LEO Satellite vs. Cellular Networks: Exploring the Potential for Synergistic Integration

Bin Hu*
University of Southern California

Xumiao Zhang*
University of Michigan

Qixin Zhang
University of Minnesota

Nitin Varyani
University of Minnesota

Z. Morley Mao
University of Michigan

Feng Qian
University of Southern California

Zhi-Li Zhang
University of Minnesota

ABSTRACT

Low-Earth-Orbit (LEO) satellite networks, such as Starlink, are transforming global network connectivity by bringing Internet access to remote and underserved areas. However, the current coverage and performance of the LEO satellite network service compared with those of cellular networks are under-explored. In this work, we present a measurement study of the Starlink LEO satellite network in comparison with cellular networks, aiming to uncover the potential for synergistic integration. Through a large-scale data collection campaign and in-depth analysis, we (1) identify the performance characteristics of two Starlink configurations, (2) evaluate the coverage of the current Starlink deployment compared to major cellular carriers, and (3) investigate the potential benefits of enabling multipath using both LEO satellite and cellular networks.

CCS CONCEPTS

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

KEYWORDS

Satellite Network, Cellular Network, Low Earth Orbit, Network Measurement, MPTCP

ACM Reference Format:
Bin Hu*, Xumiao Zhang*, Qixin Zhang, Nitin Varyani, Z. Morley Mao, Feng Qian, and Zhi-Li Zhang. 2023. LEO Satellite vs. Cellular Networks: Exploring the Potential for Synergistic Integration. In *Companion of the 19th International Conference on emerging Networking EXperiments and Technologies (CoNEXT Companion '23), December 5–8, 2023, Paris, France.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3624354.3630588

1 INTRODUCTION

Since its debut in 2020, SpaceX's revolutionary Starlink, a Low-Earth-Orbit (LEO) satellite network, has gained over 1 million users with its promise of delivering global connectivity [4]. It aims to provide services to remote and inaccessible areas where traditional wired or wireless networks fall short. However, both LEO satellite systems and cellular networks face distinctive challenges that impede their ability to consistently attain peak network performance. The seamless operation of Starlink requires an unobstructed view of the sky. On the other hand, the coverage of 4G/5G cellular networks heavily relies on the extensive deployment of base stations. Although some initial research has previously touched on Starlink's performance [21, 28, 29], there lacks a comprehensive investigation that compares the LEO satellite networks to terrestrial cellular networks. It is also necessary to explore the mobility of Starlink dishes. Additionally, the performance and accessibility of Starlink and cellular networks can often complement each other, and thus enabling multipath connections may bring superior throughput and reliability compared to relying on a single network.

To fill this gap, in this paper, we conduct a large-scale measurement study to understand the performance characteristics of Starlink satellite networks and compare the performance coverage between Starlink and cellular networks.

First, we collect a unique driving dataset comprising experimental results of both Starlink and cellular networks, including two types of Starlink configurations and three cellular carriers. During our data collection, we conduct a range of experiments on several network performance metrics. Our driving routes that span across five states in the US, consider different geographical areas, densities of infrastructure deployment, and user populations. Throughout the data collection process, we cover a total distance of over 3,800 km in over one month. The main challenge comes from how to collect data for both network types simultaneously and ensure apple-to-apple fair comparison. We address this by installing two Starlink dishes on the vehicle rooftop and carrying five smartphones set side by side. With our dataset, we aim to answer the following questions:

• **What is the performance achievable by Starlink networks, in particular under mobility?** Previous studies look at Starlink's performance using stationary hardware and none of them considers mobility. For stationary use, there are already numerous ways to access the Internet, while on the move people basically only rely on cellular networks. Therefore, we take a different view to examine Starlink under mobility and compare it with cellular networks. We evaluate key performance metrics of Starlink, such as throughput, latency, and packet loss. Through bulk transfer and ping tests conducted during driving sessions across diverse geographic locations, we evaluated the performance disparities between TCP and UDP,

* These authors contributed equally to this paper.

CoNEXT Companion '23, December 5–8, 2023, Paris, France

ACM ISBN 979-8-4007-0407-9/23/12…$15.00
https://doi.org/10.1145/3624354.3630588

uplink and downlink, and different Starlink configurations. Furthermore, we analyze the impact of different factors such as moving speed and TCP parallelism on network performance.

• **How does Starlink's performance and coverage compare to that of major cellular carriers?** To better understand whether Starlink offers a broader coverage than traditional cellular services. We compare the coverage and performance distribution of Starlink and cellular networks in different areas. We aim to provide insights into the viability of Starlink as a potential alternative to cellular networks for those who require reliable connectivity in the wild.

• **What is the potential of enabling multipath for Starlink and cellular networks?** The multipath technology was introduced to enhance network performance by utilizing multiple network paths simultaneously. Given the highly variable performance of Starlink and cellular networks, combining both networks may lead to a more satisfactory user experience. Thus, we conduct emulation using our collected network traces to compare the performance of single-path and multipath transports and understand the improvements brought by MPTCP in throughput and reliability.

We summarize our key findings as follows:

• Compared to cellular networks, Starlink suffers from elevated packet loss while the latency stays similar. We find that TCP severely suffers from such a high packet loss of Starlink, leading to only 1/5 of the throughput achieved by UDP over Starlink. TCP parallelism brings more benefits to Starlink than to cellular networks likely due to its effective handling of packet loss. Nonetheless, this finding also calls for better congestion control or Forward Error Correction (FEC) algorithms tailored for such characteristics.

• We compare two types of Starlink configurations, Roam and Mobility, both designed for use outdoors. While we do find better overall performance offered by the more expensive one (Mobility, having 2× higher mean/median throughput), its additional cost cannot be fully justified by current usages. The 75-percentile throughput of Roam, 93 Mbps, can already meet most application requirements in the wild.

• Due to the ultra-high speed operations of LEO satellites, the user's moving speed is negligible and thus poses little impact on Starlink's network performance, resulting in a similar trend across different speeds compared to cellular networks.

• Cellular networks offer better performance in urban areas thanks to the densely deployed base stations, while Starlink wins in suburban and rural areas with fewer obstructions. Since most of the time, the cellular services experiences are either low-band 5G or 4G LTE, the cellular throughput does not often reach very high. Starlink demonstrates better overall performance. However, even after combining cellular and Starlink, there are still areas with low performance (<50 Mbps), likely due to the combined effect from cellular base station deployment and obstruction to satellite connections.

• In our MPTCP experiments, we first find that, under the default OS configuration, MPTCP using Starlink and cellular networks brings marginal throughput gains compared to single-path transfer due to the high variation of both networks easily filling the buffer. After tuning the OS buffer settings, we see more significant improvements (up to 66% improvement over the better path), benefiting from the complementary characteristics between Starlink and cellular networks. This emphasizes the need for better MPTCP scheduler and congestion control algorithms.

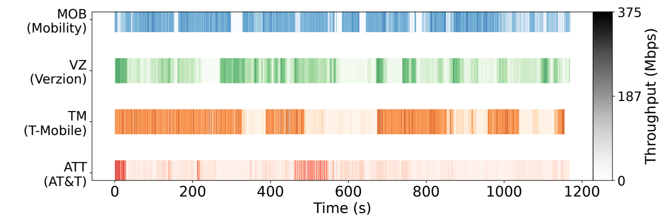

Figure 1: Download throughput of different networks.

To the best of our knowledge, this paper represents the first in-depth investigation that closely examines the performance of Starlink in comparison to cellular networks, particularly under mobility. Through our measurements and analysis, we have firmly supported the case for multi-path transport using LEO satellite networks and cellular networks. We have released all the measurement data and source code associated with this study [7].

2 MOTIVATION

LEO satellite networks, such as Starlink, utilize distinct communication and connectivity technologies that set them apart from traditional cellular networks. Unlike cellular networks which rely on terrestrial base stations, Starlink operates through a constellation of satellites. A user-side dish connects with a satellite, which in turn, communicates with a ground station. These ground stations relay data to and from the Internet. It is important to note that Starlink requires Line-of-Sight between user dishes and satellites. Obstructions such as tall buildings or trees can disrupt the satellite connections. Consequently, Starlink has better performance in open and remote areas. In contrast, cellular networks excel in densely populated areas where a dense deployment of base stations ensures reliable connectivity. Moreover, due to their differing deployment strategies and thus service availability, the two types of networks can exhibit highly varied performance and coverage characteristics.

To validate our hypothesis, we conduct a preliminary experiment using smartphones connected to Starlink routers via Wi-Fi and to cellular networks, in a moving vehicle. We perform `iPerf` data transfer tests to AWS servers and compare the throughput between Starlink and cellular networks. Our results are summarized in Figure 1, where darker colors (blue/green/orange/red) indicate periods of higher throughput. As we traversed different areas, we can observe instances where Starlink demonstrated better throughput performance compared to the cellular network, and vice versa. If combining the strengths of both networks, users can enjoy a seamless and stable high-performance experience throughout their journey. This motivates us to conduct further investigations to comprehensively understand the performance and coverage of both network types in real-world scenarios, and to assess the feasibility of multi-path transport for LEO satellite and cellular networks.

3 MEASUREMENT METHODOLOGY

3.1 Hardware and Services

We first introduce the hardware setups and network services used.

Starlink offers a variety of plans, including Residential, Business, Roam, and Mobility, among others. In this study, our focus is on the Roam (RM, for short) and Mobility (MOB) plans, specifically designed for portable and in-motion use. The Roam plan offers easy portability compared to a standard Starlink dish and provides

Figure 2: Two dishes are mounted on the rooftop and five smartphones are placed side by side in the vehicle.

Internet access while stationary or on the move. The Mobility plan, on the other hand, receives the highest priority in the network, for instance, during congestion. It is specifically designed to support critical in-motion applications, such as those used by emergency personnel. Besides, Mobility has over 4× the hardware cost and a higher monthly fee than Roam. For a direct comparison between the two plans, we have installed both Starlink dishes on the rooftop of a vehicle. However, we acknowledge the possibility of interference that may affect the results. For experiments involving cellular networks, we have selected three major commercial carriers in the US: AT&T (ATT), T-Mobile (TM), and Verizon (VZ). We utilize five Samsung Galaxy S21 smartphones. Three of these devices are connected to the cellular services, while the remaining two are connected to the Starlink dishes using Wi-Fi.

Figure 2 shows our dish placement and smartphone setup.

3.2 Software Measurement Tools

We utilize several software tools for data collection: (1) We use `iPerf` to run TCP/UDP downlink and uplink data transfer tests. (2) To measure latency accurately, we have developed an Android application that sends ping packets using UDP (`UDP-Ping`), as ICMP ping packets are often blocked by certain servers. (3) To collect information on network type, vehicle speed, GPS location, and signal strength, we employ 5G Tracker [30, 31], a monitoring toolkit for cellular networks. We have made modifications to enable its functionality under both Wi-Fi and cellular connectivity.

3.3 Data Collection

We perform extensive drive tests across major cities and interstate freeways (spanning five states) in the US. It encompasses diverse geographical regions, including densely populated urban areas with tall buildings and open rural areas with minimal obstructions. We drive at varying speeds in various areas. However, our driving speed is capped at 100 km/h due to speed limits on different road segments. The driving routes consist of both straight and curved roads, aiming to generalize our results with regard to vehicle steering. We collect data during both daytime and nighttime. Furthermore, our data collection includes not only clear weather conditions but also rainy and snowy conditions, to capture potential performance variations caused by environmental conditions. Note that, despite the breadth of factors considered, not all are explicitly discussed in the following sections. Upon analysis, certain environmental factors such as terrain and the time of day, along with network-related factors such as server locations, are found to have a minimal impact on the network performance.

Our driving trip yields a unique driving dataset, containing 1,239 network tests and 9,083 minutes of traces. Our field trip covers a total travel distance of over 3,800 km.

4 STARLINK BASIC PERFORMANCE

In this section, we analyze and compare the performance characteristics of Starlink satellite networks with cellular networks, focusing on two Starlink configurations, Roam and Mobility. We evaluate fundamental performance metrics including throughput, latency, and packet loss, and conduct an in-depth examination of the performance gaps from different aspects. Besides, we also examine the impact of vehicle speed and TCP parallelism on performance.

4.1 Throughput, Latency, and Packet Loss

We start with analyzing the key performance metrics of Starlink.

TCP vs. UDP downlink. Figure 3a plots the Cumulative Distribution Functions (CDFs) for TCP and UDP downlink throughput of both Starlink (Mobility) and cellular (AT&T, T-Mobile, and Verizon) networks. While the performance disparity between cellular TCP and UDP is minimal, the results consistently reveal that UDP outperforms TCP in satellite networks, with the mean throughput being 128 Mbps and 29 Mbps, respectively. This performance advantage of UDP can be attributed to the significant packet loss experienced by TCP in satellite networks. To substantiate this proposition, we analyze the `Tcpdump` traces collected while running `iPerf` and plot the average TCP packet loss across all networks in Figure 5. When using Starlink, there is a much higher occurrence of packet loss in both the uplink and downlink directions, compared to cellular networks. This leads to retransmissions ranging from 0.3% to 1.3%. Such elevated packet loss significantly impacts TCP performance and ultimately decreases Starlink's TCP throughput. Additionally, it is important to note that the UDP performance achieved with the Mobility plan demonstrates a level of throughput that is comparable to that of cellular networks, highlighting its effectiveness in facilitating data transfer.

Roam vs. Mobility. Figure 3b compares the network performance between the Starlink Roam and Mobility plans. The Mobility plan exhibits superior performance compared to Roam, likely because Roam's dish lacks the ability to adjust its orientation promptly under high mobility while Mobility is designed for in-motion use with a wider field of view. This may also benefit from the advertised prioritization for Mobility during network congestion. The median/mean throughput for Mobility and Roam are 197/128 Mbps and 93/63 Mbps, respectively. However, such 2x performance improvements are not that significant compared to the over 4x higher cost on the hardware [5], since the network requirements of most applications such as 1080P video streaming can already be met by Roam. Unless for critical applications with demanding requirements, the more cost-friendly Roam plan can effectively serve as a viable alternative to the Mobility plan.

Uplink vs. downlink. Comparing the UDP uplink and downlink transfer of Starlink, we find that the downlink throughput is around 10x higher than the uplink, as depicted in Figure 3c. This design choice of using FDD for dividing uplink and downlink channels [19] aligns with the inherent characteristics of network traffic, where users typically consume more data in the form of downloads rather than uploads. Additionally, Starlink's satellite dishes are optimized to prioritize downlink speed over uplink speed due to limited power resources and transmitting capacity. It is more energy-efficient to receive a signal than to transmit it [20, 36]. We also learn that, by

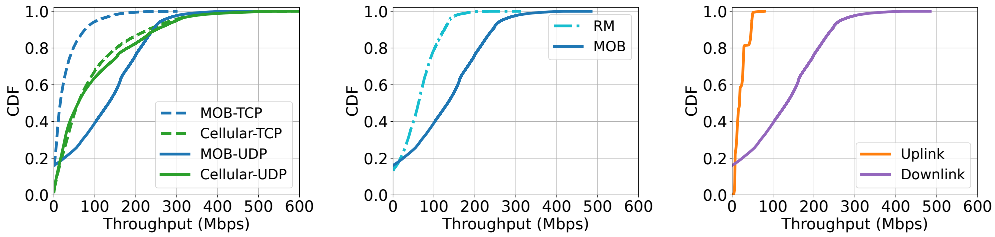

(a) TCP vs. UDP. **(b) Roam vs. Mobility.** **(c) Uplink vs. Downlink.**

Figure 3: Throughput performance comparison from different aspects.

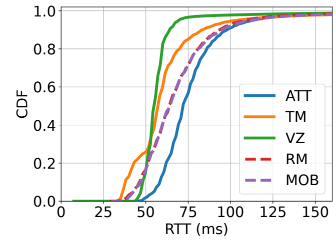

Figure 4: UDP Ping Latency.

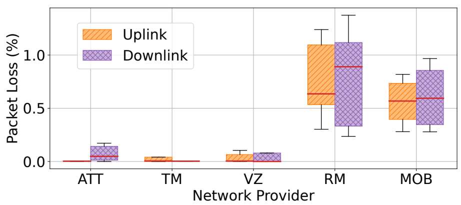

Figure 5: Packet loss in TCP transfer.

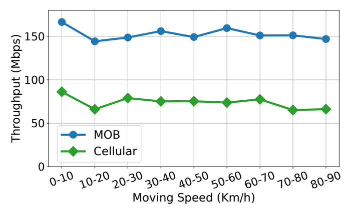

Figure 6: Impact of speed.

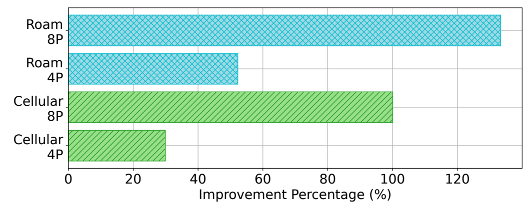

Figure 7: Impact of TCP parallelism.

design, the Starlink network offers higher downlink bandwidth than the uplink [6].

Latency. We utilize our Android application, `UDP-Ping`, to measure latency. We allocate 1024 bytes to each UDP packet and calculate the round-trip time (RTT) for each acknowledged packet. Figure 4 illustrates the CDF of latency for five different networks. Overall, the RTTs for all networks primarily fall within the range of 50 to 100ms. Verizon and T-Mobile exhibit the lowest RTT values, while Starlink Roam and Starlink Mobility plans experience comparatively higher latency. It is surprising to see that Starlink's latency is not significantly worse than that of cellular networks. Intuitively, satellite networks should incur significantly higher latency due to the long satellite-ground distance. However, the additional latency introduced by satellite data transmission is only approximately 1.8 ms one way, thanks to the "low-earth" orbiting nature. The estimation is derived from the LEO satellites' orbital altitude of around 550 km and the speed of light:

$$\text{Latency} = \left(\frac{\text{Distance}}{\text{Speed of light}}\right) = \left(\frac{550 \text{ km}}{299792 \text{ km/s}}\right) = 1.835\ ms \quad (1)$$

This acceptable level of latency indicates that satellite networks can provide reliable performance comparable to cellular networks. Notably, AT&T demonstrates the highest network latency among the tested networks, likely due to its relatively low coverage along our trip [1].

4.2 Potential Factors Affecting Throughput

We discuss the impact of two factors, moving speed and TCP parallelism, on the throughput performance of Starlink.

Moving speed. To ensure an unbiased analysis and isolate other factors, we specifically extract data collected in rural areas which offer minimal obstructions. This methodology mitigates the challenges associated with conducting high-speed driving tests in urban environments, where speed limits restrict the full exploration of network performance. More than 90% of our urban data were collected at speeds below 50 km/h. Moreover, the satellite connections can be negatively impacted by obstructions, while cellular networks may exhibit better performance in urban areas.

Figure 6 shows the average throughput grouped by speed. Notably, both satellite (Mobility) and cellular (AT&T, T-Mobile, and Verizon) network throughputs have minimal variation in relation to driving speed. This suggests that the network performance remains largely unaffected by the vehicle's speed during normal driving conditions. Considering Starlink's operation in low earth orbit at an approximate speed of 28,000 km/h, the speed of an object on the ground is negligible and can be considered stationary. For cellular networks, on the other hand, efficient handovers contribute to maintaining consistent throughput.

TCP parallelism. TCP parallelism is a technique that enables parallel TCP connections between senders and receivers to increase throughput. Our experiments compare three schemes: 1, 4, and 8 TCP connections. Figure 7 demonstrates the improvement achieved by TCP parallelism on downlink throughput for both satellite (Roam) and cellular networks. "P" denotes parallelism, where '8P' represents 8 parallel connections. Increasing the number of parallel TCP connections enhances throughput in both networks. Starlink achieves a better throughput improvement, over 50% with 4 parallel TCP connections and over 130% improvement with 8 connections. TCP parallelism optimizes bandwidth utilization by distributing data across multiple connections, thereby mitigating the impact of TCP congestion control. It also improves packet loss handling. In case of packet loss in one connection, other connections continue data transmission, minimizing the impact on overall throughput. Given the higher packet loss observed in the Starlink network (§4.1), increasing TCP parallelism enables more efficient handling of packet loss, resulting in improved throughput.

5 COVERAGE STUDY

This section focuses on the coverage area of Starlink. We discuss the network performance in different geographical regions and the proportion of coverage within each performance level.

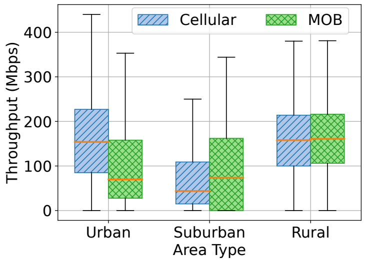

Figure 8: Downlink throughput at different area types.

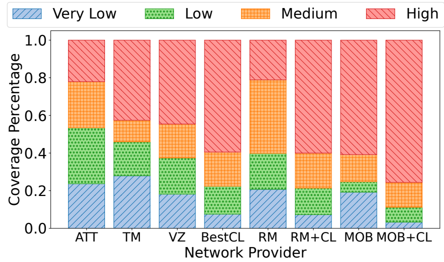

Figure 9: Comparison of network performance coverage.

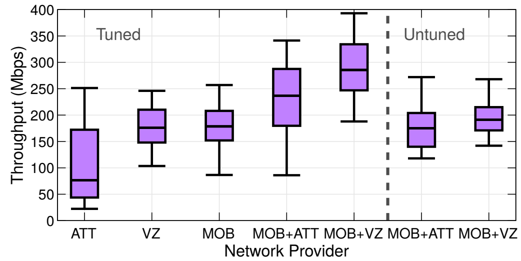

Figure 10: Single-path TCP and MPTCP data download performance.

5.1 Impact of Area Types

Deploying and operating cellular base stations in rural areas incurs much higher costs due to low population density [3], while users in urban areas enjoy more reliable cellular network connections thanks to the dense base station deployment. For Starlink, urban areas with tall buildings can obstruct satellite signal transmission. Thus, the location of the network affects both Starlink and cellular networks.

During data collection, we traversed through different types of areas, recording the latitude and longitude of each data point. Then we compile a list of all cities and towns we passed through, calculate the distances from each data point to these locations, and select the smallest distance. Subsequently, using predetermined thresholds, we categorize the data into three area types: urban, suburban, and rural. The data proportion of the three areas is 29.78%, 34.30%, and 35.91%, respectively.

Figure 8 shows the throughput distributions for both Starlink (Mobility) and cellular networks. Here, we highlight the results of UDP downlink since Starlink's downlink throughput is inherently higher and UDP is less affected by packet loss compared to TCP. They also reflect the upper limit of network bandwidth. It can be observed that the throughput of cellular networks decreases when reaching rural areas, while the throughput of Starlink networks increases in rural areas. This is because cellular network base stations are more densely deployed in populated areas, whereas in densely populated areas, the network performance of Starlink can be affected by obstacles such as tall buildings. From our analysis of UDP Downlink, we find that the throughput of Starlink networks is even higher than that of cellular networks in suburban and rural areas. We also find that the throughput of Starlink is distributed similarly in suburban and rural areas. During our driving trip, we found a lot of obstructions only in urban areas. Suburban areas such as towns have much fewer high buildings, leading to similar obstruction conditions to rural areas.

5.2 Performance Coverage

To visually represent the coverage of Starlink and cellular networks, we analyze our data and group data points of different network performance based on different performance levels. The high-performance regions are characterized by throughput exceeding 100 Mbps, while the medium-performance regions exhibit throughput ranging between 50 and 100 Mbps. The low-performance regions have a throughput between 20 and 50 Mbps. Additionally, we consider a "very-low" performance level, where the throughput is under 20 Mbps, to understand if there are regions with extremely poor coverage for the different technologies. Although our data was collected during certain periods in each region and might not reflect the long-term network performance in specific regions, it still provides insights into the coverage of network performance of Starlink and cellular networks. The results of our analysis are presented in Figure 9, showcasing the proportions of different performance regions for the five networks.

We can learn that Starlink Mobility exhibits the best overall performance, with a proportion of high-performance regions at 60.61%. Verizon and T-Mobile closely follow, with proportions of high-performance regions at 44.39% and 42.47%, respectively. Starlink Roam and AT&T, however, demonstrate the poorest performance, with proportions of low and very-low performance regions approximately at 39.88% and 53.45%, respectively. We plot a bar named BestCL which indicates the best performance of all three cellular networks. This is reasonable since many mobile virtual network operators (MVNOs) utilize the services of several mobile carriers and automatically pick the best option for users. We also combine the measurement records of Starlink and cellular as shown as the bars, RM+CL and MOB+CL, in Figure 9. Their improvements over a single cellular network are likely due to the significant presence of non-urban areas with minimal obstructions, which currently favor Starlink's performance. Noticeably, combining all cellular networks also leads to results comparable to RM+CL. As mentioned earlier, compared to Roam which is not designed for mobile use, the Mobility dish has a wider field of view and better positioning capability, resulting in Mobility having the overall best coverage of network performance regions.

We also plot two additional bars representing the best performance if a user has access to both Starlink and cellular networks (and can switch between them with zero effort). From a user experience perspective, this highlights the importance of implementing multipath for Starlink and cellular networks. Due to their inherent differences, Starlink and cellular networks exhibit significant variations in the coverage of high-performance regions. Compared with the original measurement records, these combinations achieve better high-performance network coverage. We are encouraged to explore the multipath feasibility in the next section.

6 MULTIPATH TRANSPORT

Multipath transport, in particular MPTCP, has shown its success in numerous combinations of networks [11, 26, 34] and various applications [14, 15, 33, 40]. In this section, through realistic emulation, we demonstrate the potential of enabling multipath for Starlink and cellular networks.

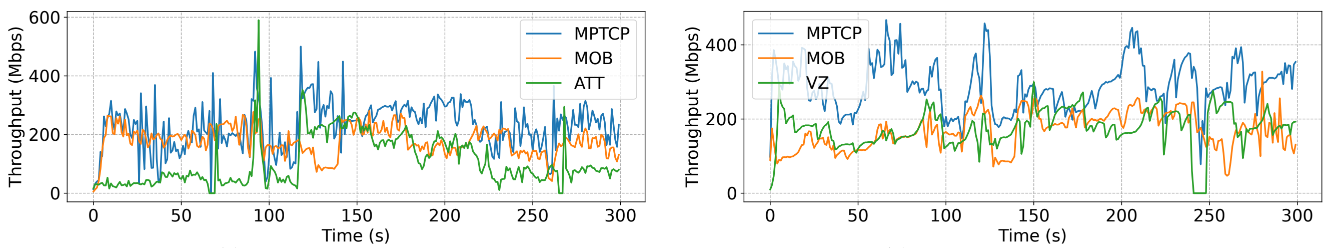

(a) Mobility and AT&T. (b) Mobility and Verizon.

Figure 11: Throughput traces for single-path TCP and MPTCP data download.

Experimental setup. We run MPTCP experiments on two Ubuntu 22.04 hosts, using `MpShell` (a variant of Mahi-mahi [12, 32]). It creates multiple virtual interfaces with controlled network conditions, using packet traces and latency statistics. To this end, we use the UDP downlink throughput traces in our driving dataset and convert them to packet traces for replay on `MpShell`. Different network traces are aligned via timestamps so that they reflect the network conditions experienced by users at the same location and time. Note that we opt for UDP data instead of TCP data to emulate the available bandwidth at each timestamp and avoid the impact of TCP congestion control. For multipath transport, we run a modified version of `iPerf` [2] that supports MPTCP. It opens MPTCP sockets instead of regular TCP sockets. To compare the performance between MPTCP and each single-path transport, we run two types of experiments in the `MpShell` environment: (1) We start two `iPerf` client instances on the client machine and two server instances on the server side. Each client downloads data using one network interface from the `iPerf` server; (2) We start an `iPerf` client with the MPTCP option enabled and a `iPerf` server. The client downloads data from the server using MPTCP.

Results. Figure 10 presents the performance of 5-min download tests. The first three boxes represent the single-path TCP transfer results under AT&T, Verizon, and Mobility networks. The next two show the MPTCP results (Mobility+AT&T, Mobility+Verizon) when concurrently using a Starlink and a cellular service. The benefits of MPTCP are clear; it improves the overall download throughput. On average, the bandwidth utilization of the two tested combinations is 81% and 84%, and the improvement over the better path reaches 30% and 66%, respectively.

We also put another two boxes showing the multipath results before tuning the system buffer. Initially, we notice that, with the default buffer sizes, MPTCP has marginal improvements over single-path transfers. In some cases, the throughput collapses to "0", leading to the failure of the `iPerf` test. Therefore, we increase the buffer size to exceed 10× the link's bandwidth-delay product to accommodate such network fluctuations.

Looking into the throughput progression over time in Figure 11, we can learn that MPTCP almost always outperforms either single-path transfer, taking advantage of the bandwidth of the faster path. For example, in Figure 11a, between 0-85s, AT&T experiences severe performance degradation, likely due to weak cellular signal strength. With MPTCP, the throughput is maintained at a much higher level. Also, in Figure 11b, when both network conditions are favorable at around 50-90s, MPTCP throughput exceeds 300 Mbps which can never be achieved by either network alone.

The current MPTCP experiments are conducted via emulation and we leave developing a MPTCP scheduler for LEO satellite networks and running real MPTCP experiments as future work. The default MPTCP scheduler implemented in the OS (kernel v5.19) is BLEST [13] which optimizes MPTCP send window occupation to avoid transport-layer head-of-line blocking. We envision that, considering the specific usage scenarios and characteristics of the two network types, further improvements can be made to future MPTCP scheduler design, such as reducing throughput fluctuations.

7 RELATED WORK

Compared to widely deployed mobile networks like 4G/5G [18, 31], LEO satellite networks are relatively new and have not been extensively studied in large-scale commercial deployments. Some works take the first step to evaluate the performance of LEO satellite networks. Michel *et al.* [29] compare Starlink with SatCom and a wired network. Kassem *et al.* [21] analyze Starlink connectivity from a browser-side view. Ma *et al.* [28] present initial measurement results on Starlink's network characteristics. Li *et al.* [25] evaluate the network impact of Starlink's self-driving for its satellites. Our study differs from these, focusing on Starlink's network performance under mobility and its comparison with cellular networks.

Other research has offered insights into future LEO satellite networks and the integration of satellite and terrestrial networks [17, 23, 27]. Some focus on routing [10, 16, 38] and network topology design [9, 35, 39], as well as handovers [8]. L2D2 [37] is a satellite ground station design for low-latency downlink transmission. Tools and platforms are proposed for evaluating LEO satellite networks [8, 22, 24]. In our paper, we explore the potential of combining Starlink and cellular networks for better network performance.

8 CONCLUSION

We present a measurement study on the performance of Starlink satellite and cellular networks. Rigorous analyses reveal that Starlink outperforms cellular networks in open areas, but suffers from higher packet loss leading to degraded TCP performance. Their complementary characteristics offer potential improvements to network connectivity but require optimizations. We hope that our findings can spur more research to improve LEO satellite networks.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable feedback. This work was supported in part by NSF under Grants CNS-1915122, CNS-1930041, CMMI-2038215, CNS-2112562, CNS-2128489, CNS-2154078, CNS-2220286, CNS-2220292, and CNS-2321531.

REFERENCES

[1] 2021. FCC Mobile Broadband Maps. https://tinyurl.com/2ayv5rz6.

[2] 2021. iPerf GitHub repository - Add MPTCP support with the –multipath flag. https://github.com/esnet/iperf/pull/1166/commits.

[3] 2022. Accelerating Rural Connectivity. https://tinyurl.com/56jk4bx7.

[4] 2022. Starlink has hit more than 1 million users despite a drop in download speeds. Here's what you need to know about the service. https://www.businessinsider.com/spacex-starlink-internet-service-elon-musk-all-you-need-know-2021-2.

[5] 2023. Starlink. https://www.starlink.com/.

[6] 2023. Starlink Specifications. https://www.starlink.com/legal/documents/DOC-1400-28829-70.

[7] 2023. Starlink-vs-Cellular GitHub repository. https://github.com/Starlink-Project/Satellite-vs-Cellular.

[8] Akram Al-Hourani. 2021. A tractable approach for predicting pass duration in dense satellite networks. *IEEE Communications Letters* 25, 8 (2021), 2698–2702.

[9] Debopam Bhattacherjee and Ankit Singla. 2019. Network topology design at 27,000 km/hour. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*. 341–354.

[10] Vaibhav Bhosale, Ahmed Saeed, Ketan Bhardwaj, and Ada Gavrilovska. 2023. A Characterization of Route Variability in LEO Satellite Networks. In *International Conference on Passive and Active Network Measurement*. Springer, 313–342.

[11] Andrei Croitoru, Dragos Niculescu, and Costin Raiciu. 2015. Towards wifi mobility without fast handover. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*. 219–234.

[12] Shuo Deng, Ravi Netravali, Anirudh Sivaraman, and Hari Balakrishnan. 2014. WiFi, LTE, or both? Measuring multi-homed wireless internet performance. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. 181–194.

[13] Simone Ferlin, Özgü Alay, Olivier Mehani, and Roksana Boreli. 2016. BLEST: Blocking estimation-based MPTCP scheduler for heterogeneous networks. In *2016 IFIP networking conference (IFIP networking) and workshops*. IEEE, 431–439.

[14] Yihua Ethan Guo, Ashkan Nikravesh, Z Morley Mao, Feng Qian, and Subhabrata Sen. 2017. Accelerating multipath transport through balanced subflow completion. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. 141–153.

[15] Bo Han, Feng Qian, Lusheng Ji, and Vijay Gopalakrishnan. 2016. MP-DASH: Adaptive video streaming over preference-aware multipath. In *Proceedings of the 12th International on Conference on emerging Networking EXperiments and Technologies*. 129–143.

[16] Mark Handley. 2019. Using ground relays for low-latency wide-area routing in megaconstellations. In *Proceedings of the 18th ACM Workshop on Hot Topics in Networks*. 125–132.

[17] Simon Heine, Florian Völk, Robert T Schwarz, and Andreas Knopp. 2022. Concept and evaluation of 5G backhauling via starlink. In *39th International Communications Satellite Systems Conference (ICSSC 2022)*, Vol. 2022. IET, 69–75.

[18] Junxian Huang, Feng Qian, Yihua Guo, Yuanyuan Zhou, Qiang Xu, Z Morley Mao, Subhabrata Sen, and Oliver Spatscheck. 2013. An in-depth study of LTE: Effect of network protocol and application behavior on performance. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 363–374.

[19] Peter Anthony Iannucci and Todd E Humphreys. 2022. Fused low-Earth-orbit GNSS. *IEEE Trans. Aerospace Electron. Systems* (2022).

[20] Nabil Jardak and Ronan Adam. 2023. Practical use of Starlink downlink tones for positioning. *Sensors* 23, 6 (2023), 3234.

[21] Mohamed M Kassem, Aravindh Raman, Diego Perino, and Nishanth Sastry. 2022. A browser-side view of starlink connectivity. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 151–158.

[22] Simon Kassing, Debopam Bhattacherjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. 2020. Exploring the" Internet from space" with Hypatia. In *Proceedings of the ACM Internet Measurement conference*. 214–229.

[23] Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. 2023. {StarryNet}: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 1309–1324.

[24] Zeqi Lai, Hewu Li, and Jihao Li. 2020. Starperf: Characterizing network performance for emerging mega-constellations. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*. IEEE, 1–11.

[25] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Wei Zhao, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, Zeqi Lai, et al. 2023. A networking perspective on starlink's self-driving leo mega-constellation. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*. 1–16.

[26] Yeon-sup Lim, Erich M Nahum, Don Towsley, and Richard J Gibbens. 2017. ECF: An MPTCP path scheduler to manage heterogeneous paths. In *Proceedings of the 13th international conference on emerging networking experiments and technologies*. 147–159.

[27] Melisa López, Sebastian Bro Damsgaard, Ignacio Rodríguez, and Preben Mogensen. 2022. An Empirical Analysis of Multi-Connectivity between 5G Terrestrial and LEO Satellite Networks. In *2022 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 1115–1120.

[28] Sami Ma, Yi Ching Chou, Haoyuan Zhao, Long Chen, Xiaoqiang Ma, and Jiangchuan Liu. 2023. Network Characteristics of LEO Satellite Constellations: A Starlink-Based Measurement from End Users. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 1–10.

[29] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. 2022. A first look at starlink performance. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 130–136.

[30] Arvind Narayanan, Eman Ramadan, Jacob Quant, Peiqi Ji, Feng Qian, and Zhi-Li Zhang. 2020. 5G tracker: a crowdsourced platform to enable research using commercial 5g services. In *Proceedings of the SIGCOMM'20 Poster and Demo Sessions*. 65–67.

[31] Arvind Narayanan, Xumiao Zhang, Ruiyang Zhu, Ahmad Hassan, Shuowei Jin, Xiao Zhu, Xiaoxuan Zhang, Denis Rybkin, Zhengxuan Yang, Zhuoqing Morley Mao, et al. 2021. A variegated look at 5G in the wild: performance, power, and QoE implications. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 610–625.

[32] Ravi Netravali, Anirudh Sivaraman, Somak Das, Ameesh Goyal, Keith Winstein, James Mickens, and Hari Balakrishnan. 2015. Mahimahi: Accurate Record-and-Replay for HTTP. In *Usenix annual technical conference*. 417–429.

[33] Ashkan Nikravesh, Yihua Guo, Xiao Zhu, Feng Qian, and Z Morley Mao. 2019. MP-H2: a Client-only Multipath Solution for HTTP/2. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–16.

[34] Swetank Kumar Saha, Shivang Aggarwal, Rohan Pathak, Dimitrios Koutsonikolas, and Joerg Widmer. 2019. MuSher: An agile multipath-TCP scheduler for dual-band 802.11 ad/AC wireless LANs. In *The 25th Annual International Conference on Mobile Computing and Networking*. 1–16.

[35] Vaibhav Singh, Akarsh Prabhakara, Diana Zhang, Osman Yağan, and Swarun Kumar. 2021. A community-driven approach to democratize access to satellite ground stations. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 1–14.

[36] Starlink Hardware. 2023. How Much Power Does Starlink Use? https://www.starlinkhardware.com/how-much-power-does-starlink-use/.

[37] Deepak Vasisht, Jayanth Shenoy, and Ranveer Chandra. 2021. L2D2: Low latency distributed downlink for LEO satellites. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 151–164.

[38] Ruibo Wang, Mustafa A Kishk, and Mohamed-Slim Alouini. 2023. Reliability Analysis of Multi-hop Routing in Multi-tier LEO Satellite Networks. *arXiv preprint arXiv:2303.02286* (2023).

[39] Yaoying Zhang, Qian Wu, Zeqi Lai, and Hewu Li. 2022. Enabling Low-latency-capable Satellite-Ground Topology for Emerging LEO Satellite Networks. In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*. IEEE, 1329–1338.

[40] Xiao Zhu, Jiachen Sun, Xumiao Zhang, Y Ethan Guo, Feng Qian, and Z Morley Mao. 2020. Mpbond: efficient network-level collaboration among personal mobile devices. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*. 364–376.