selected a typical 4K resolution video[10] with 24 frames per second (FPS) and average bitrate of 36 Mbps. We collected approximately 32 hours of measurement data in total, with a sample rate of one sample per second. The built-in monitoring tool `Stat for nerds` was used to collect video resolution, number of frame drops, connection speed, network activity and buffer health. Connection speed is measured as the actual download speed, which varies only when the client is currently downloading a video segment. Network activity represents the downloaded volume in each sample and will only be non-zero during downloading of a segment. Buffer health is measured as the length of locally buffered video in seconds, which is also most relevant to the user's experience. To ensure consistency, the browser cache was cleared after each video loop, forcing the client to download each video segment again from the server.

During the measurement for VoD service, we observed that the average connection speed of Starlink Ethernet, Starlink WiFi and terrestrial network are 67 Mbps, 62 Mbps and 381 Mbps, respectively. And the cumulative distribution function (CDF) of overall network activity is shown in Figure 2, where terrestrial network, Starlink Ethernet and WiFi have an 88%, 73% and 65% idle time, respectively (i.e., network activity is 0). This suggests that the Starlink network requires more time to download the subsequent video segment and experiences more frequent network activities compared to the terrestrial network. However, these impacts may be well concealed by a large enough buffer, which is usually the case for VoD service. During the entire measurement, we also noticed that Starlink WiFi performed closely to Starlink Ethernet. Therefore, we will focus on

[9]https://play.google.com/store/apps/details?id=com.starlink.mobile
[10]https://youtu.be/Ac07Qt84WDw

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`
Video-conferencing	Zoom	Zoom's own protocol	`Zoom Meeting API`/ `ping`

Table 2: Average statistics of Twitch service performance.

	Starlink Ethernet	Starlink WiFi	Terrestrial Network
LtB(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

analyzing the performance of Starlink Ethernet in the remainder of this subsection, as it serves as a performance upper bound of the Starlink network.

Figure 3 shows how the buffer size changes over time during one typical experiment, where the red dots indicate the network outages reported from the Starlink mobile application, with the size proportional to the outage time. It is apparent that the overall buffer size is around 15 seconds for Starlink Ethernet while it is around 25 seconds for the terrestrial network, which is because the terrestrial network has more abundant bandwidth to prefetch more segments for each buffering event. For Starlink Ethernet, we observed the average of the outage time is about 6.46 seconds, with the longest outage being 18s. Such outages can cause some impacts on the buffer health as significant buffer drops, where two typical examples can be seen at 00:15 and 01:50 in Figure 3. Nevertheless, the large buffer size usually used in VoD service is generally sufficient to compensate for these Starlink outages and deliver a seamless viewing experience.

In summary, our findings suggest that the end-to-end performance of VoD services over Starlink network is comparable to that over the conventional terrestrial network, which is consistent with previous studies [11, 13]. However, its performance may degrade significantly with high-bandwidth content, e.g., ultra-high-definition 360 videos, where the network activity may happen all the time with the connection bandwidth being largely saturated.

3.2 Live-streaming Service (Twitch)

Since the video content and live period of live-streaming are controlled by the streamer, we selected a channel[11] that is always on, and the streaming content is also relatively fixed from day to day, as our measurement target. This live-streaming channel has a resolution of 1,920 x 1,080 with 35 FPS, and the average bitrate is around 6 Mbps. We collected around 77 hours of data at a rate of one data sample per second. The built-in tool `Video Stats` provides video resolution, FPS, number of skipped frames, buffer size, latency to broadcaster (LtB) and playback bitrate. The LtB refers to the interval from the time that the video content is captured by the broadcaster to the time that it is displayed on the viewer's monitor [19]. In Twitch, the LtB can be dynamically adjusted to better match viewers' network conditions using the Low-on-Latency (LoL) method [1, 9], where the server will assign larger LtB to those viewers with poor network conditions, so that they will still have a smooth viewing experience but can suffer from higher latency when communicating with the streamer via text chat.

[11] https://www.twitch.tv/pcentinela

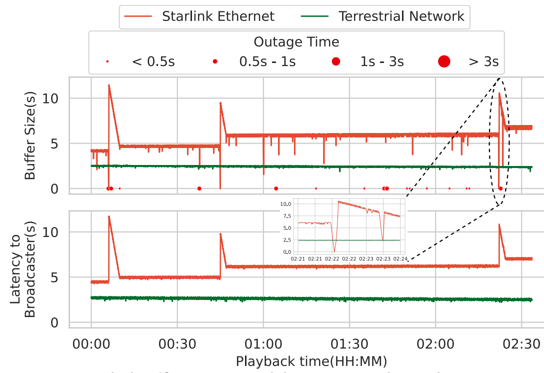

Figure 4: Twitch buffer size and latency to broadcaster over time, with red dots denoting Starlink network outages.

The overall performance of Twitch is shown in Table. 2. The terrestrial network has a significant advantage compared to Starlink Ethernet/WiFi except for the playback bitrate, which indicates the throughput is not the bottleneck for this streaming source. Figure 4 shows how the buffer size and LtB in Twitch change during a typical measurement round of about 2.5 hours. It is clear to see that the service via Starlink Ethernet encounters a higher number of buffer drops compared to the service via the terrestrial network, with some drops being so significant that they result in the streaming being paused. Fortunately, the LoL mechanism has worked to mitigate the impact on the user's experience after the LtB increases. For example, there are three severe buffer drops that consume up all preloaded video contents at 00:06, 00:45 and 02:22, respectively. After each outage, the LtB first initializes at around 10 seconds, then gradually declines until stabilizing at a higher latency level than before. It is worth noting that the LtB decrement is achieved by speeding up the video playback at a rate of 1.025, which is reported as generally unnoticeable to human [1]. Higher LtB also allows clients to prefetch more buffered content and is affordable for potential future outages. For instance, at 02:23, another long outage occurs just after the buffer drops, but the streaming does not pause because the larger buffer size helps reduce the chance of causing empty buffers. In live-streaming service, the length of the available buffered video is considerably shorter compared to VoD service. Therefore, the outages in the Starlink network are more likely to empty the current buffer, causing the streaming to freeze and directly impacting the user's experience.

Another observation from Figure 4 is that some buffer drops do not correspond to the outages reported by the Starlink mobile application. A further investigation reveals that these drops are caused by significant network jitters resulting from satellite switching, which will be further discussed in Section 4. Basically, satellite handover failure or a large delay difference between the two satellites involved in handover can cause significant network jitter within a short period, which is detrimental to live streaming services.

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

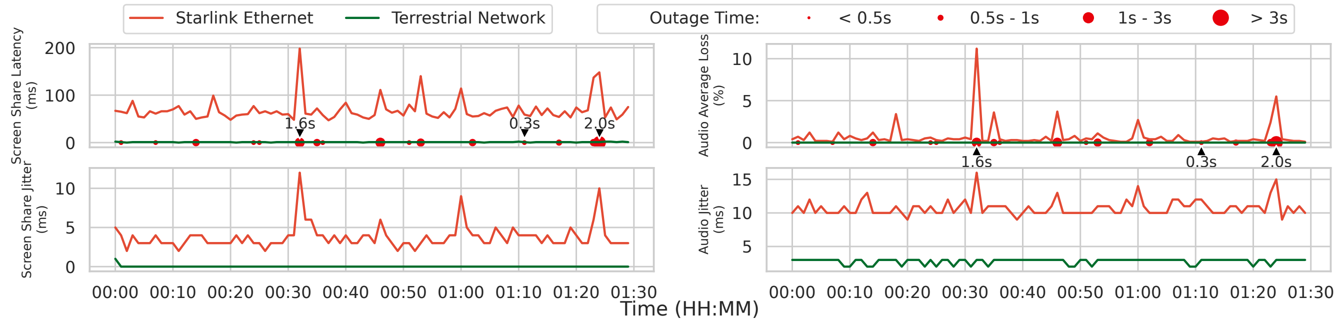

Figure 5: Screen share and audio statistics in a particular zoom meeting, with red dots denoting Starlink network outages.

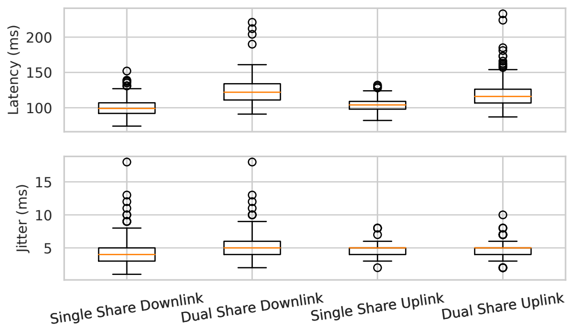

Figure 6: Latency and jitter statistics of Starlink users in a 2-participant meeting. Single share means only one user shares screen and dual share means both users share screens.

3.3 Video-conferencing Service (Zoom)

The COVID-19 pandemic has spurred numerous studies on the network performance of video conferencing applications [2, 3, 12, 14]. It was also reported that Zoom, one of the most popular video conferencing applications, has experienced a fourfold increase in traffic since 2020 [5]. As a result, our team has undertaken a performance measurement of Zoom over the Starlink network.

To measure the network performance of Zoom, we used Zoom's `Zoom Meeting API`, which allowed us to retrieve the per-minute statistical data on network metrics, such as average latency, jitter, and average loss for both audio and video. To ensure consistency across our measurements, we used Zoom's screen share function rather than a video camera. This allowed 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 bitrate of 9001 kbps, and an audio bitrate of 126 kbps. As video conferencing applications have high demands for internet connection stability, we ensured that all participants were connected to the network via Ethernet cables during Zoom meetings.

Table 3: Statistics of user interactions.

	terrestrial network	Starlink Ethernet
# of interactions	1136	869
Average of RTT (ms)	0.57	0.69
Variance of RTT (σ^2)	0.0073	0.0095

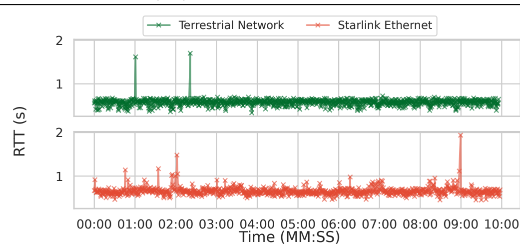

Figure 7: RTT during a 10-minute interactive Zoom session.

During our measurement, we collected over 60 hours of Zoom meeting data. On average, when the terrestrial network hosted screen sharing, the FPS and bitrate were 24 FPS and 3649 kbps, respectively, while the corresponding values for Starlink Ethernet were 13 FPS and 1512 kbps. Screen sharing hosted on the Starlink Ethernet resulted in lower FPS and bitrate of 12 FPS and 1522 kbps, respectively, for both Starlink Ethernet and terrestrial networks. These results indicate that Starlink Ethernet is the bottleneck during screen sharing. The bitrate and FPS did not show a discernible pattern of fluctuations over time, so our analysis will center on screen share latency/jitter and the average audio loss rate. Furthermore, our measurements on the cases of two and three participants showed no significant differences in network performance. This is because, according to Zoom's documentation[12], 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 with the number of users in the meeting, assuming they are not on the same LAN.

3.3.1 Screen Share from One User. We first conducted an experiment where only one of the participants was screen sharing. Figure 5 shows the measurement results of a typical Zoom meeting with two participants, with one participant video sharing using Starlink Ethernet (dish A) and the other participant watching via the terrestrial network. It is evident from the figure that the terrestrial network is stable and low for both audio and video latency/jitter. In contrast, 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 Figure 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 conducted an experiment in which a participant using the terrestrial network shared the screen and a participant using Starlink Ethernet watched the screen share. We observed that the overall pattern of performance remained consistent with our previous findings, indicating that in both scenarios, participants are sensitive to Starlink outages. Furthermore, we found

[12]https://explore.zoom.us/docs/doc/Zoom%20Connection%20Process%20Whitepaper.pdf

that the average latency for Starlink Ethernet's uplink was 61 ms, while the average latency for its downlink was 45 ms.

3.3.2 Screen Share from Multiple Users. Zoom offers a feature for users with multiple monitors that allows them to share their screen while also watching another'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 video conferencing experience. During the experiment, we had two Starlink users share their screens while simultaneously viewing each other's screen share content. Our results showed 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 Figure 6, the network latency increased for users in a dual screen share session. We conjecture that this is caused by the limited communication capacity of the Starlink dish/router, and it can only send or receive a restricted amount of data at the same time.

3.3.3 User Interactions. Since video-conferencing service often involves many realtime user interactions, we also conducted an experiment to investigate the impacts of the Starlink network on such interactions. Yet one challenge here is how to automate the user interactions and make them repeatable and as objective as possible. To this end, we carefully designed a method with each participant using two monitors, so that they can share their screens while viewing each other's screen share content. Then a Python script was run at 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 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 changes its display color, and then participant A will do the same upon detecting color change from participant B. Every time a participant detects a color change, the script will record a timestamp. Before each experiment, we used Windows' official sync time feature to ensure both systems' clocks were synced to the same server.

Figure 7 compares the results of the user interaction experiment on both the 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 terrestrial network has more stable interactions compared to the Starlink network. It is worth noting that there are two outliers in the figure for terrestrial network. We conducted multiple measurements, and each measurement contained one or two outliers. However, the overall statistics shown in Table 3 indicate that terrestrial network is more stable compared to Starlink Ethernet. The terrestrial network made about 23.5% more interactions than the Starlink network, with lower latency and variance, indicating that Zoom users with the Starlink network may have less fluent interaction experiences.

In summary, compared with the terrestrial network, although video conferencing service over the Starlink network can 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, video conferencing service is generally more sensitive to network variations. Therefore, users using Starlink for video conferencing services may experience more disruptions such as frame losses and audio cut-offs.

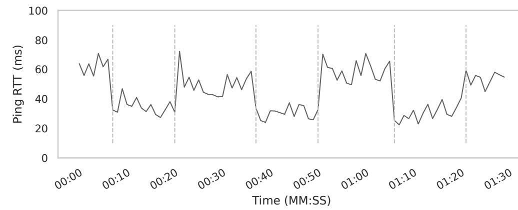

Figure 8: A typical pattern observed in `ping` RTTs, where the interval for each steady state is indicated by vertical lines.

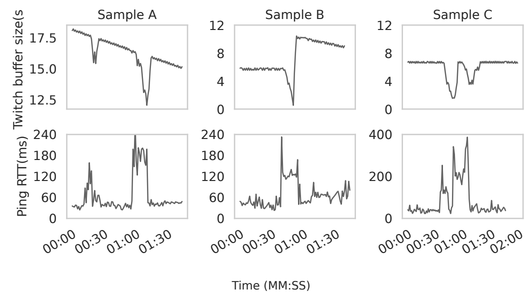

Figure 9: Impact of large jitters on Twitch's buffer size.

4 INSIGHT INTO INFLUENTIAL FACTORS

4.1 Satellite Switching

One unique characteristic of LSN compared to other network systems 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 time span, as illustrated in Figure 1. However, as *SpaceX* almost hides everything between dish and ground station to make it a blackbox (e.g., `traceroute` can barely get any meaningful route information there), it makes the investigation of the impact of satellite switching on multimedia services even more challenging. To this end, we utilized the `ping` command to collect the RTT time to ground station (GS) with a one-second interval and eventually collected 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. As an illustration, Figure 8 shows a typical example of this pattern, where starting from time *00:09*, the RTT remains steady with only small fluctuations for approximately 10 seconds and then switches to another steady state, and so on so forth. We also noticed 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

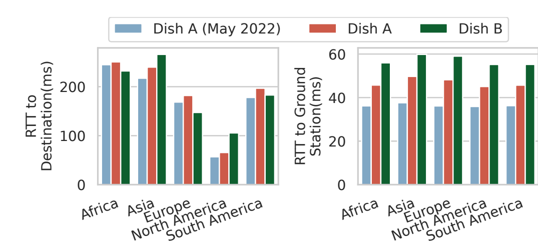

Figure 10: `traceroute` RTTs between different destinations.

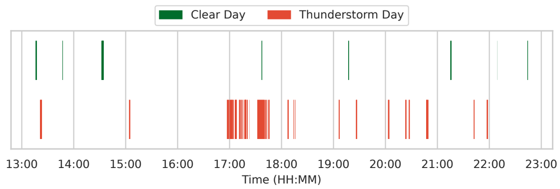

Figure 11: A 12-hour outage history, where outages less than 2 seconds are multiplied by a scaler of 80 for better visualization.

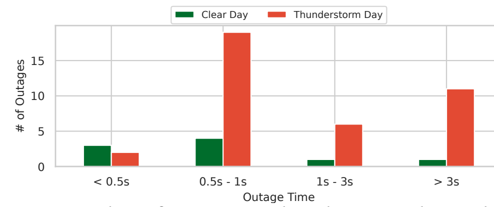

Figure 12: Number of outages in thunderstorm day and clear day.

frequency is consistent with what has been observed in the Starlink satellite simulator[13], where the primary link will shift to next nearest satellite for roughly every 15 seconds. In addition, the RTT difference between each steady state in Figure 8 is approximately 30ms, which can be even larger if handover fails or when switching to a satellite that is far from user's dish. Such large jitters can cause severe impacts on multimedia services with significant performance degradation. Figure 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 Figure 4.

4.2 Starlink Routing Strategy

Early last year, we also did some preliminary measurements on the Starlink network, where we found that the Starlink network will connect to the GS geographically closest to the dish [11]. However, we noticed that our measurement results in this work showed a different routing behavior, where all the packets were directed to a GS that may be geographically far away from our dish. This could be an indicator that *SpaceX* has changed their routing strategy or started to use Inter-Satellite Links (ISLs) since late 2022, especially considering that some Starlink users have reported that they received the ISL service enabling announcement by the end of 2022[14]. To evaluate the performance of this new routing strategy, we compared our current `traceroute` data with the previous `traceroute` data collected in May 2022 on dish A, and plotted the RTTs to different continents in Figure 10 with 25th percentile 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 use a different routing strategy from early last year. And 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 on multiple satellites.

4.3 Weather Impact on Starlink

During our measurement, we encountered thunderstorm weather in the area of dish B, which had a significant impact on Starlink's network performance, with more frequent and long-duration network outages. As an illustration, Figure 11 compares the 12-hour outage history of a thunderstorm day and a clear day. It clearly shows a significant network outage between 5 pm and 6 pm during the thunderstorm day, which matches the observation of extremely high thunderstorm activity reported during the period of that day. Figure 12 further compares the number of outages of different lengths between a thunderstorm day and a clear day. Such severe outages can have a significant impact on various multimedia services, e.g., Zoom users may find that their meeting connections are disrupted, and YouTube/Twitch videos may be paused or even unavailable under such weather conditions. We also noticed multiple occurrences of snow weather in the vicinity of dish A. However, the correlation between performance variation and snow weather is not very strong, which might be because the Starlink dish comes with a function to detect and melt snow automatically. Also, during our measurement, we did not observe any significant difference in Starlink's performance during rainy weather conditions.

5 CONCLUSION

This paper presented 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 four-month measurement covered VoD, live-streaming and video conferencing services, with data collected via two Starlink dishes located in totally different geographical areas. Our findings showed 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 video-conferencing mostly impacted due to its real-time user interactions. We believe these findings can provide useful information for future enhancements on both multimedia services and LEO satellite networks.

We will continue monitoring Starlink's network performance and explore ways to overcome the current limitations of LSNs. We think there is a discernible pattern in Starlink outages, and by utilizing weather and Starlink outage data, we could train a model that can predict Starlink outages to further improve network performance. For example, with predicted network outages, application designers could preload more video data in the buffer for video streaming applications to avoid interruptions in the video streams.

6 ACKNOWLEDGEMENT

We appreciate the constructive comments from the reviewers. This research is supported by an NSF I/UCRC Grant (1822104), a Canada NSERC Discovery Grant, an NSERC CGS M, and a British Columbia Salmon Recovery and Innovation Fund (No. 2019-045).

[13]https://starlink.sx/
[14]https://wccftech.com/starlink-turns-on-laser-satellites-for-region-with-fourth-month-long-night/

REFERENCES

[1] 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 24* (2021), 2300–2314.

[2] 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* (New York, NY, USA, Apr. 2021), ICPE '21, Association for Computing Machinery, pp. 65–71.

[3] 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* (New York, NY, USA, Nov. 2021), IMC '21, Association for Computing Machinery, pp. 216–228.

[4] Chen, L., Tang, F., and Li, X. Mobility-and load-adaptive controller placement and assignment in leo satellite networks. In *IEEE INFOCOM 2021* (2021), pp. 1–10.

[5] Choi, A., Karamollahi, M., Williamson, C., and Arlitt, M. Zoom session quality: A network-level view. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 13210 LNCS* (2022), 555–572.

[6] 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. *ACM SIGCOMM computer communication review 41*, 4 (2011), 362–373.

[7] 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* (2022), pp. 151–158.

[8] 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 58*, 2 (Apr. 2022), 1487–1491.

[9] 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 11th ACM Multimedia Systems Conference* (2020), pp. 321–326.

[10] 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 59*, 12 (2021), 36–42.

[11] 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 *IEEE INFOCOM* (2023), pp. 1–10.

[12] 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* (New York, NY, USA, Nov. 2021), IMC '21, Association for Computing Machinery, pp. 229–244.

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

[14] 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* (New York, NY, USA, Oct. 2022), IMC '22, Association for Computing Machinery, pp. 244–260.

[15] 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 58*, 3 (2022), 2606–2610.

[16] 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 647* (2021), A54.

[17] Uran, C., Horvath, K., and Wöllik, H. Analysis of a Starlink-based Internet connection – ROADMAP-5G, July 2021.

[18] Vankka, J. Performance of satellite gateway over geostationary satellite links. In *MILCOM 2013 - 2013 IEEE Military Communications Conference* (2013), pp. 289–292.

[19] Zhang, C., and Liu, J. On crowdsourced interactive live streaming: a twitch. tv-based measurement study. In *Proceedings of the 25th ACM workshop on network and operating systems support for digital audio and video* (2015), pp. 55–60.