Trajectory-based Serving Satellite Identification with User Terminal's Field-of-View

Ali Ahangarpour
aliahangarpour@uvic.ca
University of Victoria
Victoria, BC, Canada

Jinwei Zhao
clarkzjw@uvic.ca
University of Victoria
Victoria, BC, Canada

Jianping Pan
pan@uvic.ca
University of Victoria
Victoria, BC, Canada

ABSTRACT

Low-Earth-Orbit (LEO) satellite networks, such as SpaceX's Starlink, achieved global broadband Internet coverage with significantly lower latency and higher throughput than traditional satellite Internet service providers utilizing geostationary satellites. Despite the substantial advancements, the research community lacks detailed insights into the internal mechanisms of these networks. This paper presents the first systematic study of Starlink's obstruction map and serving satellite identification. Our method achieves almost unambiguous satellite identification by incorporating satellite trajectories and proposing an accurate Field-of-View (FOV) estimation approach. We validate our methodology using multiple Starlink dishes with varying alignment parameters and latitudes across different continents. We utilize Two-Line Element data to identify the available satellites within the user terminal's FOV and examine their characteristics, in comparison to those of the serving satellites. Our approach revealed a correlation between the satellite selection strategy and the user terminal to gateway latency. The findings contribute to the broader understanding of the architecture of LEO satellite networks and their impact on user experience.

CCS CONCEPTS

• **Networks** → **Network measurement**.

KEYWORDS

Low Earth Orbit Satellite Networks, Satellite Identification, Field-of-View

ACM MobiCom '24, November 18–22, 2024, Washington D.C., DC, USA

ACM ISBN 979-8-4007-0489-5/24/11
https://doi.org/10.1145/3636534.3696222

ACM Reference Format:
Ali Ahangarpour, Jinwei Zhao, and Jianping Pan. 2024. Trajectory-based Serving Satellite Identification with User Terminal's Field-of-View. In *The 2nd ACM Workshop on LEO Networking and Communication 2024 (ACM MobiCom '24), November 18–22, 2024, Washington D.C., DC, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3636534.3696222

1 INTRODUCTION

Low-Earth-Orbit (LEO) satellite networks (LSNs), exemplified by SpaceX's Starlink, Eutelsat's OneWeb, Telesat's Lightspeed, and Amazon's Kuiper, have revolutionized the truly global broadband Internet coverage around the Earth, especially connecting the remote and rural regions. With more than 6,000 active satellites and 3 million subscribers across nearly 100 countries and regions as of May 2024 [2], Starlink is the leading LEO satellite mega-constellation. However, most of the inner workings of these LSNs are unknown to the research community. Various efforts have revealed the architecture of Starlink's access and global backbone networks [5]. These studies have also explored user-perceived and externally observed performance over the past few years through both *inside-out* and *outside-in* measurements [4][11], as well as through crowdsourced data from social media [3][8]. However, detailed knowledge of the satellite selection strategy and identifying serving satellites for user terminals remain limited. A better understanding and optimization of satellite selection strategy can lead to improved service quality, especially as LEO constellations expand in size and complexity.

Previous research has observed the existence of the 15-second satellite handover interval over Starlink networks [9]. In particular, the handover events between the Starlink user terminal (UT) and satellites happen at the 12th, 27th, 42nd, and 57th seconds of each minute, synchronized globally. Extensive network measurements revealed the fluctuations of latency and throughput across different time slots, which have a significant impact on the transport layer and application layer performance [10][12]. However, the research community has not yet fully understood the time-series correlation between the satellite handover events and the user-perceived performance, namely the latency and throughput patterns. The Starlink UT previously exported the ID of the

serving satellite, as well as the service cell ID and gateway ID, through its gRPC interface. Due to various reasons, Starlink has removed such information through firmware updates. Such information is crucial to understand the satellite selection strategy and handover mechanism, as well as understanding the Starlink network architecture and performance in particular and LSNs in general. Currently, Starlink uses space tunnels to transmit the data packets from UTs through one or multiple satellites with inter-satellite links (ISLs) to reach the landing ground station (GS). From there, ground tunnels bring the data packets to the user's home point-of-presence (PoP) to access the regular Internet. The space and ground tunnel together appear to be one hop at the IP layer for any user-oriented measurement [6]. Without Starlink's exported satellite and gateway ID, inferring the serving satellite is the best we can achieve for the research community, although we still call for Starlink to reveal the ID of the serving satellite and gateway ID.

Based on our discussion, Hammas et al. [9] and Liz et al. [4] utilized the Starlink gRPC obstruction map to infer the serving satellite, and provided initial findings on the correlation between the satellite connectivity and the user-perceived performance. However, at the time of their publications, the orientation of the obstruction map obtained from UT's gRPC interface was inconsistent with the mobile app's 3D view, and the UT required periodic rebooting to clear the obstruction map, which limited the capability of in-depth and systematic measurements.

This paper is the first systematic, in-depth study of Starlink satellites in the view of user dishes. It represents the best understanding so far by the research community on the architecture of Starlink networks. We present the first comprehensive study of Starlink obstruction maps and incorporate satellite trajectories to identify the serving satellites, leading to a much more unambiguous satellite identification. Furthermore, this work is the first to unveil and analyze the functional Field-of-View (FOV) of the UT and accurately determine the available satellites within the FOV. Our analysis also revealed the correlation between the serving satellite's position and the UT to gateway latency. We evaluated the identification accuracy and validated our methodology using multiple dishes located around the world, including those in Australia, Asia, North America, and Europe, across various latitudes and ground infrastructures.

The rest of the paper is organized as follows. In Section 2, we provide a brief overview of previous research attempts to understand the FOV and identify the serving satellites. In Section 3, we detail our methodology to calculate the UT's FOV and the process to identify the serving satellites. In Section 4, we compare the characteristics of the available and serving satellites to determine the satellite selection strategy and observe the correlation between these characteristics and the UT to gateway latency. Finally, we conclude the paper and discuss future work in Section 5.

2 BACKGROUND AND RELATED WORKS

The Starlink UT's gRPC interface exposes various diagnostic data and performance metrics, including basic device information, the obstruction map, and alignment parameters. Previous Starlink UT firmware exposed the `dishGetContext` method, which included `cellId`, `satelliteId`, `gatewayId`, and other metrics related to the serving satellite and gateway. Since SpaceX disabled public access to this gRPC method through firmware updates[1], the research community has been reverse engineering to unveil the inner workings of the satellite connectivity for the UT.

Hammas et al. [9] first utilized the obstruction maps obtained from the gRPC interface and the CelesTrak [1] Two-Line Element (TLE) data for Starlink satellites to infer the serving satellite. In their paper, they assumed that a UT can connect to any satellite with an elevation angle greater than 25°, without taking into account the real-world alignment parameters or the orientation of the UT, and thus did not identify all available satellites within the FOV. Liz et al. [4] also used the gRPC obstruction map to infer the serving satellite. However, they focused solely on the satellite's location within the obstruction map and only provided a brief discussion on satellite connectivity. They only correlated the latency performance with the perceived satellite location in the obstruction map, rather than the actual satellite location and trajectory calculated from the TLE data. Both works required periodically rebooting the UT to reset the obstruction map to achieve accurate calculation, which limited in-depth and systematic measurements. Since their publications, Starlink has made improvements to the gRPC interface through firmware updates, including the ability to reset the obstruction map directly without rebooting the dish.

3 METHODOLOGY

In this section, we present our novel methodology to identify the available and serving satellites for any Starlink UT, given the alignment parameters obtained from the gRPC interface and the GPS location of the UT. These alignment parameters include the UT's tilt and boresight azimuth angle illustrated in Figure 1 along with the satellite's topocentric coordinates (elevation and azimuth). The identification of the serving satellites involves a multi-step process that begins with collecting and processing the obstruction map over time, followed by identifying the available satellites

[1] https://tysonpower.de/blog/spacex-shutdown-part-of-dishys-api-because-of-me-and-others

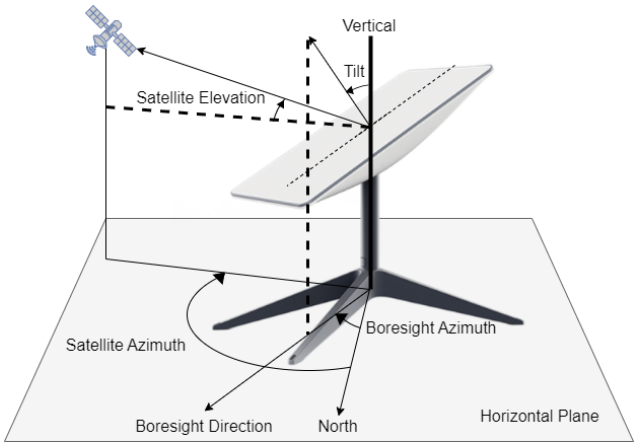

Figure 1: The illustration of alignment parameters of a Starlink UT and satellite's topocentric coordinates

within the FOV using the TLE data. Subsequently, the observed satellite trajectories are correlated with available satellites to determine the serving satellite. Although the dataset of identified available and serving satellites for the UTs used in this paper may evolve, our implementation[2] is portable and capable of running in different locations with real alignment parameters, ensuring up-to-date data collection and results.

3.1 Obstruction Data

The obstruction map of a Starlink UT can be obtained from the gRPC interface using open source tools such as `grpcurl`[3]. The obstruction map represents the UT's view of the sky, as a 123 by 123 pixel 2D image. By locating the UT's FOV reference point in this image and mapping the pixels to topocentric coordinate values, we can infer the position of the serving satellite relative to the UT's location. An example of the obstruction map is shown in Figure 2a, where the red areas represent typical obstructions such as buildings, trees, and other structures that can block the view of the sky. The white areas represent the obstruction-free sky from the UT's perspective. We call the `dishGetObstructionMap` method exposed by the gRPC interface every second and the obstruction maps of adjacent seconds are `XOR`ed to obtain the difference, which corresponds to the serving satellite's position at each second. To synchronize data collection with Starlink's scheduled handover times, we reset the obstruction map by calling the `dishClearObstructionMap` method[4] at 12-27-42-57 seconds of every minute. This approach prevents the

[2]Available: https://github.com/aliahan/SatInView
[3]https://github.com/fullstorydev/grpcurl
[4]This gRPC method was first available through firmware updates in early 2024, eliminating the need to reboot the dish to clear the obstruction map as required in previous efforts, based on our feedback and request.

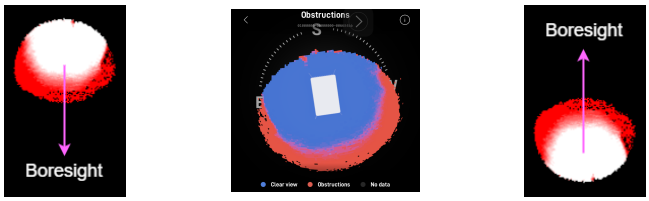

Figure 2: Starlink obstruction map obtained from (a) the gRPC interface and (b) the mobile app, and (c) the processed trajectory plot

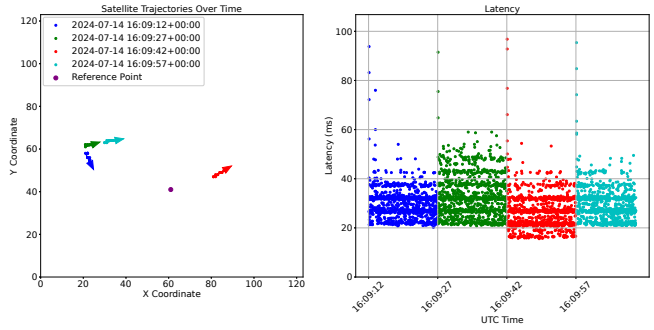

Figure 3: Serving satellites trajectory and latency within 15-second intervals

overlapping of satellite trajectories and ensures that we only obtain the trajectory of the serving satellite, tracking its precise position and direction over time. We also measure the latency from the UT to the gateway (`100.64.0.1` for standard Starlink subscribers over IPv4 [5]) for the entire duration of the trajectory data collection. Note that by measuring the latency to the gateway for the satellite link, instead of a public Internet address (such as `8.8.8.8`), we minimize the latency fluctuation of terrestrial links, also enabling our approach to be adaptable for *inactive* Starlink UTs (without active Starlink subscriptions). Figure 3 demonstrates the trajectory and the direction of serving satellites and the corresponding UT to gateway latency in one minute. It reveals the relationship between satellite trajectories and latency fluctuations, showing that latency spikes are caused by handover events between serving satellites.

When analyzing the serving satellite's trajectory with the gRPC obstruction map, we observed that the map is oriented such that the bottom-center pixel corresponds to the UT's boresight direction. In Figure 2, we purposely obstructed our UT to the north, towards its boresight direction. As illustrated in Figure 2a and Figure 2b, the obstructed region in the raw obstruction map appears at the bottom, toward

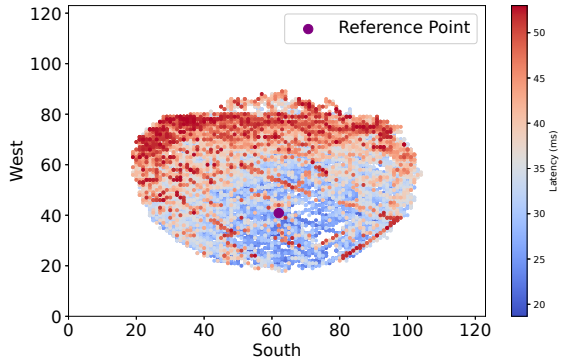

Figure 4: Latency distribution in trajectory plot

the UT's boresight direction to the north, validating our observation. To ensure consistency across geographical locations with different alignment parameters and latitudes, we vertically invert the gRPC obstruction map to produce a trajectory plot as shown in Figure 2c, ensuring the top-center of the image aligns with true north, after converting the observed pixels to topocentric coordinates.

As indicated in Starlink installation guidelines[5], the UT must have an unobstructed view above a 20° elevation for optimal performance. Consequently, the radius of the circle fitting satellite trajectories in the obstruction map, which is 62 pixels, must represent at least a 70° elevation. This ensures satellites at low elevations are accurately represented in the obstruction map when the UT is not tilted. We extensively cross-checked the observed trajectories with TLE data to determine the conversion factor (CF) between the pixel count and the elevation in degrees. This process involves overlapping the observed satellite trajectories from gRPC interface and the satellite trajectories calculated from TLE data across different 15-second intervals. Our analysis obtained an empirical CF of $(\frac{80}{62})°$, where the radius of the trajectory plot represents 80° of elevation.

We define reference point (RP) as the UT's observing position in the trajectory plot and θ as UT's tilt angle in degrees towards the boresight direction. The position of RP, influenced by θ is located at $[62, 62 - (\frac{\theta}{\text{CF}})]$ in the 123 by 123 pixel trajectory plot. All topocentric coordinates are calculated relative to this point. It is important to note that the RP may differ for UTs with different orientations and tilts.

Figure 4 shows a combined trajectory plot with UT to gateway latency distribution. We filter the latency outliers using the Interquartile Range (IQR) method to mitigate the influence of latency spikes during handover periods. The coordinates represent observed satellite positions while the

[5]https://api.starlink.com/public-files/installation_guide_standard_kit.pdf

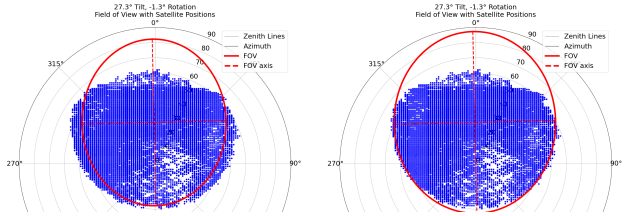

Figure 5: 110° specification FOV (left) and 120° functional FOV (right)

color gradient indicates the variation of UT to gateway latency. The dense red scatters far from the RP correspond to higher UT to gateway latency, indicating a potential correlation between the satellites' relative positions to the UT and latency. The further away a pixel is from the RP, the corresponding satellite likely has a lower elevation and a greater distance from the UT, potentially contributing to the higher latency. As previously shown in Figure 3, four consecutive satellite trajectories and their corresponding UT to gateway latency are illustrated. The red trajectory, which is relatively closer to the RP, shows a noticeably lower latency. We also observed that the 2nd and 4th satellites are actually the same satellite (STARLINK-4071), suggesting that Starlink's central scheduling controller still performs some unnecessary handovers and switching between serving satellites. This is the root cause of latency fluctuations between different 15-second intervals.

3.2 Finding Available Satellites with FOV

To determine the available satellites, we calculate the UT's FOV, which is an elliptical area expressed in degrees with a base extent (BE), representing the extent of the UT's view. A UT with no tilt has a circular FOV with a radius of BE, centered at 90° elevation. This is referred to as the initial view. To account for the UT's tilt θ, we shift the center of the ellipse by an angle equal to the tilt toward the boresight. We then compute the radii of the ellipse. The major radius remains equal to the BE, extending from the center toward and opposite the boresight direction. The minor radius is calculated using the Pythagorean theorem as $\sqrt{(BE)^2 - \theta^2}$. This results in the minor axis meeting the initial view at its widest points. The ellipse with this major and minor axis is our FOV.

We refine the base radius by tuning it with the trajectories observed from the satellites and adjust the FOV to fit the observed satellite topocentric coordinates. Topocentric coordinates of all the observed satellites are calculated by first determining the relative positions from the RP to each detected trajectory pixel. These differential coordinates are used to compute the azimuth angle using the arctangent

function. We then add the boresight azimuth angle of the UT's orientation to convert it to polar azimuth. Elevation is derived by calculating the radial distance of the trajectory pixel in pixel count and multiplying it by the CF. We utilize the topocentric coordinates of all the observed trajectories to create the polar plot illustrated in Figure 5. The center of the polar plot represents the observer, and the top center corresponds to the true north.

The Starlink UTs used in our measurement are of the hardware revision model `rev3_proto2`. The FOV of the UT listed in the specifications[6] is 110°, while our observations indicate that a 120° FOV better fits the satellite trajectories in the polar plot as shown in Figure 5. We define the wider FOV as functional FOV and verify its existence across different dishes worldwide. One possible interpretation is that LEO satellite networks often list the minimum range in their specifications to provide conservative estimates for performance and reliability, while the actual functional range may exceed the stated values.

3.3 Determining the Serving Satellite

To determine the serving satellite, we calculate the trajectories of all available satellites for each 15-second handover interval. This involves converting the orbital parameters of each satellite into topocentric coordinates relative to the RP using SGP4 propagation [7] and TLE data. The observed trajectories from obstruction data (*obs*) are then compared with the calculated satellite trajectories (*sat*) to identify the closest match. We introduce a proximity measure (D) that sums the difference in elevation and azimuth, as shown in Equation 1, at the start, middle, and end of the 15-second interval of an *obs* trajectory.

$$D = \sum_{t \in \{\text{start, middle, end}\}} \left(\frac{|elev_t^{obs} - elev_t^{sat}|}{90} + \frac{\min\left(|az_t^{obs} - az_t^{sat}|, 360 - |az_t^{obs} - az_t^{sat}|\right)}{180} \right) \quad (1)$$

where $elev$ denotes the elevation, az represents the azimuth and t represents the time of observation. Then, the directional vector difference, which quantifies how much their overall movement in elevation and azimuth differ, is added to compute the total difference. The satellite with the smallest total difference is selected as the serving satellite.

4 RESULTS AND ANALYSIS

To verify the accuracy of the proposed method, we conducted tests using UTs across geographically diverse locations, including Japan, Belgium, Canada, and Australia, each with varying alignment parameters.

[6]https://www.starlink.com/specifications

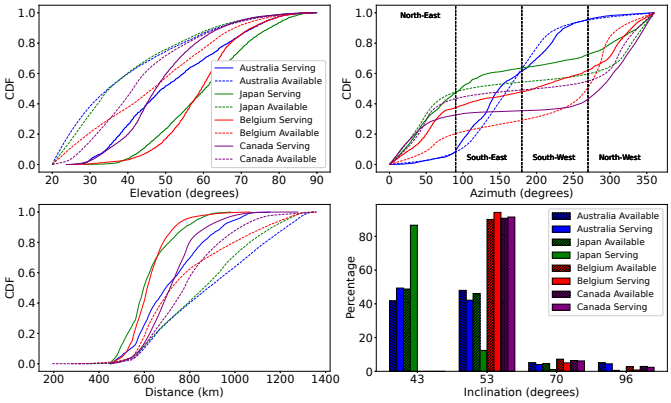

Figure 6: Evaluation and validation on different UTs

Our method accurately overlaps the serving satellite trajectory from obstruction data with the satellite's trajectory from the TLE data set in more than 99% of the 1000 visually validated gRPC and TLE trajectory sets. We observe that both elevation and azimuth deviations are predominantly low, suggesting a high level of precision in this overlapping. The mean deviations are, measuring 0.53° for elevation and 1.14° for azimuth. We observe that even if the deviation is more than expected due to our trajectory-based approach, the non-serving satellites will be discarded when compared over the 15-second interval and in their trajectory direction. However, due to geometric limitations in projecting 3D satellite positions into a 2D obstruction map, we observe increased deviations between the serving and observed satellite positions when the elevation is close to 90° or 25°, which, in a few cases, leads to uncertainty in identifying the serving satellite.

Serving vs Available: Figure 6 shows the distribution of satellite characteristics across different geographical locations and compares the serving and available satellites in their elevation, azimuth, distance, and orbital inclination. Serving satellites generally have higher elevation angles compared to available satellites, indicating a preference for satellites at higher elevations. We observe that the distances to the serving satellites are shorter, suggesting that the Starlink satellite selection strategy prioritizes satellites at higher elevations and closer distances to optimize connectivity and performance. This is especially noticeable in Japan and Belgium.

The azimuth angles exhibit significant regional variation, reflecting geographic and UT's orientational differences in satellite visibility. The dish observes more available satellites toward its boresight, and the serving satellites follow this availability trend. The Starlink satellite selection strategy favors satellites that can remain longer within the UT's FOV, such as those in the northwest for dishes tilted towards

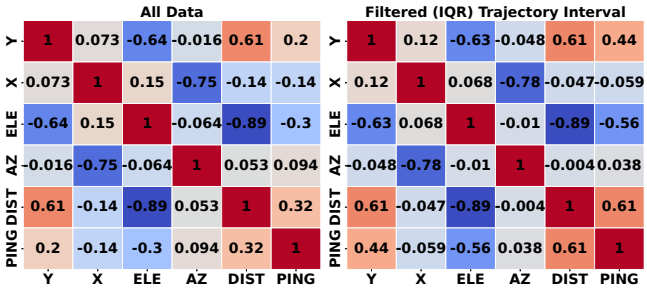

Figure 7: Unfiltered (left) and filtered (right) data correlation matrix

the north. This behavior is more evident in areas with low satellite density, such as near polar regions, where there are fewer switches to different satellites. Additionally, the satellite selection strategy follows the density of satellites' orbital inclinations, preferring 43° or 53° inclined satellites in most geographical locations based on their local availability.

Correlation Analysis: We calculate the correlation matrix for our data with and without filtering the outlier latency samples. The correlation analysis reveals a strong negative correlation between elevation and distance, indicating that as elevation increases, the distance decreases. There is also a moderate positive correlation between distance and latency. When filtering out outlier latency samples, potentially caused by handovers and the GS to PoP terrestrial network, the correlation between distance and latency nearly doubled. This indicates the considerable influence of the serving satellite on latency when these factors are excluded.

5 CONCLUSION AND FUTURE WORK

In conclusion, this paper presents a comprehensive analysis of Starlink obstruction maps and UT's FOV. We develop an accurate serving satellite identification methodology with satellite trajectories. We are also the first to compare the identified serving satellite with the actual available satellites in the FOV by incorporating the UT's orientations to calculate its FOV. Our approach has been validated across various geographic locations and latitudes, demonstrating the robustness and generalizability of this approach. The insights gained from our analysis advance the understanding of the Starlink network architecture and its performance, highlighting the correlation between serving satellites and UT to gateway latency. Future work could identify the landing GS serving the UT and investigate the impact of ground tunnels between different landing GS to user's home PoP on UT to gateway latency and overall network performance and user experience.

ACKNOWLEDGMENT

This research is supported in part by NSERC, CFI, and BCKDF. This work also would not be possible without the people such as Shun, Rory, and Francois who generously allowed us to access their dishes remotely.

REFERENCES

[1] [n. d.]. CelesTrak. https://celestrak.org/.
[2] [n. d.]. Starlink. https://en.wikipedia.org/wiki/Starlink.
[3] Manuel Bulo, Jorg Deutschmann, Kai-Steffen Hielscher, and Reinhard German. 2024. Crowdsourced Starlink Performance Measurements from https://starlinkstatus.space. In *2024 IEEE 59th International Conference on Communications (ICC '24)*. https://doi.org/10.1109/ICC51166.2024.10622394
[4] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. 2024. Democratizing LEO Satellite Network Measurement. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 8, 1, Article 13 (Feb 2024), 26 pages. https://doi.org/10.1145/3639039
[5] Jianping Pan, Jinwei Zhao, and Lin Cai. 2023. Measuring a Low-Earth-Orbit Satellite Network. In *2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC '23)*. 1–6. https://doi.org/10.1109/PIMRC56721.2023.10294034
[6] Jianping Pan, Jinwei Zhao, and Lin Cai. 2024. Measuring the Satellite Links of a LEO Network. In *2024 IEEE 59th International Conference on Communications (ICC '24)*. https://doi.org/10.1109/ICC51166.2024.10623111
[7] Brandon Rhodes. 2019. Skyfield: Generate High Precision Research-Grade Positions for Stars, Planets, Moons, and Earth satellites. https://ui.adsabs.harvard.edu/abs/2019ascl.soft07024R. Astrophysics Source Code Library (2019), ascl-1907.024.
[8] Aryan Taneja, Debopam Bhattacherjee, Saikat Guha, and Venkata N. Padmanabhan. 2023. On Viewing SpaceX Starlink through the Social Media Lens. https://doi.org/10.48550/arXiv.2307.13441 arXiv:2307.13441 [cs]
[9] Hammas Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. 2023. 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 '23)*. Association for Computing Machinery, New York, NY, USA, 37–43. https://doi.org/10.1145/3624354.3630586
[10] Shubham Tiwari, Saksham Bhushan, Aryan Taneja, Mohamed Kassem, Cheng Luo, Cong Zhou, Zhiyuan He, Aravindh Raman, Nishanth Sastry, Lili Qiu, and Debopam Bhattacherjee. 2023. T3P: Demystifying Low-Earth Orbit Satellite Broadband. https://doi.org/10.48550/arXiv.2310.11835 arXiv:2310.11835 [cs]
[11] Jinwei Zhao and Jianping Pan. 2024. LENS: A LEO Satellite Network Measurement Dataset. In *Proceedings of the 15th ACM Multimedia Systems Conference (MMSys '24)*. Association for Computing Machinery, New York, NY, USA, 278–284. https://doi.org/10.1145/3625468.3652170
[12] Jinwei Zhao and Jianping Pan. 2024. Low-Latency Live Video Streaming over a Low-Earth-Orbit Satellite Network with DASH. In *Proceedings of the 15th ACM Multimedia Systems Conference (MMSys '24)*. Association for Computing Machinery, New York, NY, USA, 109–120. https://doi.org/10.1145/3625468.3647616