



# Making Sense of Constellations

## Methodologies for Understanding Starlink's Scheduling Algorithms

Hammas Bin Tanveer  
University of Iowa

Mike Puchol  
Google X

Rachee Singh  
Cornell University

Antonio Bianchi  
Purdue University

Rishab Nithyanand  
University of Iowa

### ABSTRACT

Starlink constellations are currently the largest LEO WAN and have seen considerable interest from the research community. In this paper, we use high-frequency and high-fidelity measurements to uncover evidence of hierarchical traffic controllers in Starlink – a global controller which allocates satellites to terminals and an on-satellite controller that schedules transmission of user flows. We then devise a novel approach for identifying how satellites are allocated to user terminals. Using data gathered with this approach, we measure the characteristics of the global controller and identify the factors that influence the allocation of satellites to terminals. Finally, we use this data to build a model that approximates Starlink's global scheduler. Our model is able to predict the characteristics of the satellite allocated to a terminal at a specific location and time with reasonably high accuracy and at a rate significantly higher than baseline.

### CCS CONCEPTS

• Networks → Network dynamics.

### KEYWORDS

Satellite network; Starlink; Traffic engineering.

### ACM Reference Format:

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 Companion '23), December 5–8, 2023, Paris, France*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3624354.3630586>

### 1 INTRODUCTION

Low-earth orbit (LEO) satellite networks are expected to play an important role in achieving global broadband-like Internet connectivity since they enable low-latency, last-mile connectivity without heavy infrastructure costs (e.g., cell towers and fiber deployments). Unfortunately, prior work has shown lackluster network performance from Starlink-connected end-hosts, both via measurements [22] as well as simulations [7, 18]. Despite these findings, researchers have been unable to propose and validate methods for

improving the performance of the Starlink network. This is, in large part, due to the opacity of the network and its scheduling algorithms. Knowledge of algorithms responsible for determining which satellites route traffic from specific user terminal locations is key to engineering performance improvements for the network. In this paper, we address this gap in knowledge. In particular, we empirically uncover the scheduling algorithms used by the Starlink network by analyzing data from high-frequency measurements from four Starlink terminals (deployed both in the US and the EU) to servers co-located at their corresponding Point-of-Presence.

From our longitudinal and high-frequency measurements, we find that Starlink routes traffic from user terminals to their ground stations in a two-step process. First, a *global network controller* allocates a satellite to each user terminal based on a variety of factors including load, geospatial conditions, satellite charge, etc. Our experiments show that these allocations are made every 15-seconds, globally. Second, a *local on-satellite controller* schedules flows from the user terminals assigned to it. Taken together, these findings suggest that the hierarchical traffic engineering mechanisms, commonly deployed in terrestrial WANs [14], are also deployed by Starlink. We were able to confirm the validity of these major findings using recent FCC filings by SpaceX [5].

This work is the first to uncover hierarchical controllers and the characteristics of traffic engineering in Starlink, due to the following three key reasons. First, our high-frequency (millisecond granularity) measurements allow us to observe signatures of traffic engineering (e.g., abrupt latency changes) that cannot be observed with the coarse-grained measurements of prior work [22]. Second, by co-locating our destination server at the Starlink PoP, we minimize the influence of terrestrial latency on our measurements. Finally, our novel methodology for identifying the satellite currently serving a given terminal allows us to obtain ground-truth regarding the traffic engineering decisions made by Starlink.

**Contributions.** This paper makes four main contributions.

- We use high-frequency measurements to show evidence of hierarchical traffic engineering on Starlink (§3).
- We develop a novel technique for identifying the satellite that serves a user terminal (§4).
- We uncover the characteristics and preferences of Starlink's global scheduler – *i.e.*, the algorithm responsible for allocating satellites to specific user terminals (§5).
- We develop an approximation of the Starlink global scheduler that can predict the satellites allocated to user terminals with reasonably high accuracy (§6).



This work is licensed under a Creative Commons Attribution International 4.0 License.

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

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0407-9/23/12.

<https://doi.org/10.1145/3624354.3630586>

## 2 BACKGROUND

Starlink is a low-earth orbit satellite constellation consisting of nearly 4000 satellites in the low Earth orbit (LEO). The Starlink ecosystem has four key components: (1) In-orbit satellites, (2) User terminals or dishes (3) Ground stations and (4) Points of Presence (PoPs). Figure 1 shows how these components interact with each other to provide Internet connectivity to Starlink's end users.

**User terminals.** Terminals are deployed on user premises to connect to in-orbit satellites. Starlink user terminals are sophisticated phased-array antennas equipped with a motor that can physically reposition the angle of the dish to track fast moving satellites in the sky. User terminals can connect to any satellite at an angle of elevation (AOE) higher than  $25^\circ$  (see Figure 1). While tens of satellites satisfy the angle of elevation constraints, a terminal can connect to only one satellite at a time. Terminals forward user traffic to the satellite assigned to them. Internals of the algorithm that maps user terminals to satellites are currently known only to the operators of the Starlink network. In this work, we empirically demonstrate characteristics of this algorithm and build its approximation.

**Satellites.** A Starlink satellite connects to multiple user terminals at a time. They allocate radio frames to user terminals mapped to them to exchange data. In our work, we find evidence that this allocation is determined by a local on-satellite controller. In fact, we also find the description of this controller, referred to as the medium access control scheduler [15], in recent FCC filings from SpaceX. This controller considers factors such as user priority, current load, and per-terminal flow characteristics when forwarding the traffic from user terminals to ground stations.

**Ground stations and PoPs.** Ground stations consist of a set of phased-array antennas that receive traffic from satellites and send it through wired links to Starlink's PoPs. Like user terminals, ground stations can communicate with satellites at an angle of elevation higher than  $25^\circ$  above the horizon, relative to the ground station. A PoP is a terrestrial server with wired connectivity to a ground station. PoPs are connected to the Internet backbone. From the PoP, the traffic is routed to the destination on the Internet.

## 3 EVIDENCE OF TRAFFIC ENGINEERING

**Experiment setup: Vantage points.** We perform our measurement using four Starlink terminals – one each in Western Europe, Northeast US, Midwest US, and Northwest US. To improve the precision of our measurements, we configured the Starlink router to operate in bridge mode and connected them to a dedicated Raspberry Pi via Ethernet. These Pi's were the source of our measurements. This approach prevents the complexities which arise from using wireless routers in the measurement infrastructure. The destination of our measurements were servers co-located at the Starlink PoPs assigned to the regions of our user terminals. This choice of destination allows our measurements to be relatively unimpacted by the vagaries of terrestrial networking.

**Experiment setup: Measurements.** We conduct high-frequency measurements of the round-trip times and packet loss rates between our sources and destination servers. Packets were sent using iRTT [3] at the rate of 1 packet/20 ms and iPerf3 at a bandwidth of 50% of the upstream connection. These parameters were chosen because

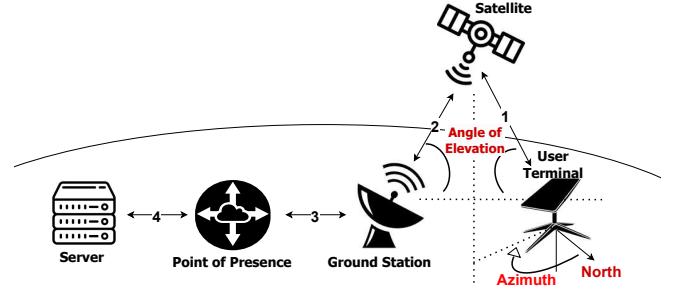


Figure 1: Starlink ecosystem.

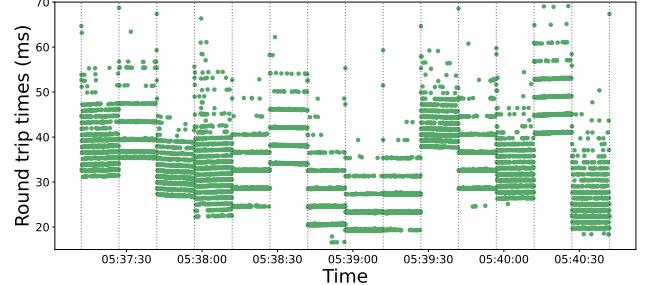


Figure 2: Measured RTT from EU terminal.

they allowed stable and reliable measurements of the Starlink network. At higher frequencies and bandwidths, the packet loss rates and measured round-trip times were highly variable even within the same measurement period. From these measurements, we recorded high-resolution round-trip times and packet loss rates. To facilitate accurate measurements of the round-trip times, the clocks of our vantage points and servers were routinely synchronized using NTP.

**Observation: Starlink relies on a global controller for satellite-to-terminal scheduling.** Figure 2 shows the changes in measured latency during a brief measurement window of two minutes for our EU terminal. It is immediately obvious that major changes in latency characteristics occur every 15 seconds – specifically, at the 12th, 27th, 42nd, and 57th second past every minute. Notably, these changes are observed from all our measured locations for all periods of time. In addition to visual tests, we are also able to confirm that the latency characteristics observed during these consecutive 15-second windows are statistically different (Mann-Whitney U test;  $p < .05$ ) from each other for all locations and over the entire period of our measurements. These drastic changes in latency are suggestive of global changes in the satellites allocated to user terminals for several reasons, including: (1) our measurements effectively nullify the impact of terrestrial networks; (2) these effects were observed, simultaneously, from all our vantage points; and (3) these effects were noticed even when our terminals were running well under capacity. Upon further investigation, we discovered an FCC filing from SpaceX which describes a global scheduler for periodically allocating terminals to satellites [5]. We conclude that our measurements have uncovered evidence of this scheduler and the periodicity of reallocations. It is important to note that this finding renders impossible the hypothesis that performance characteristics are associated with the movement of the satellite assigned to serve a terminal (e.g., Figure 7 in [17]). This is because changes in satellite allocation occur every 15 seconds which is insufficient time to

meaningfully cause impacts on performance due to change in satellite positions/distances. Next, we will uncover the characteristics of this scheduler and develop an offline approximation for it.

**Observation: Starlink uses an on-satellite controller for scheduling terminal flows.** The second peculiar characteristic of the latency measurements from our user terminals is that within the 15-second time interval, latency measurements the user terminal frequently form parallel bands that are a few milliseconds apart. These bands reflect evidence that radio frames are allocated to user terminals by an on-satellite controller in a round-robin fashion. Further investigation is required to exactly identify the characteristics of this controller, which we believe to be the on-satellite Medium Access Control scheduler described in a SpaceX FCC filing [15].

## 4 OBTAINING SATELLITE ALLOCATIONS

Our results show that the Starlink network uses a global scheduler to assign satellites to user terminals every 15 seconds (Section 2). Unfortunately, the Starlink mobile app no longer identifies the satellite that a user terminal is connected to. Not having this knowledge limits our ability to reverse-engineer the mechanics of Starlink's scheduling algorithms. In this section, we develop a novel technique that leverages Starlink's *obstruction maps* to identify the satellite allocated to a specific user terminal. At a high-level, our approach involves correlating the publicly known positions of the Starlink satellites with observations of connected satellites recorded by in the obstruction maps of each terminal.

**Data: Obstruction maps.** Obstruction maps are  $123\text{px} \times 123\text{px}$ , 2-dimensional images which mark the trajectory of satellites that recently served the user terminal. These images are used to create a 3-dimensional map made available to users via the Starlink mobile app. The 3-d map is meant to help users identify the quality of the location of their terminal, highlighting any physical obstructions between their terminal and any satellites meant to serve the terminal. Figure 3a shows an example of this map from the Starlink app. We used `starlink-grpc-tools` [4] to extract the 2-d obstruction maps every 15 seconds from each terminal. The 3-d maps are not possible to obtain programmatically. Figures 3b and 3c shows examples of these maps from two consecutive 15-second time slots.

**Data: Satellite positions.** Positions of Starlink satellites are available publicly from a variety of sources in a two-line element(TLE) format. We use CelesTrak [2] to get the TLEs for Starlink satellites. Since these files only indicate satellite positions every six hours, we use the SGP4 satellite propagation algorithm [26] to calculate satellite positions, relative to terminal, for a specific point in time.

### 4.1 Methods

We now describe our methods for identifying the satellite currently serving a user by utilizing obstruction maps and TLE files.

**Method: Uncovering gRPC obstruction map parameters.** As seen in Figures 3b and 3c, the 2-d obstruction maps are plain – only containing white pixels which indicate the trajectory of recently connected satellites without any further context (e.g., the angle of elevation and azimuth). To establish context for the 2-d image, we make use of the property that the obstruction maps are over-written with newer satellite trajectories until the terminal is rebooted. By

leaving the terminal online consecutively for a 2-day period, we allowed the terminal to connect to satellites from practically all the regions of the sky that are within its field of view. Once the 2-d image is completely filled-up, as shown in Figure 3e, we draw bounding boxes around these trajectories to identify the center and boundaries of the 2-d image. Using this method helped us find that: (1) the 2-d obstruction map is a polar plot centered at  $62 \times 62$ ; using polar plots to visualize satellite trajectory in terms of AOE and azimuth is common practice [11]. (2) the radius of the polar plot represents the angle of elevation and ranges from 25 to 90 (as terminals can only connect to satellites above  $25^\circ$ 's AOE). (3) the  $\theta$  of the polar plot represents the azimuth, where  $\theta = 0$  represents the North and increases in a clock-wise direction; and (4) the radius of the contained polar plot is 45 pixels. This context is crucial for identifying the connected satellites using the trajectories painted on the 2-d obstruction images.

**Method: Isolating satellite trajectory.** After recovering the parameters of the gRPC obstruction map, we use them to identify the trajectory of the satellite that is allocated to the terminal for a specific 15-second slot (denoted by  $x$ ). To obtain this trajectory, we perform an XOR operation on the obstruction map from  $x$  and  $x - 1$  (i.e., the prior 15-second slot). This will result in the erasure of all satellite trajectories which were common to the two figures – leaving visible only the trajectory associated with the satellite connected to the terminal during slot  $x$ . We note that for this method to work, we require that satellite trajectories are not overlapping with the trajectories of previously connected satellites (since an XOR would erase the overlaps). To ensure that this condition is met, we perform a terminal reset every 10 minutes (since resetting the terminal starts a fresh gRPC map).

**Method: Identifying serving satellite.** For each isolated satellite trajectory, we compute the AOE and Azimuth for each individual pixel. For AOE, we calculate the distance of each pixel from the center of the polar plot (at  $45 \times 45$ ) – as AOE is represented using the radius in polar plots. For Azimuth, we calculate the  $\arctan(x, y)$  – where  $x, y$  is the  $\theta$  of each individual pixel relative to the center – to determine the clock-wise angle between North at  $0 \times 0$  and every pixel in an isolated satellite's trajectory.

Although we have now computed the positions of the satellite we are connected to during a 15-second slot, we still do not know its identity. To identify the current serving satellite, we compare the AOE and Azimuths calculated above to those of all satellites in our terminal's field-of-view – calculated using TLE files (as described above) – for the given 15-second slot. For each available satellite's position, we compute the Dynamic Time Warping (DTW) [23] distance with the isolated trajectory positions; the available satellite with the lowest DTW distance is chosen as the current serving satellite. *Note:* To compute DTW distances, we first need to convert all positions from polar to Cartesian co-ordinates.

To validate our matching technique, we plot the trajectories of all available satellites on a polar plot and visually compare them to the isolated trajectory. We conduct a manual (visual) pilot test study of 500 sets of isolated trajectories and polar plots of available satellite trajectories; the DTW similarity method and our manual tests overlapped on over 99% of all outcomes.

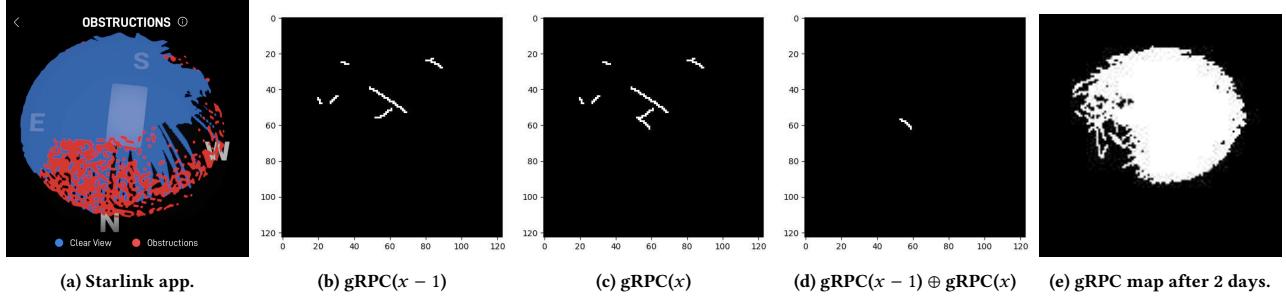


Figure 3: Obstruction maps (a) obtained from the Starlink app, (b, c) obtained from gRPC for two consecutive 15-second slots  $x - 1$  and  $x$ , (d) their XOR, and (e) the gRPC map after two days without a terminal reset.

## 5 STARLINK'S GLOBAL SCHEDULER

The Starlink network uses a global scheduler to allocate user terminals to individual satellites every 15 seconds (Section 2). However, the scheduling algorithm is not publicly known. In this section, we analyze characteristics of satellites that were allocated to our terminals with the goal of reverse engineering the algorithm of the global scheduler. On average, there are  $\sim 40$  satellites in the field of view of a user terminal during a 15 second slot. Having identified the satellite that a terminal is connected to during a slot (Section 4), we compare properties of the satellites that are selected by the scheduler with those that were available but not chosen.

### 5.1 Impact of satellite position

**Angle of elevation.** First, we compare the positions in the sky of satellites that were available but not selected to the positions of satellites that were selected by the global scheduler. The position of a satellite in the sky is defined by its AOE and azimuth with respect to the user terminal. Figure 4 shows that the median angle of elevation of selected satellites (solid lines) is  $22.9^\circ$  higher than that of the available but unselected satellites (dotted lines). Although only 30% of all available satellites had their AOE in  $45^\circ$  to  $90^\circ$  range, the global scheduler picked 80% of satellites from the range (averaged over all locations).

**Direction.** Figure 5 shows the distribution of azimuths of the two sets of satellites. The plot is divided into 4 quadrants which represent the direction – stated at the top of each quadrant – of the satellites relative to the face of the user terminal. Although the azimuths of available satellites (dotted lines) are evenly distributed throughout the 4 quadrants, azimuths of selected satellites (solid lines) are skewed towards the north of the dish, except for the dish in Ithaca, NY. Further investigation revealed that our user terminal in Ithaca picked fewer satellites from the north west as this region was severely obstructed by trees. The terminal in Ithaca was assigned only 9.7% of the satellites from the region compared to 55.4% on average by user terminals in other locations. In other locations, 58% of satellites on average were available towards the north of the user terminals, however, the user terminal was mapped to satellites from the north 82% of the times.

**Rationale:** The International Telecommunication Union has imposed a mandatory geo-stationary orbit exclusion zone, which prohibits LEO satellites from transmitting to or receiving from a ground station while being in the protected part of the sky [1]. This mandate forces terminals at latitudes more than  $40^\circ\text{N}$ , the

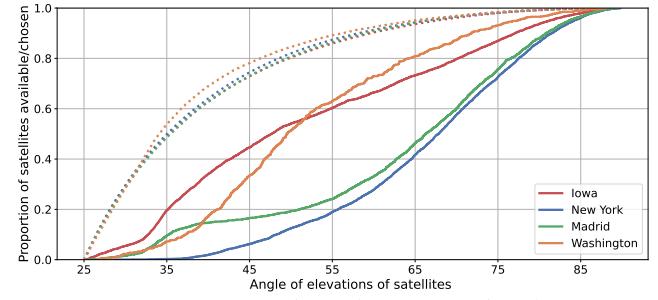


Figure 4: AOEs of available (dotted) vs. selected (solid) satellites.

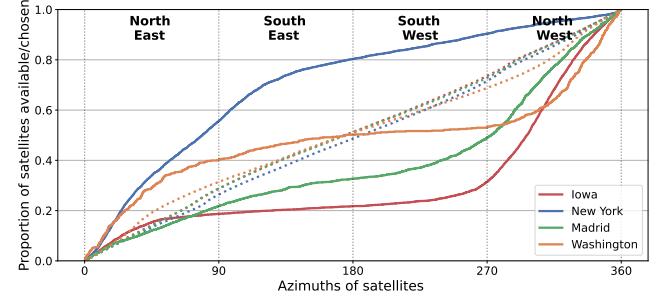
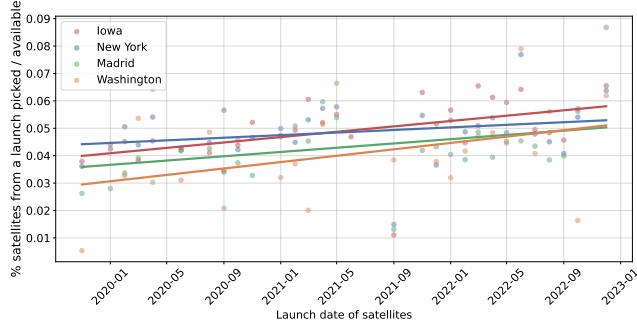


Figure 5: CDF of AZs of satellites available (dotted) vs. AZs of satellites chosen (solid). The x-axis is divided into 4 quadrants of 90 degrees each; the direction of the quadrant relative to the user terminal's face is listed on top of each quadrant.

approximate latitude of our terminals, to point much higher than required by the minimum angle of elevation constraint. This is why the global scheduler assigns satellites higher up in the fields-of-view of our terminals. Moreover, satellites with a higher AOE can communicate with terminals in a more energy efficient way. The distance between the user terminal and satellite increases inversely with AOE. As radio frequency (RF) power decreases inversely with distance, satellites farther away need to use significantly more power to communicate with user terminals.

### 5.2 Impact of satellite launch dates

The Starlink constellation consists of more than 4,000 satellites which were released in batches since 2018. We analyze whether the global scheduler prefers satellites from certain batches over others. To achieve this, we bin satellites by the year and month of their launch batch. We then compute the percentage of number of slots a satellite from a launch was picked vs. the number of slots a satellite from a launch was available – for all 15 second slots



**Figure 6: Correlation of probability of picking a satellite from a launch vs. launch date**

in our observation. Figure 6 shows a positive correlation between launch date and probability of a satellite being picked; the Pearson correlation, averaged over all locations is 0.41 (discarding the New York location because of significant obstructions). Figure 6 shows an increase in the probability of being picked of 0.02 when comparing the latest launch to the earliest (for Iowa).

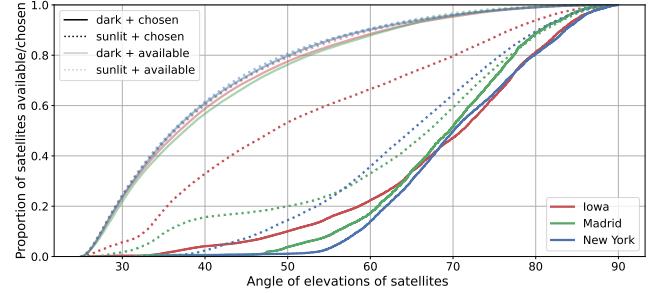
**Rationale:** As Starlink satellites are launched in batches, the difference in service time between the latest and oldest satellites can differ by years. This can lead to some parts of the constellation going out of service years before others because each Starlink satellite only has a lifetime of  $\sim 5$  years; requiring consistent replacement efforts to maintain the constellation's coverage. Hence, using satellites launched later in time can help offset the discrepancy in service times of satellites allowing the constellation to stay stable for longer. However, as the global scheduler needs to maintain performance/efficiency of the network for every 15 second slot, other decisions e.g., position of the satellite, might take precedence over longer term goals like stabilization of the constellation. This is evident in the low absolute values of the increase in probability caused by later launch dates. Regardless, a positive correlation indicates a preference for using satellites launched later.

### 5.3 Impact of being sunlit

Starlink satellites use on-board solar panels to provide power for their operations. However, as satellites orbit the earth, they periodically go in and out of sunlight. To serve user terminals during times when a satellite is not under sunlight, it has to conserve energy to stay functional until it gets sunlight again. Hence, we analyze if Starlink's global scheduler has preference for sunlit satellites.

For every 15 second slot, we calculate positions of all available satellites relative to the sun using the SkyField library. We then extract 15 second slots during which at least one sunlit and at least one dark satellite was available. During such slots, the global scheduler opts for the sunlit satellites 72.3% of the time averaged over all locations. We also find that the global scheduler only picks dark satellites during 15 second slots where the  $\%(\text{dark}/\text{available})$  satellites is  $\geq 35\%$  (averaged over all locations). Essentially, the global scheduler prefers sunlit satellites when both, sunlit and dark satellites are available in roughly equal majority.

Next, we compare the positions of the dark and sunlit satellites that were picked by the global scheduler. Figure 7 shows that although the distribution of AOE of the available, dark and sunlit satellites, is similar, the distribution of their selected counterparts is significantly different. Dark satellites are picked when they are at



**Figure 7: AOE of dark available, dark selected, sunlit available and sunlit selected satellites.**

a higher AOE compared to picked sunlit satellites; 82% of all dark satellites picked had an AOE above  $60^\circ$  compared to 54% of sunlit satellites. Averaged over all locations, chosen dark satellites had AOE  $29^\circ$  higher than their sunlit counterparts.

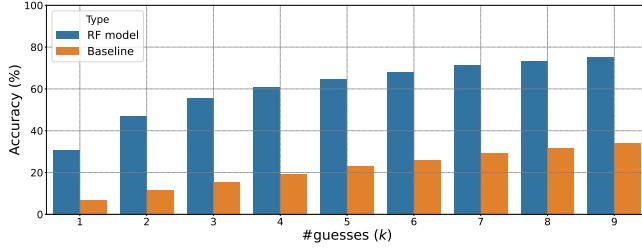
**Rationale:** The RF power decreases inversely with distance causing satellites to expend more energy in communicating with the user terminal. As dark satellites have limited battery to stay functional, the global scheduler assigns satellites to user terminals that are higher up in the field-of-view to conserve energy.

## 6 MODELING THE GLOBAL SCHEDULER

Our analysis has shown that the Starlink global scheduler has specific preferences *i.e.*, it selects newer satellites that are sunlit, located towards the North West, and at a high angle of elevation. We now use these preferences and our data to build an offline model of the scheduler. The goal of this model is to *predict* characteristics of the satellite allocated to serve a terminal in a given location and time.

**Feature selection.** Given a specific location and time, we first identify the satellites that are available to serve the user terminal using our TLE dataset and the SGP4 algorithm. Next, we cluster the available satellites based on their azimuth ( $\theta$ ), angle of elevation ( $\phi$ ), age ( $a$ ), and sunlit status ( $\epsilon$ ) as follows: Given a set of satellites ( $S$ ) available at time  $t$  for location  $l$ , the satellite  $s \in S$  with parameters  $(\theta_s, \phi_s, a_s, \epsilon_s)$  is placed in the cluster:  $(\frac{\theta_s - \mu(\theta)}{\sigma(\theta)}, \frac{\phi_s - \mu(\phi)}{\sigma(\phi)}, \frac{a_s - \mu(a)}{\sigma(a)}, \epsilon)$ . Here  $\mu(x)$  and  $\sigma(x)$  denote the mean and standard deviation of the feature  $x$  (computed from  $S$ ). Effectively, this approach clusters satellites by how many standard deviations away from the group mean each of their features are. For example, a satellite in the cluster  $(1, 0, 2, 1)$  is further than 1 standard deviation away from the mean azimuth and 2 standard deviations away from the mean age of all the satellites in  $S$ . As features for our scheduler, we present the local time ( $t_l$ ) and the count of satellites available in each cluster at the start of the nearest 15-second interval.

**Training, testing, and validating our model.** Our goal is to construct a model which takes the above features as input and returns the cluster to which the allocated satellite belongs. We train a random forest model because of its robustness to over-fitting and the explainability of its predictions. We got the parameters of this model using grid-search and five-fold cross-validation. We use the data gathered from each location and measured satellite allocations (Section 4) to construct our model – 80% of the data is used to create a training/testing data-set for a five-fold cross-validation evaluation and the remaining 20% is used to create a holdout data-set to validate our model's robustness to over-fitting.



**Figure 8: Accuracy of our model compared to baseline using the top- $k$  accuracy metric.**

**Evaluating the model.** To measure the accuracy of our model, we use a top- $k$ -accuracy metric – *i.e.*, we analyze the accuracy of the  $k$  most likely predictions made by the model. To put this accuracy in context, we compare it to the baseline model which simply returns the (top- $k$ ) cluster(s) with the most number of available satellites as its prediction. Our results obtained over our holdout set are shown in Figure 8. The proposed model significantly outperforms the baseline model and predicts the correct allocated satellite characteristics 65% of the time ( $k=5$  guesses), in comparison to the baseline of 22%.

**Feature importance.** We use the the gini importance score to understand which features contribute the most to the learning of our model. *local hour* stands out in all locations with a gini importance score of 0.04; this can be explained by the change in global scheduler’s decision making process during night time when it picks dark satellites (which have a higher AOE than their sunlit counterparts). The feature tuples  $(-1, -1, -1, 1)$  and  $(1, -1, -1, 1)$  consistently show up as important features for all locations, indicating that the model is sensitive to azimuths within 1 standard deviation, on either side of the mean. This can be explained by the global scheduler’s preference to not pick satellites towards the south as explained in Section 5. As azimuths (of available satellites) are distributed evenly from 0 to 360°s, the range of azimuths indicated by these features cover azimuths between 90°s and 180°s (south).

We also observe important features with the pattern  $(x, y, -1, 1)$  indicating that the model is sensitive on clusters with sunlit satellites having age one standard deviation less than the mean implying scheduler’s preference for newer sunlit satellites. We also observe important features of the form  $(x, 2, y, z)$ , indicating that clusters with AOE 2 standard deviations above the mean (satellites higher in the sky) are favored by the model.

**Limitations.** Our model is constructed using only publicly measurable data related to Starlink’s satellites. However, based on disclosures in SpaceX’s FCC filings and patents [5, 15], we expect that other publicly-unavailable features such as satellite load characteristics will also impact the global scheduler. Therefore, the performance of our model is constrained by the unavailability of data. Despite this, our model is demonstrably robust to over-fitting and is able to make reasonably accurate and explainable predictions at a rate far higher than baseline.

**Model release.** To facilitate future simulations and evaluations of the Starlink network, our model and data is available at [this link](#).

## 7 RELATED WORK

Previous work has studied Starlink’s performance along the axes of latency and throughput with coarse-grained measurements [22, 28] highlighting some of the issues that the current Starlink network

faces in terms of performance and recommending possible fixes [21]. Performance of other kinds of constellations *e.g.*, Geo-stationary orbit, has also been analyzed [25]. Researchers have found geographic variability [17] in Starlink performance. Others have developed simulations of LEO constellations [18, 29] and uncovered Starlink’s physical network topology [24] to suggest improvements to the network [8]. Researchers have proposed to improve end host performance in Starlink using better routing protocols [12, 13], better beam-forming techniques [19], satellite network topology reconfiguration [6], clean-slate protocol design [10] and efficient satellite hand-off techniques [16, 20]. Security researchers have also found potential attack surfaces in Starlink [9, 27]. To the best of our knowledge, our work is the first which conducts fine-grained measurements to partly reverse-engineer the inner workings Starlink’s global and traffic schedulers.

## 8 LIMITATIONS AND FUTURE WORK

**Generalizability of our findings.** Our results show that the global scheduler exhibits similar behavior for all of our geographically distributed vantage points. Although our findings are generalizable to a wide range of longitudes, we understand that the global scheduler can be forced to make different decisions in other latitudes *e.g.*, in the southern hemisphere, because of a change in the GSO exclusion zone [1]. To counter this limitation, we plan on including vantage points from a wide variety of latitudes in our future work.

**Impact of performance related features.** The purpose of this work was to understand if publicly available data could be used to analyze the global scheduler’s decision making process. Although we were able to get valuable insights, we understand that performance related features *e.g.*, traffic load on satellites, can possibly play a significant role in the decision making process. However, measuring this relationship would require a more elaborate experimental setup, which is part of our future work.

## 9 CONCLUSIONS

In this paper, we leveraged high-frequency and high-fidelity measurements from four globally distributed vantage points to uncover the presence of a hierarchical traffic controller – a global controller for allocating satellites to user terminals and an on-satellite local controller for scheduling user flows. Using a novel technique for identifying the satellites allocated to terminals, we identified the characteristics and preferences of the global scheduler and developed an offline approximation. Taken all together, our work is a first step towards understanding Starlink’s traffic engineering decisions and supporting future efforts to model and evaluate scheduling algorithms for the Starlink network.

**Acknowledgements.** We thank the ITS and CLAS Linux groups at the University of Iowa (JJ Urich, Scott Allendorf, Brian Bacher and Dan Holstad) for their help in experimentation setup for this study and Jianping Pan (University of Victoria) for his valuable recommendations. We also thank our reviewers for their valuable feedback. This research was funded by the National Science Foundation under award #1953983. The opinions, interpretations, conclusions, and recommendations are those of the authors and are not endorsed by the funding bodies.

## REFERENCES

[1] [n. d.]. 47 CFR § 25.289 - Protection of GSO networks by NGSO systems. <https://www.law.cornell.edu/cfr/text/47/25.289>. ([n. d.]). (Accessed on 06/25/2023).

[2] [n. d.]. CelesTrak. <https://celesttrak.org/>. ([n. d.]). (Accessed on 06/25/2023).

[3] [n. d.]. IRTT. <https://iperf.fr/>. ([n. d.]). (Accessed on 06/25/2023).

[4] [n. d.]. sparky8512/starlink-grpc-tools: Random scripts and other bits for interacting with the SpaceX Starlink user terminal hardware. <https://github.com/sparky8512/starlink-grpc-tools>. ([n. d.]). (Accessed on 06/25/2023).

[5] [n. d.]. Starlink Services LLC Application for ETC Designation | PDF Host. [https://pdfhost.io/v/BnYWSSR-wq\\_Starlink\\_Services\\_LLC\\_Application\\_for\\_ETC\\_Designation](https://pdfhost.io/v/BnYWSSR-wq_Starlink_Services_LLC_Application_for_ETC_Designation). ([n. d.]). (Accessed on 06/27/2023).

[6] Debopam Bhattacharjee 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.

[7] 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.

[8] Shkelzen Cakaj. 2021. The parameters comparison of the "Starlink" LEO satellites constellation for different orbital shells. *Frontiers in Communications and Networks* 2 (2021), 643095.

[9] Giacomo Giulari, Tommaso Ciussani, Adrian Perrig, Ankit Singla, and E Zurich. 2021. ICARUS: Attacking low Earth orbit satellite networks.. In *USENIX Annual Technical Conference*. 317–331.

[10] Giacomo Giulari, Tobias Klenze, Markus Legner, David Basin, Adrian Perrig, and Ankit Singla. 2020. Internet backbones in space. *ACM SIGCOMM Computer Communication Review* 50, 1 (2020), 25–37.

[11] L Gratton, S Khanafseh, B Pervan, S Pullen, J Warburton, and William J Hughes. 2004. Experimental observations and integrity monitor applications of LAAS IMLA carrier phase measurements. In *Proceedings of the 17th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2004)*. 2259–2270.

[12] Mark Handley. 2018. Delay is not an option: Low latency routing in space. In *Proceedings of the 17th ACM Workshop on Hot Topics in Networks*. 85–91.

[13] 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.

[14] Chi-Yao Hong, Subhasree Mandal, Mohammad Al-Fares, Min Zhu, Richard Alimi, Kondapa Naidu B., Chandan Bhagat, Sourabh Jain, Jay Kaimal, Shiyu Liang, Kirill Mendelev, Steve Padgett, Faro Rabe, Saikat Ray, Malveeka Tewari, Matt Tierney, Monika Zahn, Jonathan Zolla, Joon Ong, and Amin Vahdat. 2018. B4 and after: Managing Hierarchy, Partitioning, and Asymmetry for Availability and Scale in Google's Software-Defined WAN. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18)*. Association for Computing Machinery, New York, NY, USA, 74–87. <https://doi.org/10.1145/3230543.3230545>

[15] Jayasuryan V Iyer, Khasim Shaheed Shaik Mohammad, Yashodhan Dandekar, Ramakrishna Akella, Chen Chen, Phillip E Barber, and Peter J Worters. 2022. System and method of providing a medium access control scheduler. (Dec. 27 2022). US Patent 11,540,301.

[16] Enric Juan, Mads Lauridsen, Jeroen Wigard, and Preben Mogensen. 2022. Performance evaluation of the 5G NR conditional handover in LEO-based non-terrestrial networks. In *2022 IEEE Wireless Communications and Networking Conference (WCNC)*. IEEE, 2488–2493.

[17] 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.

[18] Simon Kassing, Debopam Bhattacharjee, 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.

[19] Rajnish Kumar and Shlomi Arnon. 2022. DNN Beamforming for LEO Satellite Communication at Sub-THz Bands. *Electronics* 11, 23 (2022), 3937.

[20] Yaoqi Liu, Xiaogang Tang, Yiqing Zhou, Jinglin Shi, Manli Qian, and Shaoyang Li. 2022. Channel Reservation based Load Aware Handover for LEO Satellite Communications. In *2022 IEEE 95th Vehicular Technology Conference (VTC2022-Spring)*. IEEE, 1–5.

[21] 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.

[22] 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.

[23] Meinard Müller. 2007. Dynamic time warping. *Information retrieval for music and motion* (2007), 69–84.

[24] Jianping Pan, Jinwei Zhao, and Lin Cai. 2023. Measuring a low-earth-orbit satellite network. *arXiv preprint arXiv:2307.06863* (2023).

[25] Daniel Perdigues, Gianluca Perna, Martino Trevisan, Danilo Giordano, and Marco Melia. 2022. When satellite is all you have: watching the internet from 550 ms. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 137–150.

[26] Brandon Rhodes. 2019. Skyfield: High precision research-grade positions for planets and Earth satellites generator. *Astrophysics Source Code Library* (2019), ascl-1907.

[27] Joshua Smailes, Edd Salkield, Simon Birnbach, Martin Strohmeier, and Ivan Martinovic. 2023. Dishing out DoS: How to Disable and Secure the Starlink User Terminal. *arXiv preprint arXiv:2303.00582* (2023).

[28] C Uran, K Horvath, and H Wöllrik. 2021. *Analysis of a Starlink-based Internet connection*. Technical Report. tech. rep., Carinthia University of Applied Science, Research Group ROADMAP-5G.

[29] Aiden Valentine and George Parisis. 2021. Developing and experimenting with leo satellite constellations in omnet++. *arXiv preprint arXiv:2109.12046* (2021).