# Inferring Satellite Network Architectures and Orbital Regimes from End-to-End Traffic

Sukwha Kyung
Arizona State University
skyung1@asu.edu

Jaejong Baek
Arizona State University
jbaek7@asu.edu

Gail-Joon Ahn
Arizona State University
gahn@asu.edu

***Abstract*—The rapid expansion of emerging satellite communication networks has transformed global connectivity by extending broadband access to geographically isolated and infrastructure-deficient regions. However, it simultaneously enlarges the attack surface by introducing distinct network behaviors, routing architectures, and control mechanisms. In this work, we present a multi-stage traffic fingerprinting approach that can distinguish satellite network traffic from terrestrial flows, inferring the underlying routing architecture and identifying the orbital regime of the satellite network. We construct a controlled simulation environment using NS-3, emulating dynamic satellite network behaviors, including variable latency, intermittent connectivity, and handover-induced disruptions, while capturing end-to-end packet traces for analysis. We also analyze network-level features to determine their discriminative power in differentiating satellite traffic from ground-based flows. Our experimental results show measurable statistical deviations in temporal and structural traffic characteristics across link types, demonstrating the feasibility of traffic-based fingerprinting for satellite communication systems.**

## I. Introduction

The advancement of large-scale Low Earth Orbit satellite constellations, such as Starlink [1], OneWeb [2], and Kuiper [3], is transforming global communication infrastructures by extending broadband connectivity to geographically isolated and infrastructure-deficient regions. Beyond providing Internet access, these systems serve as a fundamental component of next-generation space–terrestrial networks supporting 5G/6G services [4], [5], intelligent transportation [6], and aerospace communication backbones networks [7].

However, despite this rapid evolution, it remains unclear to what extent the architectural and orbital properties of satellite networks are observable from network-layer measurements. End-to-end traffic inherently captures the cumulative effects of propagation delays, orbital motion, handover dynamics, and in-orbit routing strategies. If such characteristics manifest as measurable flow-level signatures, passive observations can reveal the presence of satellite links, their routing architectures, and even the orbital regimes of the constellation. As a result, the inference of infrastructure characteristics from encrypted traffic introduces security concerns, as it may enable reconnaissance, targeted disruption, or infrastructure-aware traffic manipulation [8], [9]. These considerations highlight the need for a systematic and quantitative understanding of how satellite network structures are reflected in observable traffic dynamics.

In this work, we present a study that empirically characterizes the observability of satellite network architectures and orbital regimes from flow-level traffic. We introduce a multi-stage inference approach that distinguishes satellite paths from terrestrial paths, identifies in-orbit routing architectures, and classifies orbital regimes using flow-level observables from an end-to-end vantage point. Using a data-driven simulation environment with real-world Starlink traffic traces, we extract the dominant traffic features that expose satellite network structures and quantify their impact on inference performance.

Our evaluation shows that orbital regimes are highly distinguishable, achieving up to 98.1% accuracy under bent-pipe configurations, while the inference of routing architectures remains more challenging in ISL-enabled networks (77.1%). These results indicate that propagation delay dynamics provide strong and persistent signatures at the network layer, whereas in-orbit routing and gateway level optimizations partially obscure link-specific characteristics. Furthermore, feature analysis reveals that the variance of round-trip time (RTT) is the dominant contributor to the inference outcome, highlighting temporal variability as the primary source of satellite network observability from end-to-end measurements.

Overall, the contributions of this work are three-fold:

1. We present a systematic study that empirically characterizes the observability of satellite network architectures and orbital regimes from flow-level traffic.
2. We demonstrate multi-stage inference of satellite versus terrestrial paths, routing architectures, and orbital regimes from an end-to-end vantage point.
3. We identify RTT variability as the dominant indicator of orbital dynamics, with implications for traffic monitoring and management in space–terrestrial networks.

## II. Background

### A. Orbital Regimes and Network-Layer Implications

Satellite systems are classified by orbital altitude, which directly determines propagation delay, satellite visibility time, and handover frequency. First, *Low Earth Orbit* (LEO) satellites (160 – 2,000 km) move rapidly relative to Earth's surface, resulting in short visibility windows and frequent handovers. Their low altitude yields lower propagation delays, while dynamic topology and handovers introduce pronounced temporal variability. *Medium Earth Orbit* (MEO) systems (2,000 km up

to GEO altitude) exhibit longer visibility periods and higher propagation delays than LEO, leading to more stable but still exhibit time-varying latency characteristics. Finally, *Geostationary Earth Orbit* (GEO) satellites (35,786 km) appear stationary to ground observers, largely eliminating handovers and producing consistently large round-trip delays. As a result, different orbital regimes induce distinct patterns in latency magnitude and variability that can be reflected in flow-level traffic.

### B. Relay and In-Orbit Routing Architectures

End-to-end traffic characteristics are also shaped by the satellite routing architecture. Specifically, architectural differences determine how prominently orbital dynamics and routing structures manifest in end-to-end measurements.

In a *bent-pipe* architecture, the satellite acts as a transparent relay without in-orbit forwarding; packets are downlinked to a ground gateway for terrestrial routing [10]. Accordingly, end-to-end delay is strongly influenced by the propagation distance between the ground segment and the satellite, often preserving distance-dependent latency signatures in observed traffic.

In contrast, *inter-satellite links* (ISLs) enable in-orbit forwarding, allowing packets to traverse multiple satellites before downlinking to a gateway [11]. Such spaceborne routing can reduce reliance on intermediate ground relays and may alter or partially mask link-specific delay signatures through dynamic path selection and gateway scheduling.

## III. Approach

We present our approach for inferring satellite network characteristics from end-to-end flow-level observables. We first discuss our threat model, followed by our approach including network simulation and feature extraction.

### A. Threat Model

We consider a passive adversary that monitors traffic at a ground gateway where terrestrial and non-terrestrial paths converge. The adversary can observe packet timing and flow-level statistics but cannot access or decrypt the payload. This setting reflects an end-to-end vantage point in which only traffic metadata is available. The threat model is illustrated in Figure 1. In our threat model, we formulate satellite traffic fingerprinting as a hierarchical inference problem over observed flows. Given traffic captured at the monitoring point, the objective is to infer: (1) whether the path includes a satellite segment, (2) the underlying routing architecture (bent-pipe or ISL), and (3) the orbital regime (LEO, MEO, or GEO). This formulation enables systematic analysis of how orbital dynamics and relay strategies manifest in flow-level observables.

### B. Satellite Network Simulation

We implement a data-driven satellite network simulation in NS-3 to generate end-to-end flow-level traffic under controlled orbital and routing configurations. Satellite nodes are equipped with mobility models that approximate orbital motion, while ground stations and user terminals are placed at fixed geographic locations.

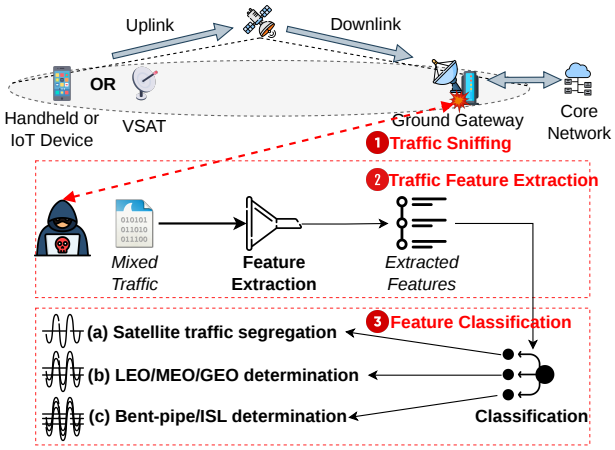


Fig. 1: Threat Model for Satellite Traffic Fingerprinting.

We configure representative orbital regimes to capture distinct latency characteristics: LEO at 500 km (90-minute period), MEO at a representative medium-altitude orbit, and GEO at 35,000 km as a fixed-position satellite. Satellite handovers are implicitly modeled through time-varying link availability as satellites move in and out of ground-station visibility. This simulated connectivity behavior introduces latency fluctuations and flow disruptions, thereby capturing the operational characteristics observed in real-world satellite network environments.

Satellite–ground and inter-satellite links are modeled as point-to-point channels with propagation delays derived from geometric distance. Bent-pipe architectures are realized by disabling in-orbit forwarding and routing traffic through ground gateways, whereas ISL-enabled configurations allow multi-hop spaceborne paths prior to downlink.

Traffic is generated at the application layer using both TCP and UDP with configurable patterns (constant bit-rate, bursty, and adaptive flows) to emulate broadband access and data delivery workloads. Packet traces are collected at gateway observation points, and flows are segmented into fixed-length time windows to extract flow-level feature vectors that capture the temporal dynamics of satellite communication.

### C. Feature Extraction

To jointly perform traffic classification and feature importance analysis, we extract flow-level observable features from packet traces captured at the gateway. Packets are first aggregated into flows using a 5-tuple identifier $f$(src IP, dst IP, src port, dst port, protocol) and then segmented into fixed-length time windows of duration $T$.

*Flow windowing.* Let a packet be represented as $p_i = (t_i, \ell_i, \delta_i, \sigma_i)$, where $t_i$ is the timestamp, $\ell_i$ is the packet length (bytes), $\delta_i \in \{\uparrow, \downarrow\}$ indicates direction (uplink/downlink), and $\sigma_i$ denotes a header-level type (e.g., DATA or ACK). For window index $w \in \mathbb{Z}_{\geq 0}$, the interval

can be defined as

$$I_w = [t_0 + wT,\ t_0 + (w+1)T). \tag{1}$$

The set of packets belonging to flow $f$ within window $w$ is

$$\mathscr{P}_{f,w} = \{\, p_i : \mathrm{fid}(p_i) = f,\ t_i \in I_w \,\}. \tag{2}$$

From $\mathscr{P}_{f,w}$ we compute a fixed-dimensional feature vector

$$\mathbf{x}_{f,w} = \phi(\mathscr{P}_{f,w}) = \left[\mu_{\mathrm{rtt}},\ \sigma^2_{\mathrm{rtt}},\ S,\ \mathbf{h}_\ell,\ \lambda_{\mathrm{ho}},\ \sigma^2_\rho,\ \eta_{\mathrm{ack}},\ \eta_{\mathrm{idle}}\right]_{(f,w)}. \tag{3}$$

The above features jointly capture (i) latency magnitude and temporal variability ($\mu_{\mathrm{rtt}}, \sigma^2_{\mathrm{rtt}}$), (ii) traffic volume and rate dynamics ($S, \sigma^2_\rho$), (iii) disruption/continuity behavior associated with mobility ($\lambda_{\mathrm{ho}}, \eta_{\mathrm{idle}}$), and (iv) transport-layer feedback structure ($\eta_{\mathrm{ack}}$), which are expected to vary across terrestrial paths, satellite routing architectures, and orbital regimes.

Consequently, we extract the following features from each flow window: (1) *mean and variance of round-trip time*, (2) *flow size per window*, (3) *packet size distribution*, (4) *handover frequency*, (5) *throughput variance*, (6) *ACK-to-data ratio*, and (7) *idle-time ratio*. These features capture latency magnitude, temporal variability, and flow continuity, which are directly influenced by orbital altitude and routing strategy. We also calculated the Gini importance score [12] for each feature to quantify the contribution of each observable to the inference process.

The training dataset consists of both simulated and real-world traces. Simulated data are generated using the NS-3 environment for terrestrial, LEO, MEO, and GEO scenarios under bent-pipe and ISL configurations. We also included real-world timing features for LEO traffic in our data including RTT statistics and inter-packet arrival dynamics from publicly available Starlink measurements [13].

## IV. Evaluation

To evaluate our approach, we consider three different traffic classification scenarios: (1) we aim to *segregate satellite network traffic from mixed traffic*; (2) we investigate the feasibility of *fingerprinting satellite links* determining if the traffic is routed through either bent-pipe or ISL; and (3) classification regarding the type of satellite communications system. Specifically, we *distinguish LEO from MEO or GEO systems*.

In our analysis, we use Random Forest, an ensemble learning algorithm to extract traffic features and classify the traffic. Random Forest can directly exploit structural regularities in raw or minimally processed data, making it particularly suitable for training without manually defining features to focus on. We particularly focused on the features such as packet RTT, length, traffic throughput, ACK/Data ratio and data packet ratio, handover frequency, and idle time ratio as presented in Section III-C.

**Data Collection.** We generated synthetic traffic by simulating LEO, MEO, and GEO communications in NS-3. For LEO and MEO, we simulated both bent-pipe and ISL while GEO was simulated with bent-pipe relay only, because GEO satellites act as fixed relay nodes between ground terminals and terrestrial core networks. In other words, due to the stationary orbital geometry of GEO satellites relative to the Earth's surface, GEO satellites maintain continuous visibility with designated ground gateways.

For simulation, it is difficult to collect real-world traffic data from operational satellite networks and conduct real analysis on operational satellite networks. Therefore, data-driven simulation has become a mainstream methodology that is widely used in satellite network research [14]–[17]. Similarly, we simulated the link based on Starlink constellation model derived from its two-line element set [18]. We then generated synthetic end-to-end traffic from simulated user terminals to remote web servers. The traffic is routed based on the simulated link and the traffic is captured and recorded as Ethernet packets in PCAP files [19]. Through our simulation, we confirm that the testbed simulated packets can generate the same metrics as the real-world ones. We captured the traffic in our simulation continuously over 72 hours to ensure comprehensive temporal coverage.

Overall, the complete dataset comprises over $12,480$ labeled flow samples distributed across seven class configurations: terrestrial (2,400 flows), LEO bent-pipe (1,920 flows), LEO ISL (1,920 flows), MEO-Low bent-pipe (1,080 flows), MEO-Low ISL (1,080 flows), MEO-High bent-pipe (1,080 flows), and GEO bent-pipe (1,000 flows). Finally, the dataset is partitioned using stratified sampling into training (70%, $n$ = 8,736), validation (15%, $n$ = 1,872), and test (15%, $n$ = 1,872) subsets.

### A. Segregation of Satellite from Mixed Traffic

The first goal of our fingerprinting is to simply distinguish satellite traffic from terrestrial ones. Figure 2 shows the cumulative distribution function (CDF) of Terrestrial, LEO, MEO, and GEO traffic in simulation. Specifically, Figure 2 (a) shows the CDF of RTT for bent-pipe traffic from Terrestrial, LEO, MEO, and GEO networks, while Figure 2 (b) shows the CDF of RTT for ISL from Terrestrial, LEO, and MEO. Note that GEO was not simulated with ISL due to the geometry of GEO satellites.

In addition, although we mainly used Random Forest as our classifier in this work, we compared the classification performance of different machine learning models. The models we compared include Decision Tree, SVM, XGBoost, and Random Forest. The binary classification task of distinguishing satellite flows from purely terrestrial flows achieves the highest accuracy of 96.2% with the Random Forest classifier. Table I presents the comparative performance of all evaluated classifiers on this task.

In bent-pipe links, all satellite communication traffic was distinguishable from terrestrial networks with high accuracy. As shown in Table II, satellite traffic can be classified in mixed satellite-terrestrial settings with an accuracy of 96.2% and an F-1 score of 0.952. It is also evident in Figure 2 (a) that RTT can be utilized as a factor in distinguishing satellite network traffic from terrestrial traffic.

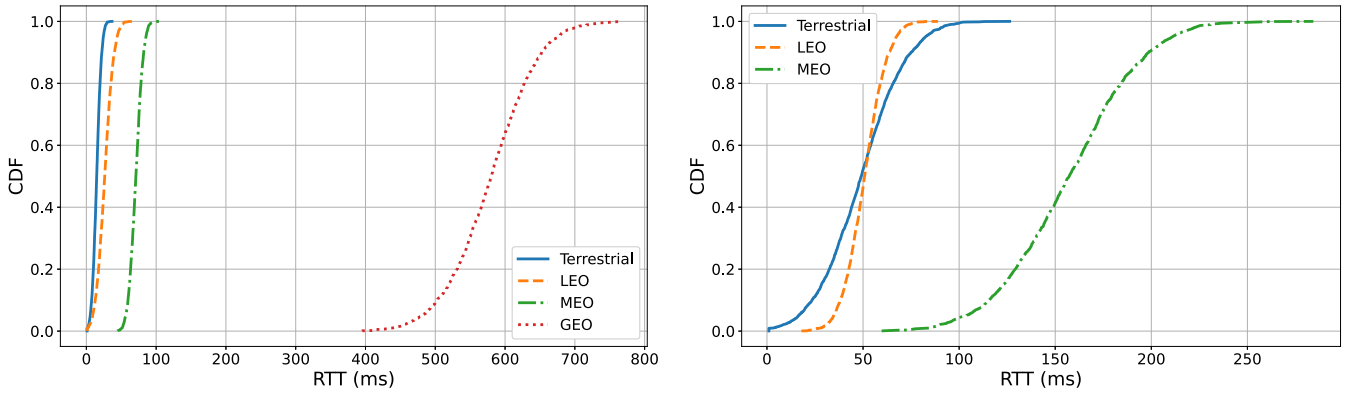


(a) CDF of bent-pipe traffic.

(b) CDF of ISL traffic.

Fig. 2: CDF of RTT for Terrestrial, LEO, MEO, and GEO Network traffic in Simulation.

TABLE I: Classifier Model Comparison (Terrestrial vs. Satellite)

| Classifier | Accuracy | Precision | Recall | F1-Score |
|---|---|---|---|---|
| Decision Tree | 0.911 | 0.903 | 0.897 | 0.900 |
| SVM (RBF Kernel) | 0.938 | 0.931 | 0.924 | 0.927 |
| XGBoost | 0.958 | 0.949 | 0.941 | 0.945 |
| Random Forest | 0.962 | 0.955 | 0.948 | 0.952 |

TABLE II: Performance of Random Forest-Based Satellite Network Fingerprinting

| Classification Task | Accuracy | Precision | Recall | F1-score |
|---|---|---|---|---|
| Terrestrial vs. Satellite | 0.962 | 0.955 | 0.948 | 0.952 |
| Bent-pipe vs. ISL | 0.771 | 0.740 | 0.722 | 0.731 |
| LEO vs. MEO | 0.942 | 0.936 | 0.921 | 0.928 |
| LEO vs. GEO | 0.981 | 0.974 | 0.962 | 0.968 |

Table II presents the overall classification performance of the proposed fingerprinting framework using a Random Forest classifier. The results indicate that passive traffic features provide meaningful signatures of satellite network characteristics. The terrestrial-satellite classification task achieves high accuracy, demonstrating that satellite traffic can be reliably segregated even under heterogeneous traffic conditions. From our observation, misclassifications occurred primarily in scenarios where satellite gateways were geographically proximate to users or when terrestrial paths exhibited elevated latency due to congestion, reducing the contrast between the two classes.

### B. Bent-Pipe vs. ISL

The bent-pipe and ISL classification tasks exhibited lower performance, with an F1-score of 73.1%. This result reflects the architectural similarity of the two satellite configurations at the physical layer and highlights the challenge of inferring in-orbit routing behavior from end-to-end observations alone. In addition, although ISL-enabled networks can introduce smoother rate adaptation and reduced directional asymmetry [20], raising a distinctive traffic feature, these effects are partially masked by ground-station scheduling. As a result, the traffic from the two routing systems shows low distinguishability.

### C. LEO vs. MEO vs. GEO.

The LEO versus MEO or GEO classification tasks achieved F1-scores of 92.8% and 96.8% respectively, indicating that orbital regime inference is highly feasible. Feature analysis reveals that the primary factors affecting the performance of fingerprinting are latency variance and average RTT. These effects arise from the distinct physical geolocation of satellites for each system. Also, handover rates were another obvious indication of LEO traffic because of the architectural difference between LEO and other satellite systems.

### D. Feature Analysis

It is evident that the traffic features of each satellite system possess distinctive characteristics that can be fingerprinted. Therefore, we analyzed how these features affect the performance of fingerprinting. To do so, we calculated the Gini importance scores for the features used in our Random Forest classifier. Gini importance, also referred to as the mean decrease in impurity, is a measure of feature relevance in tree-based learning models such as Random Forest. It is derived from the concept of Gini impurity, which quantifies the heterogeneity of class distributions within a node. For a classification task, the Gini impurity of a node is defined as $G = 1 - \sum_{k=1}^{K} p_k^2$, where $p_k$ denotes the empirical probability of class $k$ in the node. In a Random Forest, this quantity is aggregated over all trees and typically normalized so that the importance scores sum to 1.0, yielding a relative measure of feature influence. In other words, Gini importance offers a global explanation of the model by quantifying the relative contribution of each feature to classification performance in terms of impurity reduction.

Figure 3 illustrates the Gini importance scores obtained from the trained Random Forest classifier. Higher values indicate a stronger influence on the decision-making process of the model. The results show that *RTT Variance* (0.24) is the most influential feature, accounting for 24% of the total impurity reduction. This suggests that variability in RTT plays a dominant role in distinguishing traffic patterns within the dataset. The next most significant features are *Mean*

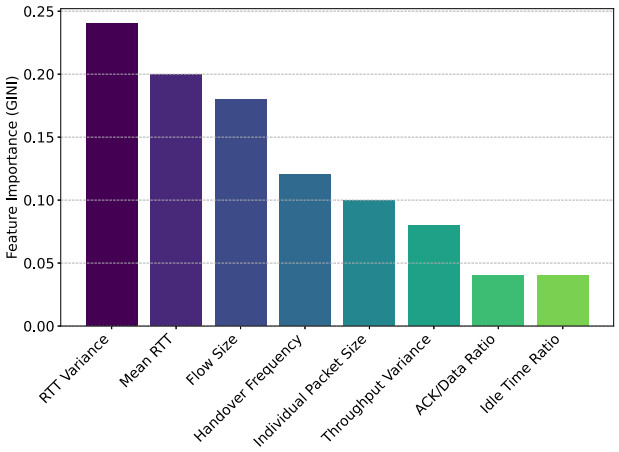


Fig. 3: Features affecting the performance of fingerprinting.

*RTT* (0.20) and *Flow Size* (0.18), both of which contribute substantially to classification performance. Collectively, these three temporal and volume-related characteristics account for 62% of the total importance scores, indicating that latency dynamics and aggregate traffic magnitude are primary discriminative factors for the model. *Handover Frequency* (0.12) and *Individual Packet Size* (0.10) exhibit moderate influence. Their contributions imply that mobility-related behavior and packet-level structural characteristics can be used as secondary and discriminatory factors. Overall, the distribution of importance values shows that the classifier primarily exploits latency variability, average delay, and traffic volume characteristics to construct decision boundaries.

## V. Related Work

We provide an overview of the related work and briefly compare our work with them. Those works emphasize the threat of unauthorized fingerprinting with further attacks.

Deng et al. [14] reframed dynamic satellite network topology, which is typically viewed as a resilience advantage, as a structural security risk requiring dedicated defenses by identifying time-varying bottleneck links as a fundamental and previously underexamined vulnerability in LEO satellite networks. Through simulation and the proposed SKYFALL analysis framework, the authors demonstrated that adversaries controlling geographically distributed user terminals can exploit dynamically shifting congestion points to launch stealthy link-flooding attacks that significantly degrade network throughput. The study further shows that traditional terrestrial defense mechanisms are ill-suited to such highly dynamic topologies and argues for LEO-specific mitigation strategies combining traffic engineering, routing adaptation, and resource control. This study provides empirical evidence supporting our argument that the exploitation of traffic fingerprinting poses significant security threats.

Oligeri et al. [21] introduced SatPrint, a physical-layer authentication framework designed to detect satellite signal spoofing by fingerprinting the unique noise and fading characteristics of satellite communication channels. In their study, the authors argue that satellite links exhibit propagation patterns that fundamentally differ from terrestrial transmitters commonly used by adversaries, enabling reliable discrimination without retraining models when new transmitters appear. Empirical evaluation showed that the method achieved detection accuracy exceeding 0.99 across tested scenarios while remaining computationally lightweight compared to deep-learning-based approaches. While this work mainly considered timing-based physical features of the communication, our approach adopts a multi-dimensional methodology with ensemble learning that considers a set of distinct features from flow-level traffic.

Another related work by Singh et al. [22] also focused on fingerprinting in satellite networks. The authors investigated website fingerprinting attacks on Tor traffic transmitted over Starlink's LEO satellite Internet compared to traditional fiber connections. In addition to observing Starlink's distinct networking characteristics, including 33% higher latency, increased packet exchanges, and more frequent TCP retransmissions, the authors also found that Tor traffic over satellite links remained equally vulnerable to state-of-the-art fingerprinting attacks. The unique characteristics of Starlink traffic were largely absorbed by the inherent latency and jitter of multi-relay routing of Tor, resulting in comparable attack success rates. This work has explored application-level fingerprinting within a specific satellite overlay network, while our work differs in its focus on architectural inference. Consequently, it highlights the significance of our study that can serve as a foundation for subsequent, more targeted attacks.

## VI. Discussion

This section elaborates on potential mitigation strategies for the satellite traffic fingerprinting, followed by a discussion of the limitations of this study and potential directions for future research in satellite network fingerprinting.

**Mitigation.** Traffic fingerprinting can still succeed in satellite settings because many discriminating features are embedded in communication metadata rather than payload contents [22]. Orbit-regime inference is particularly challenging to mitigate because propagation delay creates a strong, physics-driven signal. For example, GEO RTT has a hard lower bound and often appears substantially higher in operational services [23].

Mitigations therefore aim to reduce separability across multiple metadata features, with explicit trade-offs. Some mitigations include traffic shaping and padding. Fixed-rate regularization [24] and randomized padding [25] are evaluated on Starlink traces and shown to reduce classifier accuracy relative to undefended traffic, though often with non-trivial latency overhead [22]. For link-type inference, defenses can stabilize paths through Point-of-Presence (PoP) mapping or use overlays to present stable virtual paths. These mitigations are motivated by large-scale Starlink measurements showing that latency is bounded by nearby PoP availability and that the use of ISL can significantly increase routing path length and latency [26].

**Limitations.** In this work, the majority of timing-based features are derived from an open-source Starlink dataset, while other features, particularly non-timing-related attributes for MEO and GEO systems, are obtained from our simulated environment. Consequently, the experimental setup does not incorporate additional physical factors, such as weather conditions, that may influence satellite network performance [13], [27]. Nevertheless, the data-driven simulation is dominant in prior studies, as conducting controlled experiments on operational satellite networks remains largely impractical due to cost, accessibility, and deployment constraints. In addition, network jitter and other real-world environmental dynamics cannot be perfectly replicated within a simulation-based framework. However, our simulation design and empirical findings are consistent with prior studies on satellite traffic characteristics as well as observations derived from real-world datasets, lending credibility to the validity of our results [13], [28].

**Future Work.** Building upon our current research, we aim to investigate the practical application of our analysis in real-world. This includes scenarios where an adversary is located at different network vintage points such as user terminal and their impacts on relevant feature availability.

## VII. Conclusion

In this work, we demonstrated that satellite network traffic can be characterized at a finer granularity than simple satellite traffic segregation by examining the distinguishability of satellite links (bent-pipe and ISLs) and different satellite constellations (LEO, MEO, and GEO). Using a data-driven simulation approach, we showed that our ensemble learning-based fingerprinting method achieved an accuracy exceeding 94% in differentiating among various satellite network systems. Furthermore, we analyzed the key features of satellite flows that most significantly influence the accuracy of fingerprinting. Our analysis indicated that the variance of the RTT was the most influential feature, which is consistent with previous studies considering the characteristics of satellite networks. By investigating the feasibility of fingerprinting satellite traffic along with traffic feature analysis, our work offers practical utility in network monitoring and traffic management in addition to its security implications.

## References


[1] Starlink, "Starlink Services, LLC," https://starlink.com/, [Online; accessed 09-Feb-2026].

[2] Eutelsat, "Eutelsat OneWeb," https://www.eutelsat.com/, [Online; accessed 09-Feb-2026].

[3] Amazon, "Project Kuiper," https://leo.amazon.com/, [Online; accessed 12-Feb-2026].

[4] X. Fang, W. Feng, T. Wei, Y. Chen, N. Ge, and C.-X. Wang, "5G embraces satellites for 6g ubiquitous iot: Basic models for integrated satellite terrestrial networks," *IEEE Internet of Things Journal*, vol. 8, no. 18, pp. 14 399–14 417, 2021.

[5] X. Zhu and C. Jiang, "Integrated satellite-terrestrial networks toward 6G: Architectures, applications, and challenges," *IEEE Internet of Things Journal*, vol. 9, no. 1, pp. 437–461, 2021.

[6] K. Çelikbilek, Z. Saleem, R. Morales Ferre, J. Praks, and E. S. Lohan, "Survey on optimization methods for leo-satellite-based networks with applications in future autonomous transportation," *Sensors*, vol. 22, no. 4, p. 1421, 2022.

[7] O. Kodheli, E. Lagunas, N. Maturo, S. K. Sharma, B. Shankar, J. F. M. Montoya, J. C. M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff *et al.*, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2020.

[8] H. Cui, G. O. Karame, F. Klaedtke, and R. Bifulco, "On the fingerprinting of software-defined networks," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 10, pp. 2160–2173, 2016.

[9] A. Al-Shawabka, F. Restuccia, S. D'Oro, T. Jian, B. C. Rendon, N. Soltani, J. Dy, S. Ioannidis, K. Chowdhury, and T. Melodia, "Exposing the fingerprint: Dissecting the impact of the wireless channel on radio fingerprinting," in *INFOCOM*. IEEE, 2020, pp. 646–655.

[10] M. Koerner, "A bent pipe design for relaying signals received by an orbiting deep space relay station to a ground station," *DSN Progress Report*, vol. Jan-Feb, pp. 76–84, 1980.

[11] Z. Han, G. Zhao, Y. Xing, N. Sun, C. Xu, and S. Yu, "Dynamic routing for software-defined leo satellite networks based on isl attributes," in *GLOBECOM*. IEEE, 2021, pp. 1–6.

[12] H. Han, X. Guo, and H. Yu, "Variable selection using mean decrease accuracy and mean decrease Gini based on random forest," in *7th International Conference on Software Engineering and Service Science*. IEEE, 2016, pp. 219–224.

[13] J. Zhao and J. Pan, "Lens: A leo satellite network measurement dataset," in *the 15th ACM Multimedia Systems Conference*, 2024, pp. 278–284.

[14] Y. Deng, Q. Wu, Z. Lai, C. Gu, H. Li, Y. Li, and J. Liu, "Time-varying bottleneck links in leo satellite networks: Identification, exploits, and countermeasures." in *NDSS*, 2025.

[15] J. Vanlyssel, E. Sobrados, R. Anwar, G.-C. Roman, and A. Anwar, "Spychain: Multi-vector supply chain attacks on small satellite systems," *arXiv preprint arXiv:2510.06535*, 2025.

[16] E. Jedermann, M. Strohmeier, V. Lenders, and J. Schmitt, "Record: A reception-only region determination attack on leo satellite users," in *33rd USENIX Security Symposium*, 2024, pp. 6113–6130.

[17] Z. Lai, H. Li, Y. Deng, Q. Wu, J. Liu, Y. Li, J. Li, L. Liu, W. Liu, and J. Wu, "{StarryNet}: empowering researchers to evaluate futuristic integrated space and terrestrial networks," in *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2023, pp. 1309–1324.

[18] CelesTrak. (2026) Starlink tle. [Online]. Available: https://celestrak.org/NORAD/elements/gp.php?GROUP=starlink&FORMAT=tle

[19] G. Combs and M. Richardson, "The pcap capture file format," IETF Internet-Draft draft-ietf-opsawg-pcap-02, 2019.

[20] Y. Zhang, Y. Feng, L. Guo, Y. Gao, and Z. Chen, "Era-leo: An efficient rate adaptation with probabilistic constellation shaping for leo satellite networks," in *33rd International Conference on Network Protocols (ICNP)*. IEEE, 2025, pp. 1–12.

[21] G. Oligeri, S. Sciancalepore, and A. Sadighian, "Satprint: Satellite link fingerprinting," in *Proc. of the 39th ACM/SIGAPP Symposium on Applied Computing*, 2024, pp. 177–185.

[22] P. Singh, D. Barradas, T. Elahi, and N. Limam, "Connecting the dots in the sky: Website fingerprinting in low earth orbit satellite internet," in *NDSS*, 2024.

[23] G. Huston, "Using LEOs and GEOs," https://labs.apnic.net/index.php/2022/04/28/using-leos-and-geos/, [Online; accessed 01-Feb-2026].

[24] X. Cai, R. Nithyanand, T. Wang, R. Johnson, and I. Goldberg, "A systematic approach to developing and evaluating website fingerprinting defenses," in *ACM Conference on Computer and Communications Security*, 2014, pp. 227–238.

[25] M. Juarez, M. Imani, M. Perry, C. Diaz, and M. Wright, "Toward an efficient website fingerprinting defense," in *European Symposium on Research in Computer Security*. Springer, 2016, pp. 27–46.

[26] L. Izhikevich, M. Tran, K. Izhikevich, G. Akiwate, and Z. Durumeric, "Democratizing leo satellite network measurement," *ACM on Measurement and Analysis of Computing Systems*, vol. 8, no. 1, pp. 1–26, 2024.

[27] C. Beckman, P. Sjödin, C. Celine, F. Emil, and O. Robert, "Throughput analysis of starlink satellite internet: Study on the effects of precipitation and hourly variability with tcp and udp," in *NexComm Congress*. IARIA, 2025.

[28] S. Ma, Y. C. Chou, H. Zhao, L. Chen, X. Ma, and J. Liu, "Network characteristics of leo satellite constellations: A starlink-based measurement from end users," in *INFOCOM*. IEEE, 2023, pp. 1–10.