

Starlink Performance through the Edge Router Lens

Sarah-Michelle Hammer

TU Berlin
Berlin, Germany

Max Franke

TU Berlin
Berlin, Germany

Vamsi Addanki

TU Berlin
Berlin, Germany

Stefan Schmid

TU Berlin
Berlin, Germany

ABSTRACT

Low-Earth Orbit satellite-based Internet has become commercially available to end users, with Starlink being the most prominent provider. Starlink has been shown to exhibit a periodic pattern with a characteristic throughput drop on the boundaries of 15s intervals. A multitude of prior works hypothesize various root causes for this pattern, such as re-ordering and packet loss. Some works have attributed these effects to the edge router, advocating for explicit feedback to the transport layer. However, with the edge router being a proprietary Starlink device, it raises questions about the extent of its influence on periodic throughput drops, losses, and jitter, leaving us to wonder if we fully understand the underlying issues.

This paper presents the first measurement study with a vantage point that is by far the closest (last hop) to the core Starlink network. We use a Generation 1 dish, which allows us to bypass the proprietary Starlink router and connect a Linux server directly to the dish. We investigate the impact of the edge router on the observed periodic pattern in Starlink performance. Our results are primarily negative in terms of any significant buffer buildup and packet losses at the edge router, suggesting that the causality of the observed patterns lies entirely in the core network, a proprietary space that cannot be fixed by the end user. Interestingly, we observe similar patterns even with a constant bitrate UDP sender, likely indicating that the periodic drop in throughput is not an inherent limitation of existing TCP implementations but rather a core network characteristic!

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

LEO-NET 24, November 18–22, 2024, Washington D.C., DC, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1280-7/24/11.

<https://doi.org/10.1145/3697253.3697273>

CCS CONCEPTS

• **Networks** → **Network measurement**; *Network performance analysis*.

KEYWORDS

Starlink, Low-Earth Orbit Satellite, Internet measurements

ACM Reference Format:

Sarah-Michelle Hammer, Vamsi Addanki, Max Franke, and Stefan Schmid. 2024. Starlink Performance through the Edge Router Lens. In *2nd International Workshop on LEO Networking and Communication (LEO-NET 24), November 18–22, 2024, Washington D.C., DC, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3697253.3697273>

1 INTRODUCTION

The first wide-area packet-switched network dates back to the late 1960s [9], and it took more than five decades for the internet to grow from a few computers at universities on the West Coast to a billion devices connected across the globe. Several breakthroughs in technology have enabled this growth, including the development of the TCP/IP protocol suite, broadband connectivity to end-user homes, and wireless technologies such as WiFi and cellular networks [9]. Yet, even after five decades of technological advancements, providing high-speed internet access to all corners of the world remains an economic and technical challenge[3].

Low-Earth Orbit (LEO) satellite networks have emerged as a promising solution to this challenge, offering internet access to remote areas, even during disasters [19]. Service providers such as Starlink [18], OneWeb [14], and Amazon’s Project Kuiper [7] have deployed thousands of satellites [11], transforming LEO satellite networks into the next generation of ISPs. Starlink is currently the most prominent provider, with over three million users worldwide [17].

A multitude of recent measurement studies have examined the performance and network characteristics of the Starlink network in comparison to cellular and wired networks. End-to-end performance measurements of Starlink show its potential to challenge and replace traditional networks in terms of reachability, low latency, and high data rates [5, 12, 13], albeit with the caveat of short-duration

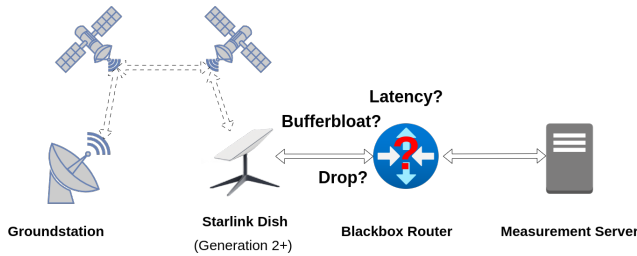


Figure 1: Prior works rely on vantage points behind the Starlink router, a black box that leaves several questions. In contrast, we present measurement results taken directly behind the dish providing further insights into the periodic throughput drop.

throughput drops occurring approximately every 15s periodically [2, 13]. Interestingly, this periodic pattern in both uplink and downlink is a synchronized event across users, orchestrated by a global scheduler at 12, 27, 42, and 57 seconds past every minute [13, 20]. From the user’s networking stack perspective, the periodic pattern manifests as a drop in TCP’s congestion window (and consequently throughput) in response to potential spikes in round-trip time, packet reordering, or high loss rates at each 15s mark. Mitigating this undesirable reaction of TCP is challenging, as it requires a deep understanding of the root causes, raising the question of whether and to what extent the user-facing network is responsible for the problem.

Unfortunately, existing works leave one important aspect of the user-facing network unexplored: the edge router. The edge router is a proprietary device (essentially a black box) provided by Starlink to end users. It serves as the last hop in the user’s network before packets enter the core Starlink network as shown in Figure 1. Given the periodic pattern in throughput and the disruptions caused by the global scheduler, it is natural to question whether any buffer buildup occurs at the edge router, leading to packet drops and subsequently influencing TCP’s reaction. Many questions remain regarding how the router behaves during the 15s reconfiguration intervals and how it impacts the observed patterns in throughput, delay, and loss.

Fortunately, we have access to a Generation 1 Starlink dish, which allows us to physically bypass the proprietary Starlink router and connect a Linux server directly to the dish. Thereby our measurement machine occupies the same point of view as the Starlink router would, as a network node directly connected to the dish, and as such we consider our measured results to be from a router’s perspective or through the lens of the edge router.

We conduct a series of measurements to understand the role of the edge router in the observed patterns in TCP’s reactions, focusing on the following key questions:

(Q1) Is there any buffer buildup at the edge router during upload?

Our measurements show that there is in fact *no* buffer buildup at the edge router during upload, suggesting that packet losses, if any, occur either at the dish or in the core network beyond the edge router.

(Q2) Does the edge router experience interruptions on its port transmissions?

We do not observe any interruptions to the edge router’s transmissions; all our `ethtool` measurements show a consistent rate of data transmission. Further strengthening recent observations, our measurements indicate that the periodic drops in throughput stem directly from the core network.

(Q3) Is the periodic pattern in throughput inevitable?

We show that the periodic pattern in throughput is not an inherent limitation of the current TCP implementation in relation to the Starlink network. In fact, we observe the same pattern for UDP! This indicates that it is likely not explicitly solvable, neither from the transport layer nor with sophisticated edge routers.

As we consider the implications of these periodic throughput drops, intriguing questions arise. How can applications adapt to these brief yet impactful interruptions without compromising user experience, particularly in real-time scenarios such as video conferencing and online gaming? We aim to explore these questions in our future work.

2 BACKDROP

We first provide a brief background on the state-of-the-art understanding of the performance of the Starlink network and in particular its periodic reconfigurations. We then highlight the challenges and motivate the need for a deeper understanding of the edge router’s role in the observed patterns in throughput and packet drops.

2.1 Starlink Performance: Known & Unseen

Recent studies have highlighted specific network performance characteristics of Starlink. Notably, the existence of multiple latency “bands” [20] and the periodic performance drops in TCP throughput, occurring at 15-second intervals, significantly affecting performance. These drops have been observed to cause severe path under-utilization, with only 46% utilization of the estimated available bandwidth [2]. Unfortunately, none of the widely available congestion control algorithms provide substantial improvement; even the best-performing BBR reaches only about 50% estimated path utilization [5]. Mohan et al. [13] note that these significant throughput drops occur during both uploads and downloads, independent of any changes in the User-Dish-to-Satellite mapping. The repeating 15-second patterns have been attributed to global reconfiguration intervals [13], orchestrated

by a global scheduler at 12, 27, 42, and 57 seconds past every minute [20]. Other metrics have also exhibited periodic patterns due to these reconfiguration intervals, including round-trip time (RTT) shifts in 15-second windows [13], as well as significant peaks in RTT [20], packet loss [10], and one-way delay variation [10]. Simulation work has long proposed that packet reordering, due to path changes in the highly dynamic LEO networks, is a significant issue [1, 6]. This may be strongly related to the aforementioned momentary peaks in delay. If a packet is significantly delayed while subsequent packets experience much lower delays due to path changes, the latter may arrive before the first transmitted one. In fact, several recent works suggest that many of the losses perceived by TCP may be attributed to reordering events causing erroneously assumed congestion states, thus leading to throughput deterioration [1, 16]. The existence of the 15-second periodic pattern in TCP throughput, alongside patterns in metrics of latency and loss, is well established. However, the interaction between latency, loss, and reordering in the causality of TCP throughput degradation requires further investigation.

2.2 Challenges of Edge Vantage Points

Measurement works on the Starlink network have so far relied on vantage points that are two hops away from the satellite network, with the dish and Starlink router along the path. Unfortunately, recent generations of Starlink dishes can only be connected to the proprietary router provided by Starlink, effectively making it a black box from a measurement perspective. The proprietary dish connector inevitably requires the Starlink router as an extra black box on the path. While a bypass mode and Ethernet adapter for the proprietary router are available, careful reverse engineering of this adapter has shown it to essentially be a passthrough to 2-port Ethernet switch circuitry on the router, with one port exposed through the Ethernet adapter [8]. Therefore with Starlink equipment of Generation 2 and above, even in bypass mode, a direct connection of a third party device to the dish without extra switching circuitry is not possible.

As a result, several questions remain unanswered and hard to reason about, e.g., whether there is any periodic buffer buildup, interruptions in transmission, or packet losses *at the edge*, that may potentially influence TCP’s congestion response. Several recent works focus on improving the end-host TCP stack tailored to the Starlink network [10, 16], but the root causes of the periodic throughput drops and the role of the edge router still remain unclear.

2.3 Our Setup: Bypass the Starlink Router

Starlink’s Generation 1 dish was originally shipped with an RJ45 power injector, allowing for direct connection to any end point, e.g., a custom router. Thankfully, Generation 1

Potential causes	Our measurement result
Packet drops (edge)	No
Dish↔Router disruptions	No
Reordering	No
Timeouts	No
Packet drops (core)	Likely yes

Table 1: Prior works hypothesize various potential root causes for the periodic throughput drop pattern. Our measurements rule out the most common on and indicate a likely yes for packet drops in the core network.

dishes are still operational (although extremely rare to find) and similar in performance[15]. We were able to obtain one to conduct measurement studies through the edge router lens. For this, we connect the dish directly to a Linux server equipped with an Intel X550 NIC. This setup allows us to investigate the role of the edge router in the observed patterns in throughput, latency, and loss. Our direct-connect measurement setup enables us to measure sending rates to and from the dish. Furthermore, it allows us to observe the queuing behavior of the edge router and the impact of the global reconfiguration intervals on the edge router’s transmission rates. Unless explicitly specified, all measurements in this paper involve communication between the Starlink edge node (our Linux server) and a node (an Azure VM) in the terrestrial network positioned near the closest Starlink PoP in Frankfurt, Germany.

3 RESULTS

We now report our measurement results. In line with the observations made by various measurement studies in the recent past, Figure 2 shows the characteristic periodic throughput drop in download with TCP in our directly connected setup. In this section, we perform a series of measurements starting at link level and moving up to the transport layer socket statistics and eBPF traces to better understand the performance issues and the root causes. Table 1 summarizes our main results.

3.1 Edge Router → Dish Link Utilization

In order to understand whether there are any disruptions to the link utilization between the edge router and the dish (due to periodic reconfigurations), we launch a UDP flow from the edge router towards the terrestrial node. We use `iperf2` for transmission and `ethtool` to measure link-level utilization. In this upload scenario, we do not observe *any* disruptions to the link utilization, as shown in Figure 3. Furthermore, we measure the queuing behavior at the edge router using Linux `tc`. We launch a TCP flow from the edge router towards the terrestrial node while setting different bandwidth limits

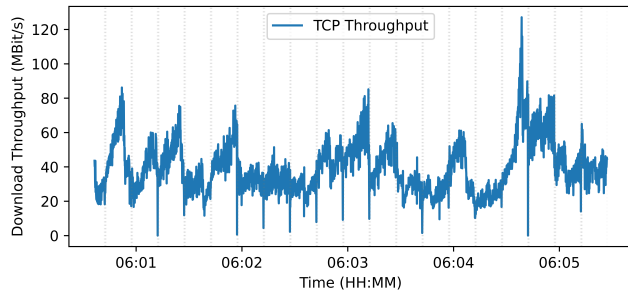


Figure 2: Download throughput of TCP connection showing the characteristic periodic throughput drops at 15s intervals (vertical grey dotted lines).

ranging from 0.1 Mbps to 100 Mbps and log the queue buildup towards the dish. From Figure 4, we observe that queue buildup only occurs when the edge router is severely rate limited (< 50 Mbps). A NIC connected to the Starlink dish typically autonegotiates to 1 Gbps unless the connecting device has a lower bandwidth. Our results confirm that queue buildup at the edge router plays no significant role in relation to the observed periodic patterns in performance.

3.2 Ping latency in Relation to minRTO

Given that the edge router does not cause any packet drops, we now turn to latency measurements to understand the triggering factors for TCP’s reaction, as shown in Figure 2. Latency peaks during reconfiguration can present a serious challenge to TCP if they increase the round-trip time (RTT) above the retransmission timeout (RTO), causing TCP to consider the segment lost. Notably, segments deemed lost due to RTO typically elicit a more severe reaction from TCP than packets considered lost because of duplicate acknowledgments (Fast Recovery). In many TCP stacks, the default minimum RTO is set to 200ms, which we confirmed on our measurement machine. To collect fine-grained latency statistics, we conducted multiple 1-hour measurement runs of UDP “pings” with the tool `irtt` [4] over several days and at different times of the day. Packets were sent at an interval of 20ms to ensure stable timer behavior. To confirm that no packets with significant delays were overlooked by the sampling interval, we also conducted runs with 1ms intervals. Figure 5 shows a representative measurement of RTT relative to the 200ms minimum RTO value. Out of the 897,899 RTT values collected, only 22 exceeded the minimum RTO of 200ms, indicating that Starlink’s latency rarely triggers timeouts and that latency by itself does not represent the root cause of TCP’s periodic throughput drops.

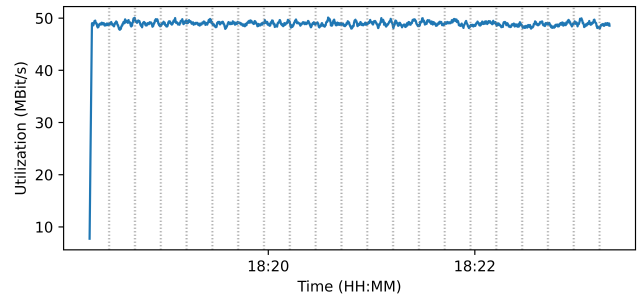


Figure 3: Transmission rate reported by ethtool for the link between the edge router and the dish during a UDP upload towards the terrestrial node showing no disruptions during the global reconfiguration intervals.

3.3 TCP Retransmissions

We now move further up in the network stack and measure TCP socket statistics as well as trace each function call using eBPF tracing. Figure 6 shows the socket statistics of the sender (terrestrial node) during a TCP download at the Starlink edge router. In particular, “retrans” value reported by `ss` (Linux socket statistics tool) shows the number of retransmissions triggered at the sender, while “reord_seen” shows the low levels of reordering observed. Given that our previous measurements in Figure 5 indicate that timeouts are unlikely, we dig deeper by tracing every function call in the TCP stack using eBPF tracing. We omit these traces due to space constraints. For the duration of the iperf transmission in Figure 6, our eBPF traces show only one retransmission timeout. This indicates that the retransmissions are caused primarily due to packet losses. To confirm whether packet losses consistently occur in the Starlink network, we performed several measurements using ICMP ping and logged the sequence numbers of ICMP replies arriving at the measurement host (edge router). Figure 7 shows our results. All positive values greater than 1 indicate gaps in the sequence numbers, suggesting packet losses, compared to negative values indicating reordering. This further confirms that the periodic throughput drops in TCP are primarily related to packet losses and not retransmissions triggered by RTO timeouts or reordering.

3.4 Beyond Starlink Core Network

Our measurement results so far indicate packet losses to be the root cause of TCP’s periodic throughput drop but at the same time we observe no packet losses and buildup at the edge router. This begs the question whether certain packet losses are inevitable. To answer this, we observe the arrival rate at the receiver (terrestrial node) corresponding

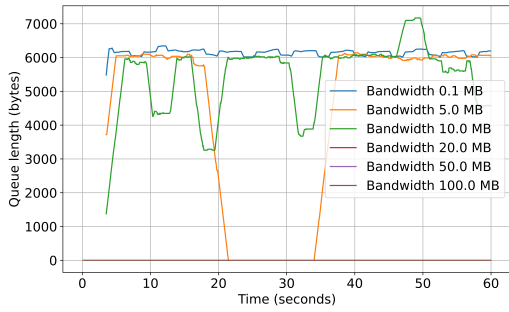


Figure 4: Queuing at the edge router does not show any significant buffer buildup unless the router→ dish link has significantly low bandwidth (< 50Mbps).

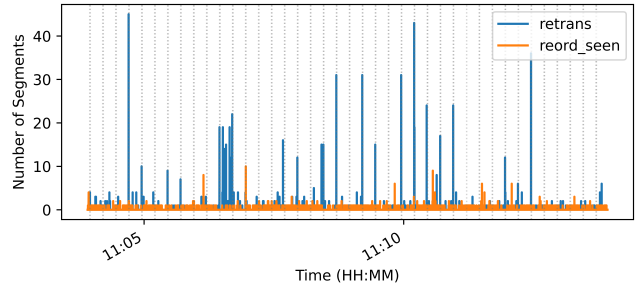


Figure 6: Socket Statistics of the sender (terrestrial node) during a TCP download at the Starlink edge router. The retrans value shows the number of retransmissions triggered at the sender.

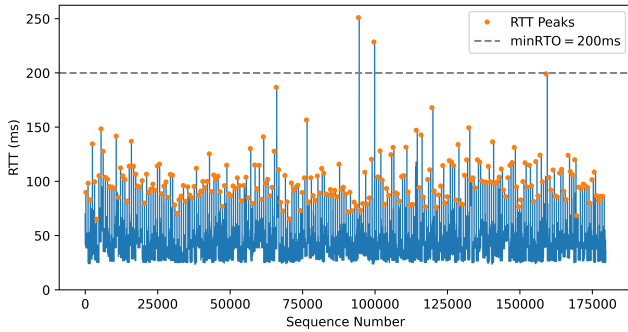


Figure 5: Round-trip-time (RTT) in relation to TCP’s minimum retransmission timeout (minRTO= 200ms)

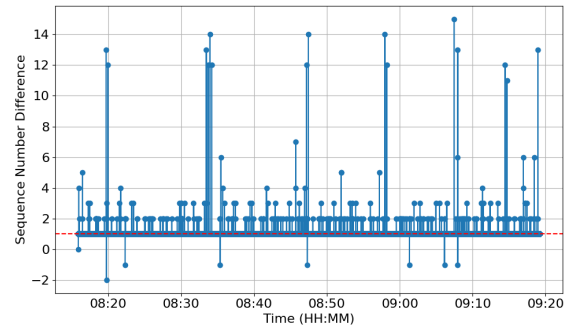


Figure 7: Difference in the Sequence Numbers of ICMP replies arriving at the edge router. Positive values (> 1) indicate gaps in the sequence numbers, suggesting packet losses.

to our iperf UDP test in Figure 3 where the edge router (Starlink edge node) acts as the sender achieving nearly constant bitrate transmission towards the dish. Figure 8 shows the corresponding bitrate received at the receiver (terrestrial node) reported by ethtool. Surprisingly, the same periodic throughput pattern observed for TCP in Figure 2 reappears for UDP as well! The throughput pattern even for UDP aligns with the 15s intervals. Our results confirm that the periodic throughput pattern is not an inherent limitation of existing TCP implementations. While undesirable overreaction of TCP to packet losses can be mitigated to some extent, the periodic pattern appears to be a core network characteristic and not an artifact of the transport layer.

4 DISCUSSION

On the one hand, our results are positive in the sense that the Starlink router (and any edge router) is not a potential bottleneck and does not play a significant role in the short-duration performance interruptions observed in the Starlink network. On the other hand, our results are negative in that Starlink appears to have core network characteristics that cause periodic throughput drops, which are likely inevitable

unless transmissions are explicitly paused i.e., by trading a little latency for packet drops. In particular, certain loss in throughput appears to be inevitable. This raises the question of how applications react to, and can adapt to, these short-duration interruptions without compromising user experience, particularly in real-time scenarios such as video conferencing and online gaming. Given that the performance dips are globally synchronized events at explicit times, the networking stack can leverage this predictability to mask the effects of these interruptions. For example, transport layer could pause transmitting new data and buffer more data during the reconfiguration intervals. We are currently exploring a Linux traffic-control tc qdisc that periodically pauses and resumes transmissions at the reconfiguration intervals. This approach could potentially reduce the impact of the periodic throughput drops on TCP performance without any invasive modifications to the hardware or Kernel. Adaptively synchronizing the periodic pauses to the global Starlink’s events and avoiding drifts in the synchronization remains a challenge.

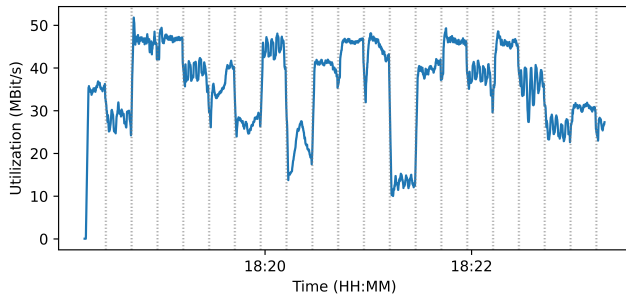


Figure 8: Receive rate at the receiver (terrestrial node) during constant rate UDP upload from the edge router connected directly to Starlink dish.

5 CONCLUSION

We presented a measurement study on Starlink performance from a unique vantage point directly behind the dish. Our findings uncover several insights into the root causes of the periodic throughput drop pattern consistently observed in recent works. Notably, we demonstrate that the edge router does not play a significant role in these performance dips. Our results suggest that the periodic throughput drop pattern is likely an inherent characteristic of the Starlink core network, rather than an artifact of the transport layer or the edge router. These findings open up new research directions for exploring how applications can adapt to these short-duration interruptions without compromising user experience.

ACKNOWLEDGMENTS

The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the programme of “Souverän. Digital. Vernetzt.” Joint project 6G-RIC, project identification number (PIN): FKZ 16KISK030.

REFERENCES

- [1] Rasmus Sibbern Frederiksen, Thomas Gundgaard Mulvad, Israel Leyva-Mayorga, Tatiana Kozlova Madsen, and Federico Chiariotti. End-to-end delivery in leo mega-constellations and the reordering problem. arXiv:2405.07627, 2024.
- [2] Johan Garcia, Simon Sundberg, Giuseppe Caso, and Anna Brunstrom. Multi-timescale evaluation of starlink throughput. In *Proceedings of the 1st ACM Workshop on LEO Networking and Communication, LEO-NET '23*, page 31–36, New York, NY, USA, 2023. Association for Computing Machinery.
- [3] Mark Graham, Scott Hale, and Monica Stephens. Featured graphic: Digital divide: The geography of internet access. *Environment and Planning A: Economy and Space*, 44(5):1009–1010, 2012.
- [4] Paul Heist. irtt: A network latency and loss measurement tool. <https://github.com/heistp/irtt>, 2024. Accessed: 2024-08-02.
- [5] Mohamed M. Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. A browser-side view of starlink connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference, IMC '22*, page 151–158, New York, NY, USA, 2022. Association for Computing Machinery.
- [6] Simon Kassing, Debopam Bhattacharjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. Exploring the “internet from space” with hypatia. In *Proceedings of the ACM Internet Measurement Conference, IMC '20*, page 214–229, New York, NY, USA, 2020. Association for Computing Machinery.
- [7] Project kuiper. <https://www.aboutamazon.com/what-we-do/devices-services/project-kuiper>, 2024. Accessed: 2024-08-01.
- [8] Oleg Kutkov. Reverse engineering of the starlink ethernet adapter; oleg kutkov blog. <https://olegkutkov.me/2022/03/07/reverse-engineering-of-the-starlink-ethernet-adapter/>, 2022. Accessed: 2024-08-02.
- [9] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. A brief history of the internet. *SIGCOMM Comput. Commun. Rev.*, 39(5):22–31, October 2009.
- [10] Jihao Li, Hewu Li, Zeqi Lai, Qian Wu, Yijie Liu, Qi Zhang, Yuanjie Li, and Jun Liu. Satguard: Concealing endless and bursty packet losses in leo satellite networks for delay-sensitive web applications. In *Proceedings of the ACM on Web Conference 2024, WWW '24*, page 3053–3063, New York, NY, USA, 2024. Association for Computing Machinery.
- [11] Jonathan McDowell. Jonathans space pages: Starlink statistics. <https://planet4589.org/space/con/star/stats.html>, 8 2024. Accessed: 2024-08-1.
- [12] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. A first look at starlink performance. In *Proceedings of the 22nd ACM Internet Measurement Conference, IMC '22*, page 130–136, New York, NY, USA, 2022. Association for Computing Machinery.
- [13] Nitinder Mohan, Andrew E. Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K. Marina, and Jörg Ott. A multi-faceted look at starlink performance. In *Proceedings of the ACM on Web Conference 2024, WWW '24*, page 2723–2734, New York, NY, USA, 2024. Association for Computing Machinery.
- [14] Onweb: Connect with ease. https://onweb.net/connect_with_ease, 2024. Accessed: 2024-08-01.
- [15] Jianping Pan, Jinwei Zhao, and Lin Cai. Measuring a low-earth-orbit satellite network. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–6, 2023.
- [16] Refzan Pawirotaroen. Space internet congestion control (SICC). Master’s thesis, University of Twente, 2022.
- [17] Starlink Official X Page @Starlink. “Starlink is connecting more than 3M people with high-speed internet across nearly 100 countries, territories and many other markets. ...”. <https://x.com/Starlink/status/1792678386353213567>, 5 2024. Accessed: 2024-08-1.
- [18] Starlink. <https://www.starlink.com>, 2024. Accessed: 2024-08-01.
- [19] Aleksandr Stevens, Blaise Iradukunda, Brad Bailey, and Ramakrishnan Durairajan. Can leo satellites enhance the resilience of internet to multi-hazard risks? In Philipp Richter, Vaibhav Bajpai, and Esteban Carisimo, editors, *Passive and Active Measurement*, pages 170–195, Cham, 2024. Springer Nature Switzerland.
- [20] Hammam Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. Making sense of constellations: Methodologies for understanding starlink’s scheduling algorithms. In *Companion of the 19th International Conference on Emerging Networking Experiments and Technologies, CoNEXT 2023*, page 37–43, New York, NY, USA, 2023. Association for Computing Machinery.