Predicting Connectivity and Link Performance in Dynamic LEO Satellite Networks

Bin Hu
University of Southern California
hubin@usc.edu

Feng Qian
University of Southern California
fengqian@usc.edu

Ahmad Hassan
University of Southern California
ahmadhas@usc.edu

Z. Morley Mao
University of Michigan
zmao@umich.edu

Zhi-Li Zhang
University of Minnesota
zhzhang@cs.umn.edu

Abstract

Although LEO systems such as Starlink provide low-latency access through high-speed, low-altitude orbits, they also experience frequent handovers and highly variable link conditions that undermine transport-layer performance. In this work, we present a multi-granularity prediction framework that learns to predict future satellite connections and link quality directly from network telemetry and satellite geometry. At the coarse scale, our model predicts upcoming satellite handovers and connection windows. At the fine scale, it predicts per-second performance within each window. Through real-system measurements on Starlink, we demonstrate that LEO link dynamics exhibit strong temporal and spatial correlations that can be learned and predicted accurately. Our results highlight prediction as a key enabler for proactive, satellite-aware optimization across transport, routing, and application layers in future LEO Internet systems.

CCS Concepts

• **Networks** → **Network performance modeling**; *Network measurement*; *Network dynamics.*

Keywords

Satellite Network, LEO, Satellite Behavior, Prediction

ACM Reference Format:
Bin Hu, Feng Qian, Ahmad Hassan, Z. Morley Mao, and Zhi-Li Zhang. 2026. Predicting Connectivity and Link Performance in Dynamic LEO Satellite Networks. In *The 27th International Workshop on Mobile Computing Systems and Applications (HotMobile '26), February 25–26, 2026, Atlanta, GA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3789514.3792058

1 Introduction

The deployment of large-scale Low Earth Orbit (LEO) satellite constellations, pioneered by SpaceX's Starlink, has introduced a fundamentally new paradigm for global Internet connectivity. Unlike terrestrial networks, LEO satellites operate at low altitudes and high velocities, offering low latency but causing frequent handovers between satellites. Their rapidly changing geometry and limited capacity lead to significant temporal variations in link quality that challenge existing network protocols and prediction models [7, 11, 15, 16, 24]. Such dynamics often result in unreliable throughput and frequent performance degradation [25], where TCP performance on Starlink can drop to around 25 Mbps while UDP remains near 120 Mbps [7]. This instability directly affects latency-sensitive and bandwidth-hungry applications such as video conferencing and adaptive video streaming [6, 9, 14, 29]. Although recent work has sought to mitigate handover-induced disruptions [1], packet loss in LEO systems is not confined to switching periods. These dynamics highlight the need not only for measurement but also for prediction, enabling systems to predict satellite transitions and link degradation.

Recent advances in LEO networking have made impressive progress in network optimization under dynamic satellite conditions, as demonstrated by SaTCP [1], SatPipe [28], and LeoCC [12]. These advancements show that incorporating satellite awareness into congestion control and link adaptation can significantly improve throughput and stability. However, their designs still inherit the logic of terrestrial networks, reacting to measured impairments rather than predicting them. The underlying models primarily rely on packet-level feedback or link-layer indicators, while the influence of satellite-specific factors such as orbital geometry, elevation, or handover periodicity remains underutilized. A few recent efforts have explored performance prediction using statistical or learning-based methods [13, 21, 27], yet they often ignore the direct coupling between satellite motion and network behavior, limiting prediction accuracy and generalization. To address these gaps, we propose a two-stage prediction framework that jointly models satellite connectivity and performance evolution. Our approach leverages precise satellite information obtained through gRPC [22] and TLE [3] alignment, and integrates them into a joint learning model that predicts throughput and latency conditioned on predicted connections. By explicitly learning the interaction between orbital dynamics and link performance, our framework provides a predictive foundation for future satellite-aware optimization systems.

In this work, we focus on prediction as the foundation for proactive LEO networking. We present a hierarchical prediction framework that jointly models satellite connectivity and network performance across multiple time scales. Our approach moves beyond simple satellite awareness: it learns to predict which satellite a user terminal will connect to next, and how throughput and latency will evolve within and across connection periods. By accurately predicting both coarse-grained and fine-grained dynamics, our framework

HotMobile '26, February 25–26, 2026, Atlanta, GA, USA

ACM ISBN 979-8-4007-2471-8/2026/02
https://doi.org/10.1145/3789514.3792058

provides a crucial building block for future satellite-aware optimization, from congestion control to routing and application-layer adaptation.

We make the following contributions:

• **Exploratory analysis of satellite network dynamics.** Through preliminary experiments on real Starlink systems, we observe the impact of satellite switching events and reveal strong correlations between orbital dynamics and end-to-end performance. These findings motivate our prediction design and underscore the tight coupling between satellite motion and link variability.

• **Multi-granularity performance prediction.** We design a hierarchical prediction framework that captures both long-term connection transitions and short-term performance fluctuations. At the coarse granularity, we predicts the next connection and temporal-level link quality; at the fine granularity, we predicts per-second performance variations, achieving a throughput MAE of 8.37 Mbps.

• **Implementation and real-system evaluation.** We implement the proposed framework on operational Starlink terminals and evaluate its prediction accuracy using real trajectories and traffic traces. Our model achieves up to 20% improvement over state-of-the-art satellite networking optimization solutions. These results demonstrate the feasibility of low-latency, high-fidelity prediction in live LEO environments and highlight its potential as a foundation for prediction-driven optimization in satellite Internet systems.

2 Background and Motivation

2.1 Limitations of Existing Prediction

Existing models generally ignore the deterministic influence of satellite geometry and connection-level discontinuities. As a result, they cannot predict abrupt link transitions or capture fine-grained performance variations necessary for latency and throughput sensitive applications.

Our preliminary experiments illustrate this limitation concretely. A baseline model built solely on an LSTM architecture without incorporating any satellite information achieves a throughput mean absolute error (MAE) of 41.13 Mbps, as shown in Table 1. This error is substantial given that the measured average throughput is approximately 200 Mbps, indicating that purely temporal models fail to capture the key factors governing performance dynamics in LEO networks.

LEO satellite networks, unlike terrestrial networks, rely on fast-moving satellites that relay data over hundreds to thousands of kilometers. This dynamic topology leads to frequent handovers, as user terminals continuously switch connections among satellites passing overhead. Figure 1 presents time-series throughput measurements across multiple Starlink connection periods, where each window represents the duration of a single satellite connection. Red dashed lines denote handover events. We observe that throughput fluctuates significantly across windows, suggesting heterogeneous performance among satellites, while within a single window, performance remains relatively stable. These observations highlight the importance of accurately identifying the connected satellite, as it provides essential context for interpreting link performance variations and predicting future connectivity dynamics.

These limitations and observations motivate our two-stage prediction framework: (1) predicting the next connected satellite and (2) predicting link performance conditioned on the predicted connection. By leveraging orbital regularity and satellite-specific features, our design transforms prediction from a reactive process into a proactive mechanism that predicts network dynamics.

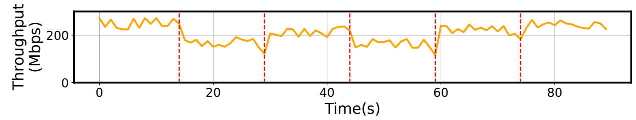

Figure 1: Performance Heterogeneity.

Model	Throughput MAE (Mbps)
Model without satellite information	41.13
Model without connection prediction	27.54

Table 1: Impact of Satellite Information on Prediction.

2.2 Identifying the Connected Satellite

To effectively exploit the relationship between satellites and network performance, it is necessary to identify which satellite the user terminal is connected to, along with the corresponding satellite properties. However, it is challenging to directly obtain satellite connection information through any publicly available interface. One potentially useful signal is the 2D image obtained from the user terminal via starlink-grpc-tools [22]. As illustrated in Figure 2, this 2D image depicts the trajectory of the currently connected satellite in polar coordinates. For example, Figure 2a shows a 15-second satellite trajectory corresponding to a single connection period.

Nevertheless, the information provided by this 2D image is limited: azimuth and elevation are encoded only as angular coordinates and radial distance, without explicit geometric or orbital context, making it insufficient for accurately inferring the satellite's 3D position or connectivity dynamics. Fortunately, global satellite position information is available through publicly accessible two-line element (TLE) data [3]. TLEs provide a compact orbital description for Earth-orbiting satellites by encoding key Keplerian elements and their time of validity, and can be propagated to obtain time-resolved 3D satellite trajectories.

By propagating TLEs to the corresponding measurement timestamps and projecting the resulting satellite positions into the local azimuth-elevation frame, we can correlate the dish-reported 2D image with candidate satellite trajectories. This correlation enables us to infer richer serving satellite information, such as satellite identity and handover dynamics, that is not directly exposed by the dish telemetry alone. However, it is difficult to accurately perform the matching using the method proposed in [23], due to ignoring dish geolocation and position, as well as gradual orbital drift over time [19, 20].

In our preliminary data collection, we conduct long-term measurements of Starlink's 2D trajectory images. These measurements reveal that the center of the observed trajectories does not coincide with the geometric center of the image, and that the distribution of visible satellite tracks is highly uneven. The satellite's visibility region depends on both the global constellation geometry and the dish's pointing direction. When the dish faces north (azimuth $\approx 0°$), the visible region is symmetric and centered; as the dish rotates eastward, the dense trajectory region shifts westward, reflecting the relative motion of satellites within the local orbital plane.

As illustrated in Figure 2c, even when the dish points eastward from north, most visible tracks cluster on the west side in the 2D projection. This occurs because the dish's boresight (i.e., the physical pointing direction) defines the center of the 2D projection, rather than the true zenith. Starlink's phased-array dishes typically maintain a fixed elevation (e.g., $\sim 70^{\circ}$) and a slight azimuth offset (e.g., $\sim 10^{\circ}$), limiting their observable sky region to a tilted partial dome relative to the local vertical. Consequently, connected satellites appear concentrated within this offset region, and the overall trajectory pattern becomes displaced and anisotropic. Furthermore, the non-uniform distribution of orbital planes at mid-latitudes leads to directional density asymmetry, with tracks accumulating toward the azimuth where orbital planes intersect the horizon.

Figure 2b further demonstrates that the 2D satellite image is not referenced to an absolute global frame, but is instead centered on the user dish's position and orientation. To validate this, we intentionally rotate the user dish during a single satellite connection window (e.g., a 15-second window), thereby altering the dish orientation. As a result, the dish-centric coordinate system changes, causing what was previously a continuous satellite track in the 2D image to split into two segments. Notably, only the white pixels correspond to the actual satellite trajectory. This observation indicates that the projected satellite trajectory in the 2D image is directly coupled with the dish's position and orientation.

Together, these effects explain why long-term averaged trajectory patterns are shifted and non-uniform, rather than centered and circular. They also motivate the design of a more reliable satellite identification method that accounts for dish orientation, positional uncertainty, and orbital geometry, forming the foundation for our satellite information alignment model.

2.3 Beyond Satellite Awareness: Toward Next Connection Prediction

While satellite awareness helps correlate link performance with orbital geometry, accurately predicting network performance in dynamic LEO networks remains challenging. Knowing only which satellite the user terminal is currently connected to is insufficient for reliable performance prediction, as link quality is also influenced by imminent handovers and rapidly evolving satellite positions. Moreover, in the absence of any open interface exposing the provider's handover policy, future satellite connections and satellite positions relative to the user cannot be directly observed, further complicating the prediction of near-term network performance.

Specifically, incorporating historical satellite features (e.g., elevation angle, distance, or identifier) can improve short-term prediction, but such approaches remain fundamentally reactive. Once the user terminal switches to a new satellite, prior correlations often no longer hold, and the performance context effectively resets. Because each satellite differs in orbit, visibility, and load, these transitions introduce significant discontinuities that static or per-satellite models cannot capture. In addition, it is possible that after a handover event is triggered, the user dish continues to connect to the same satellite. As illustrated in Figure 2d, the user dish remains connected to a single satellite for three consecutive connection windows (e.g., 45 seconds), which is reflected as a long continuous trajectory in the 2D image. Consequently, network performance exhibits significantly less variation compared to cases where the connection switches to a different satellite. In such scenarios, historical measurements provide stronger predictive power for forecasting near-term network performance.

In our preliminary experiments, we also implement an enhanced LSTM-based model [26] that incorporates satellite-related features, including satellite connection status and geometric position. As shown in Table 1, incorporating satellite information substantially improves prediction accuracy, achieving an MAE of 27.54. However, residual prediction errors remain non-negligible, highlighting the need for prediction mechanisms that are more robust to uncertainty when used to guide network policy adaptation. The primary reason for this limitation is Starlink's proactive switching of satellite connections approximately every 15 seconds, which forces the network into a new performance state after each handover. These frequent and abrupt transitions make it difficult for traditional time-series models to capture long-term dependencies using historical data alone. This challenge motivates us to more effectively leverage satellite dynamics in the prediction process. Accordingly, we incorporate both satellite handover behavior and real-time satellite status into our framework, decomposing the overall prediction task into multiple subtasks: connection prediction, window-level performance prediction, and per-second fine-grained prediction. These subtasks are designed to operate cooperatively, enabling the models to learn complementary relationships rather than functioning independently.

3 Multi-Granularity Prediction

Understanding and predicting satellite network dynamics is crucial for improving performance stability in LEO Internet systems. Due to frequent handovers between satellites and rapid spatiotemporal variation in link conditions, accurately predicting network performance remains challenging. Building on next-connection prediction, we design a multi-granularity prediction framework that captures satellite behavior and network quality across different time scales. At the satellite level, we predict the next connected satellite by modeling orbital geometry and dish orientation. At the temporal level, we predict the average network performance within each connection period (e.g.,15s), capturing coarse-grained performance trends. Finally, at the per-second level, building on the temporal-level prediction, we estimate short-term fluctuations in throughput and latency, providing fine-grained insights into instantaneous performance. Together, these layers form a unified predictive view that bridges physical satellite motion and end-to-end network experience, enabling proactive understanding and potential optimization of LEO Internet systems.

3.1 Satellite-Level Prediction

To obtain a precise insight into future network performance, it is essential to determine which satellite the user terminal will connect to next. Put simply, our goal is to predict the satellite that a user terminal will connect to in the upcoming connection period. User terminals typically switch to a new satellite approximately every 15 seconds, depending on satellite visibility and orbital dynamics. Our approach is to identify the next serving satellite from the set

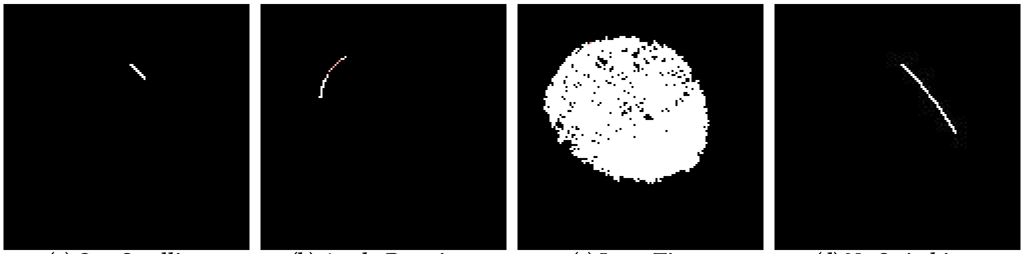

(a) One Satellite. (b) Apply Rotation. (c) Long Time. (d) No Switching.

Figure 2: 2D trajectories collected from Starlink Debug data.

of candidate satellites currently overhead. This approach is reliable because satellite motion is continuous, and the next connected satellite must appear among recently visible candidates. Before predicting the next connection, the first step is to identify the current serving satellite. We start from the gRPC 2D image and match it against candidate satellites to infer detailed satellite information through the following steps. **Converting the gRPC 2D image to global coordinates.** Each pixel (x_i, y_i) in the gRPC 2D image corresponds to a direction vector in the dish's local coordinate system. The pixel is first mapped to a local azimuth a_i and elevation e_i according to its position in the image. This mapping yields the local direction vector:

$$\mathbf{v}_i^{\text{loc}} = \begin{bmatrix} \cos(e_i)\sin(a_i) \\ \cos(e_i)\cos(a_i) \\ \sin(e_i) \end{bmatrix} \tag{1}$$

Given the dish orientation matrix $\mathbf{R}_{\text{dish}}$, the corresponding direction in the global coordinate system is obtained as

$$\mathbf{v}_i^{\text{glob}} = \mathbf{R}_{\text{dish}}^{\top}\mathbf{v}_i^{\text{loc}} \tag{2}$$

Matching gRPC Observations to TLE-Derived Tracks. For each candidate satellite s from the TLE dataset, its predicted global direction vector at time t_i is denoted by $\mathbf{u}_s^{\text{glob}}(t_i)$. The difference between the observed and predicted directions is measured by the angular distance:

$$\alpha_i = \arccos\big(\mathbf{v}_i^{\text{glob}} \cdot \mathbf{u}_s^{\text{glob}}(t_i)\big) \tag{3}$$

The satellite with the smallest mean angular distance over all observation samples is selected as the most likely match:

$$s^* = \arg\min_s \frac{1}{N}\sum_i \alpha_i \tag{4}$$

With accurate satellite positioning and enriched features such as distance, elevation, and azimuth, we construct a supervised dataset for predicting the next connected satellite. We employ a classification model to identify the most probable candidate from this refined set, enabling the system to predict upcoming connectivity transitions. Rather than training this model in isolation, we integrate it with the performance prediction framework, allowing joint learning of both connection dynamics and link quality evolution.

3.2 Temporal-Level Prediction

To accurately model and predict Starlink network behavior, we propose a joint prediction architecture that simultaneously learns satellite connectivity status and link performance trends. As shown in Figure 3, the architecture consists of two interconnected prediction branches: one for predicting the next satellite connection, and another for predicting average network performance over the upcoming connection window.

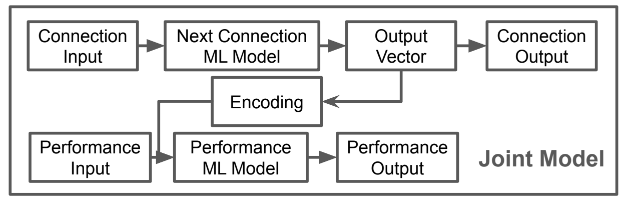

Figure 3: Joint Prediction Model.

The connection prediction branch takes as input the current user dish orientation and a set of visible satellite candidates. We use a Gated Recurrent Unit (GRU) [4]-based temporal encoder and combine it with discrete embeddings for satellite identifiers and hardware versions. The resulting output also serves as a learned representation that captures satellite-specific behavioral patterns.

The performance prediction branch incorporates both past performance indicators (e.g., throughput, jitter, and handover history) and the shared representation from the connection prediction branch. By conditioning performance prediction on the likely next satellite, the model is better equipped to predict performance shifts caused by upcoming handovers or spatial transitions. We observe that user dishes often reconnect to the same satellite across consecutive handover events, suggesting that similarities between the current connection and candidate satellites can provide useful cues for predicting future connections. Motivated by this observation, we apply an attention mechanism over the satellite candidate set to emphasize the most relevant candidates, enabling more accurate and robust performance prediction.

The two branches are trained jointly with a shared encoder. The overall loss function is a weighted combination of the connection prediction loss and the performance regression loss:

$$\mathcal{L}_{\text{total}} = \alpha \cdot \mathcal{L}_{\text{connection}} + \beta \cdot \mathcal{L}_{\text{performance}} \tag{5}$$

Here, α and β balance the influence of each task. This joint training strategy encourages the encoder to learn shared, generalizable representations that are useful for both connection prediction and performance estimation, thereby improving the robustness and accuracy of the prediction system.

3.3 Second-Level Prediction

While average performance prediction provides coarse-grained guidance over each satellite connection period, latency-sensitive applications and real-time traffic steering require finer temporal resolution. To address this need, we develop a per-second performance predictor that estimates throughput in the next second within each active satellite window.

This predictor builds on the output of the average performance model and leverages recent historical data collected within the current connection window. Specifically, we use a sequence of past per-second performance metrics including throughput, latency, and packet loss over a fixed-length horizon as the input.

We implement the model using a GRU architecture, which effectively captures temporal dependencies in the performance signals. The GRU processes the sequential input and outputs a single scalar representing the predicted throughput for the next second. Instead of predicting the absolute performance value, we predict the performance fluctuation relative to the temporal-level prediction.

4 Evaluation Results

We implement our framework on a real Starlink Standard Kit (Generation 3). To collect network metrics, we develop a custom UDPPing program to measure latency and packet loss, and use iperf3 for throughput evaluation. We also utilize starlink-grpc-tools [22] and Celestrak [3] to gather satellite telemetry and orbital information. In our experimental setup, the Starlink user terminal and client are deployed in Los Angeles, CA, while a remote server located in Minnesota is used to generate and capture end-to-end network traffic.

4.1 Prediction Model Performance

We evaluate the effectiveness of our prediction framework using two key metrics: mean absolute error (MAE) for throughput prediction and connection error rate (CER) for connection prediction. Table 2 reports the prediction performance for throughput and latency across different model configurations.

We replicate a self-similar model [27] that captures long-term temporal correlations in network data by feeding multi-scale time-series structures into an LSTM. Our proposed multi-granularity model achieves a throughput MAE of 8.37 Mbps and a CER of 0.1685. As shown in Figure 4, our model outperforms the self-similar baseline by first predicting satellite-level performance and then capturing fine-grained fluctuations conditioned on this prediction, following the principle of multi-granularity modeling. These results highlight the advantage of jointly learning satellite-aware representations that integrate both connection dynamics and path-specific characteristics. By sharing learned features across tasks, the model generalizes more effectively and achieves higher robustness in Starlink's highly dynamic environment.

Model	CER	Throughput MAE (Mbps)	Latency MAE (ms)
Self-Similar	N.A.	29.54	26.31
Multi-Granularity	0.1685	8.37	14.25

Table 2: Performance and Connection Prediction.

4.2 Multi-Granularity Prediction into Scheduling

To demonstrate how our prediction results can enhance existing satellite networking solutions, we reimplement three representative transport-layer models—SaTCP [1], SatPipe [28], and LeoCC [12]—on our real Starlink testbed. We then integrate our prediction outputs into each algorithm to evaluate how predictive information can improve their performance. The results are shown in Figure 5. The blue bars indicate the original performance achieved by the respective models, while the red bars show the additional improvement obtained after incorporating our prediction framework. For SaTCP, which freezes the congestion window during satellite switching, we apply our channel capacity prediction to adjust the maximum congestion window size before the last loss event in the CUBIC algorithm. This strategy allows the sender to better utilize available bandwidth when switching to a higher-capacity satellite, while mitigating packet loss when transitioning to a weaker link. For SatPipe, beyond its deterministic handover detection mechanism, we incorporate our fine-grained (per-second) throughput prediction into the BBR [2] model to refine the estimated bottleneck bandwidth, enabling more responsive and efficient sending-rate adjustments. For LeoCC, we replace the aggressive and moderate estimators with *prediction+uncertainty* and *prediction−uncertainty* variants, respectively. This design improves adaptation when the user terminal switches to a different satellite with substantially different link capacity. Overall, these experiments illustrate that integrating predictive information into existing transport algorithms significantly enhances robustness and throughput under highly dynamic LEO conditions.

5 Discussion and Future Work

By revealing how satellite transitions, geometry, and path heterogeneity jointly shape throughput and latency, we establish prediction as a critical foundation for future satellite-aware networking. In future work, we will extend our evaluation to Quality-of-Experience (QoE) analysis, quantifying how accurate prediction of upcoming connections and short-term performance fluctuations affects real applications such as video streaming and interactive services. This direction will bridge the gap between physical-layer satellite dynamics and user-perceived quality, guiding the design of future predictive and adaptive satellite networking systems.

Moreover, LEO systems support a wide range of mobility scenarios, including vehicular, maritime, and other use cases in which user movement is inherent. Such mobility introduces additional challenges for network performance prediction, as the user terminal's coordinates continuously change, directly affecting link geometry, propagation conditions, and, consequently, network performance. In future work, we plan to investigate how user terminal mobility impacts the accuracy and robustness of network performance prediction.

6 Related Work

Recent research on Low Earth Orbit (LEO) satellite systems has primarily focused on measurement, characterization, and system extensions [5]. Studies such as [8, 10, 17, 18] analyze Starlink's latency, throughput, and handover dynamics, revealing strong temporal and spatial variability in link performance. While these works provide valuable empirical insights, they remain largely observational and do not address how such dynamics can be predicted or proactively exploited. Recent advances in LEO transport protocols, including SaTCP [1], SatPipe [28], and LeoCC [12], aim to make congestion

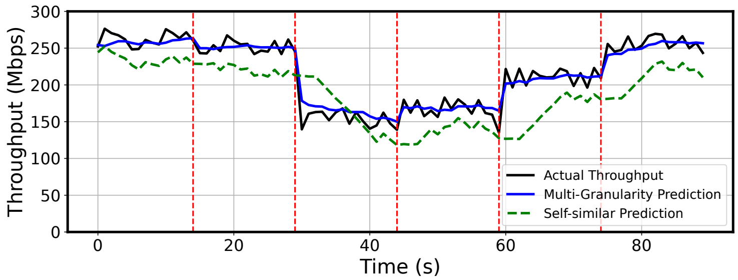

Figure 4: Prediction Track.

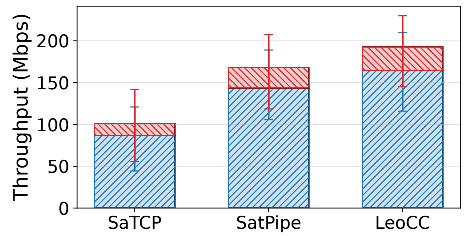

Figure 5: Throughput Improvement.

control more adaptive to satellite-induced dynamics. These systems rely on link-layer feedback or deterministic timing models to mitigate performance degradation during frequent handovers. However, their effectiveness is limited by the lack of accurate prediction of upcoming satellite transitions and link fluctuations. Our work complements these efforts by developing a multi-granularity prediction framework that forecasts both satellite connectivity and fine-grained link performance.

7 Conclusion

We design a multi-granularity prediction framework that predicts both upcoming connections and performance variations across multiple time scales. Our results demonstrate that accurate prediction is essential for understanding and mitigating the dynamics of LEO Internet systems, and lay the groundwork for future satellite-aware applications and network optimization.

Acknowledgments

We thank the anonymous reviewers for their valuable feedback. This work was supported in part by NSF under Grants CNS-2333489, CNS-2409269, CCF-2212318, CNS-2220286, CNS-2220292, CNS-2321531, CNS-2323174, and CNS-2321532.

References

[1] Xuyang Cao and Xinyu Zhang. 2023. SaTCP: Link-layer informed TCP adaptation for highly dynamic LEO satellite networks. In *IEEE INFOCOM 2023-IEEE Conference on Computer Communications*. IEEE, 1–10.
[2] Neal Cardwell, Yuchung Cheng, C Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016. Bbr: Congestion-based congestion control: Measuring bottleneck bandwidth and round-trip propagation time. *Queue* 14, 5 (2016), 20–53.
[3] CelesTrak. 2025. Starlink Two-Line Element Set. https://celestrak.org/NORAD/elements/gp.php?GROUP=starlink&FORMAT=tle. Accessed: 2025-10-10.
[4] Rahul Dey and Fathi M Salem. 2017. Gate-variants of gated recurrent unit (GRU) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*. IEEE, 1597–1600.
[5] Geneva Ecola, Bill Yen, Ana Banzer Morgado, Bodhi Priyantha, Ranveer Chandra, and Zerina Kapetanovic. 2025. SARLink: Satellite Backscatter Connectivity using Synthetic Aperture Radar. In *Proceedings of the 23rd ACM Conference on Embedded Networked Sensor Systems*. 398–410.
[6] Hao Fang, Haoyuan Zhao, Feng Wang, Yi Ching Chou, Long Chen, Jianxin Shi, and Jiangchuan Liu. 2025. Streaming media over LEO satellite networking: A measurement-based analysis and optimization. *ACM Transactions on Multimedia Computing, Communications and Applications* 21, 9 (2025), 1–24.
[7] 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*. 45–51.
[8] Tianze Huang, Qing Li, Chenren Xu, Mengwei Xu, Shangguang Wang, Gang Huang, and Xuanzhe Liu. 2025. From Earth to Orbit: Launch Sequence Optimization for LEO Mega-Constellations. *IEEE Transactions on Mobile Computing* (2025).
[9] Liz Izhikevich, Reese Enghardt, Te-Yuan Huang, and Renata Teixeira. 2024. A global perspective on the past, present, and future of video streaming over starlink. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 8, 3 (2024), 1–22.
[10] 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 (2024), 1–26.
[11] Sangeetha Abdu Jyothi. 2021. Solar superstorms: planning for an internet apocalypse. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 692–704.
[12] Zeqi Lai, Zonglun Li, Qian Wu, Hewu Li, Jihao Li, Xin Xie, Yuanjie Li, Jun Liu, and Jianping Wu. 2025. LeoCC: Making Internet Congestion Control Robust to LEO Satellite Dynamics. In *Proceedings of the ACM SIGCOMM 2025 Conference*. 129–146.
[13] Kexuan Liu, Yasheng Zhang, and Shan Lu. 2025. A Dynamic Spatio-Temporal Traffic Prediction Model Applicable to Low Earth Orbit Satellite Constellations. *Electronics* 14, 5 (2025), 1052.
[14] Yu Liu, Qiao Jin, Zejun Zhang, Bo Han, Svetlana Yarosh, and Feng Qian. 2024. HoloClass: Enhancing VR Classroom with Live Volumetric Video Streaming. In *Adjunct Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–3.
[15] Steven H Low, Fernando Paganini, and John C Doyle. 2002. Internet congestion control. *IEEE control systems magazine* 22, 1 (2002), 28–43.
[16] 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.
[17] 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.
[18] Nitinder Mohan, Andrew E Ferguson, Hendrik Cech, Rohan Bose, Prakita Rayyan Renatin, Mahesh K Marina, and Jörg Ott. 2024. A multifaceted look at starlink performance. In *Proceedings of the ACM Web Conference 2024*. 2723–2734.
[19] Joanna Najder and Krzysztof Sośnica. 2021. Quality of orbit predictions for satellites tracked by SLR stations. *Remote Sensing* 13, 7 (2021), 1377.
[20] Hao Peng and Xiaoli Bai. 2018. Improving orbit prediction accuracy through supervised machine learning. *Advances in space research* 61, 10 (2018), 2628–2646.
[21] Liangang Qi, Enqiang Wang, Tianfang Xu, Yuan Zhu, and Yun Zhao. 2025. Improving satellite network efficiency with terminal traffic prediction and SQP-SRA algorithm. *Computer Communications* (2025), 108293.
[22] sparky8512. 2025. starlink-grpc-tools: Random scripts for interacting with the SpaceX Starlink user terminal via gRPC. https://github.com/sparky8512/starlink-grpc-tools. GitHub repository, accessed January 24, 2026.
[23] 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*. 37–43.
[24] 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.
[25] Chenren Xu, Yuhan Zhou, Jing Wang, Ruihan Li, Lingyang Song, and Guangyu Zhu. 2025. HiMo: End-to-End Congestion Control for High Speed Rail Data Networking. *IEEE Transactions on Intelligent Transportation Systems* 26, 6 (2025), 7715–7725.
[26] Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. 2019. A review of recurrent neural networks: LSTM cells and network architectures. *Neural computation* 31, 7 (2019), 1235–1270.
[27] Yan Zhang, Yong Wang, Haotong Cao, Yihua Hu, Zhi Lin, Kang An, and Dong Li. 2025. Self-similar traffic prediction for LEO satellite networks based on LSTM. *IET communications* 19, 1 (2025), e12863.
[28] Ding Zhao, Xinyu Zhang, and Myungjin Lee. 2025. SatPipe: Deterministic TCP Adaptation for Highly Dynamic LEO Satellite Networks. In *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*. IEEE, 1–10.
[29] 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*. 109–120.