TVR Working Group
Internet-Draft
Intended status: Informational
Expires: 3 September 2026

Q. Wu
Y. Weng
Z. Lai
H. Li
Tsinghua University

2 March 2026

Path Verification in LEO Satellite Networks
draft-wu-tvr-path-verification-00

Abstract

Emerging satellite Internet constellations such as SpaceX's Starlink deploy thousands of broadband satellites and construct LEO satellite networks (LSNs) in space, significantly expanding the boundaries of today's terrestrial Internet. However, due to the unique global LEO dynamics, satellite routers inevitably pass through uncontrolled areas, suffering from security threats. It is important for satellite network operators (SNOs) to enable verifiable risk-avoidance routing to identify path anomalies.

This document specifies StarVeri, a network path verification framework tailored for emerging LSNs. StarVeri addresses the limitations of existing crypto-based and delay-based verification approaches and accomplishes efficient and accurate path verification through: (i) a segment-based verification protocol that divides paths into verifiable segments using dynamic satellite relays; and (ii) a hybrid verification approach combining cryptographic authentication with adaptive delay thresholds to verify each segment.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 3 September 2026.

Table of Contents

1. Introduction

Low-Earth Orbit (LEO) Satellite Networks (LSNs) are gaining tremendous popularity in recent years, with leading providers like SpaceX and Amazon Kuiper deploying thousands of satellites powered by laser inter-satellite links (ISLs) to provide global Internet services. Starlink, the largest operational LSN today, has launched more than 10,000 LEO satellites and attracted more than 12 million global subscribers.

However, unlike static terrestrial infrastructures, LEO satellites continuously orbit the Earth and inevitably traverse uncontrolled regions with potential security threats. These include electromagnetic interference , traffic hijacking, and information leakages through satellite hacking. For satellite network operators (SNOs), it is critical to not only enforce forwarding paths that bypass risk areas, but also verify that actual paths comply with avoidance policies.

Existing path verification approaches face significant challenges in dynamic LSNs. Crypto-based methods [CoNext11pv] [SIGCOMM14pv] [sec20pv] require per-hop authentication, imposing high computation overhead on resource-constrained satellites. Delay-based methods [SIGCOMM15pv] [sec17pv] [NDSS19pv] assume linear delay-distance relationships, which do not hold in LSNs due to frequent topology changes [INFOCOM22bg] [CoNext23bg] [HotNets18routing].

This document analyzes the impact of LEO dynamics on existing path verification approaches, demonstrating through case studies on Starlink that delay-based methods suffer from high verification inaccuracy in dynamic satellite environments. We present StarVeri, a path verification framework that specifies a segment-based verification protocol with cryptographic authentication and adaptive delay thresholds to achieve both efficiency and accuracy.

2. Characteristics of LEO Satellite Networks

2.1. Preliminaries for LEO Satellite Networks

Today’s LSNs like SpaceX’s Starlink consist of hundreds to thousands of LEO satellites that work as "routers in space", together with many geo-distributed ground stations connecting satellites and terrestrial network infrastructures. Satellites communicate with ground entities (e.g., ground stations or satellite terminals) via ground-satellite radio links (GSL). Besides, many satellite constellations (e.g., Starlink and Kuiper) leverage high-speed laser inter-satellite links (ISLs) for inter-satellite networking and communication. Satellites are typically organized in a +Grid topology, where each satellite

connects to two neighboring satellites in the same orbit and two in adjacent orbits. As satellites move at a high velocity in various orbital directions over the Earth, the entire LSN experiences frequent topology changes, especially in space-ground connections.

LSN traffic between two terrestrial nodes (e.g., from a Starlink's satellite terminal to a remote ground station) is forwarded by a network path constructed by uplink/downlink and ISLs. To deal with the unique LSN topology fluctuation, space routing mechanisms dynamically build and maintain paths for any two terrestrial nodes connected to the LSN.

2.2. Potential Risks in Uncontrolled Areas

As satellites move globally, they might enter an uncontrolled risk area, posing routing security threats in LSNs. In this document, we define a "risk area" as a geographical region where malicious attackers may exist and launch the following attacks on satellites above the area:

* Information stealing: Attackers in the risk area can eavesdrop on traffic and extract or speculate private data. Attackers in risk areas can eavesdrop on traffic forwarded between satellites.

* Traffic hijacking: More powerful attackers in a risk area can hijack the route and cause path inconsistency. They propagate error routing advertisements to allure surrounding satellites to forward packets to them and redirect traffic to specific nodes for censorships, or other man-in-the-middle attacks like packet injection, modification, and counterfeit.

Therefore, to avoid the above risks, it is essential for satellite network operators (SNOs) to: (i) apply avoidance policies to force their traffic to bypass potential risk areas, and more importantly, (ii) verify that the actual paths are consistent with the planned avoidance policies in practice.

3. Impact of LEO Dynamics on Path Verification

3.1. Today's Network Path Verification

Over the past decade, the network community has had a body of efforts working on network path verification. In particular, existing path verification efforts can be classified into two main categories: crypto-based and delay-based approaches.

3.1.1. Crypto-based path verification

One classic path verification approach is to embed the planned path (e.g., a sequence of nodes that build the network path) into the header of each packet. Then intermediate nodes authenticate and update related fields before sending to the next node [CoNext11pv] [SIGCOMM14pv] [sec20pv]. These approaches have three limitations in LSNs. First, each intermediate node needs to perform cryptographic operations like calculating MACs hop by hop, which involves high computation overhead that could be unaffordable for resource-constrained satellites. Second, the increased length of the verification header also leads to extra packet header overhead, resulting in goodput ratio reduction. Third, per-hop authentication inevitably involves additional processing delay on each node. Note that the high mobility of LEO satellites incurs frequent path changes [INFOCOM22bg] [CoNext23bg] [HotNets18routing], if the packet processing delay is too high, the pre-calculated path might be invalid on the route.

Although some recent works like PPV [IWQoS18pv] and MASK [ToN23pv] propose probabilistic authentication instead of verifying every packet hop by hop, the computation overhead can be still high when the traffic volume is large.

3.1.2. Delay-based path verification

Alibi Routing [SIGCOMM15pv] proposes a new path verification approach, exploiting the key idea that a detour from the planned path to any node in the risk area will breach the maintained delay by incurring significant extra delay. To verify whether a network path from a source (Src) to its destination (Dst) passes through the risk area, Alibi Routing pre-computes a target area that is far away from the risk area where possible verifiable relays (called alibis) are located. It enforces that the path from Src to Dst must pass through the alibis. Because the relay is very far from the risk area, if the Src to Dst path passes through the risk area, the observed end-to-end delay should be significantly higher than the maintained value.

However, these approaches are based on a fundamental assumption that the delay between two terrestrial nodes has a linear relationship with their great-circle distance. Although this assumption exists in many scenarios in today’s terrestrial Internet [SIGCOMM09linear], it is not applicable in LSNs with highly dynamic topology. An increase in path delay may not necessarily be caused by traversing the risk area but path changes due to topology fluctuation.

3.2. A Case Study: LEO Dynamics Affect the Verification Accuracy

We conduct a quantitative analysis to demonstrate how the unique topology fluctuations in LSNs affect the verification accuracy of existing delay-based approaches. Our analysis is based on the real Starlink constellation information and an LSN simulation based on StarryNet [StarryNet]. We build a virtual LSN consisting of 1584 satellites with 72 orbits and 22 satellites per orbit (Starlink Phase 1, Shell 1) and 165 geo-distributed ground stations around the world. We simulate about 2000 city pairs communicating over the LSN using shortest path routing.

First, we observe that the linear relationship assumption [SIGCOMM09linear] does not hold in LSNs. By plotting the relationship between great circle distance and Round-Trip Time (RTT) for LSN paths, we find significant scatter and deviation from a linear relationship. On our further investigation, we confirm that the delay increase is caused by frequent topology fluctuations and path changes in LSNs even if the source and destination are static on the ground.

Second, we observe the non-linear relationship can lead to inaccuracy for delay-based verification approaches. We simulate the Alibi Routing [SIGCOMM15pv] mechanism and set a static ground relay for each city pair. We calculate the ratio of false positive (FP, i.e., the real path does not traverse the risk area but is considered as passing through it) and false negative (FN, i.e., the opposite of FP) of each city pair in 24 hours with an interval of 1 second. These error ratios are defined as the proportion of the amount of time slots that FP or FN occurs to the total amount of time slots.

Figure 1 shows the CDF of error ratios across all city pairs:

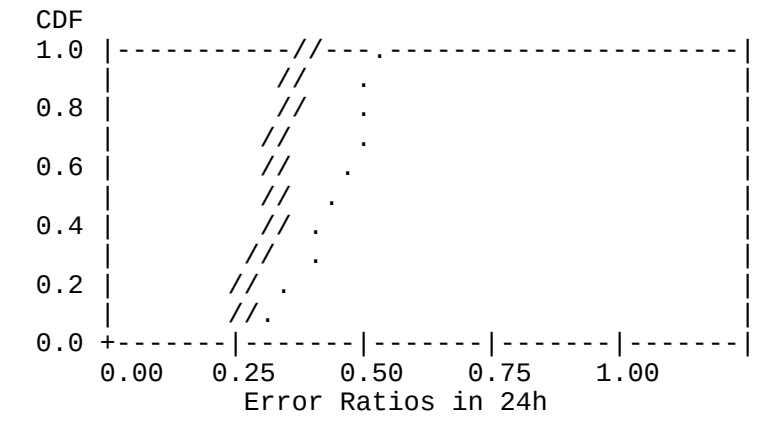

Figure 1: CDF of verification error ratios for Alibi Routing in LSNs

The results demonstrate that delay-based verification methods suffer from high inaccuracy. As shown in Figure 1, most city pairs experience false negative (FN) ratios around 25% and false positive (FP) ratios between 25-35%, indicating that existing delay-based approaches frequently fail to correctly identify whether paths traverse risk areas.

The root cause of this inaccuracy lies in the violation of the linear delay-distance assumption in dynamic LSNs. Unlike terrestrial networks where routing paths remain relatively stable, satellite networks experience continuous topology changes due to orbital motion. When a satellite moves or an ISL switches, the forwarding path changes even though the source and destination remain fixed on the ground. These path changes introduce delay variations that are indistinguishable from detours through risk areas, causing the delay-based verification to produce false alarms or miss actual violations. The fundamental challenge is that delay increases in LSNs can result from either legitimate topology changes or malicious path deviations, making delay alone an unreliable verification metric.

Therefore, accomplishing accurate path verification in dynamic LSN environments requires new approaches that account for topology dynamics while maintaining verification efficiency.

4. Path Verification Framework for LEO Satellite Networks

4.1. Framework Overview

StarVeri provides satellite network operators (SNOs) with the ability to verify path compliance between planned packet delivery paths and actual forwarding paths in dynamic LSN environments. The framework can be deployed on existing Path-Aware Networking (PAN) or source routing architectures (e.g., Segment Routing [RFC8402]) that allow path enforcement through pre-calculated satellite relays.

StarVeri adopts a segment-based verification approach that addresses the limitations of existing methods. At a high level, to verify a network path from source (Src) to destination (Dst), StarVeri uses a collection of dynamic satellite relays to split the entire path into a series of consecutive segments. Thus, verifying the entire path is equivalent to verifying each segment path. To reduce verification overhead, StarVeri only requires relays (instead of all nodes on the path) to authenticate forwarded packets. While the end-to-end delay of a network path fluctuates due to LSN dynamics, the delay fluctuation of a segment (e.g., between two adjacent relays) is less dramatic than the delay fluctuation of the entire path. StarVeri verifies each segment by comparing the real inter-relay delay with an estimated upper bound.

The framework consists of three key components:

* SNO Operation Center: Calculates and distributes relay configurations to end users
* Satellite Relays: Selected satellites that authenticate forwarded packets and verify segment paths
* End Nodes (Src/Dst): Source initiates verification fields; destination performs final path validation

Figure 2 illustrates the overall architecture:

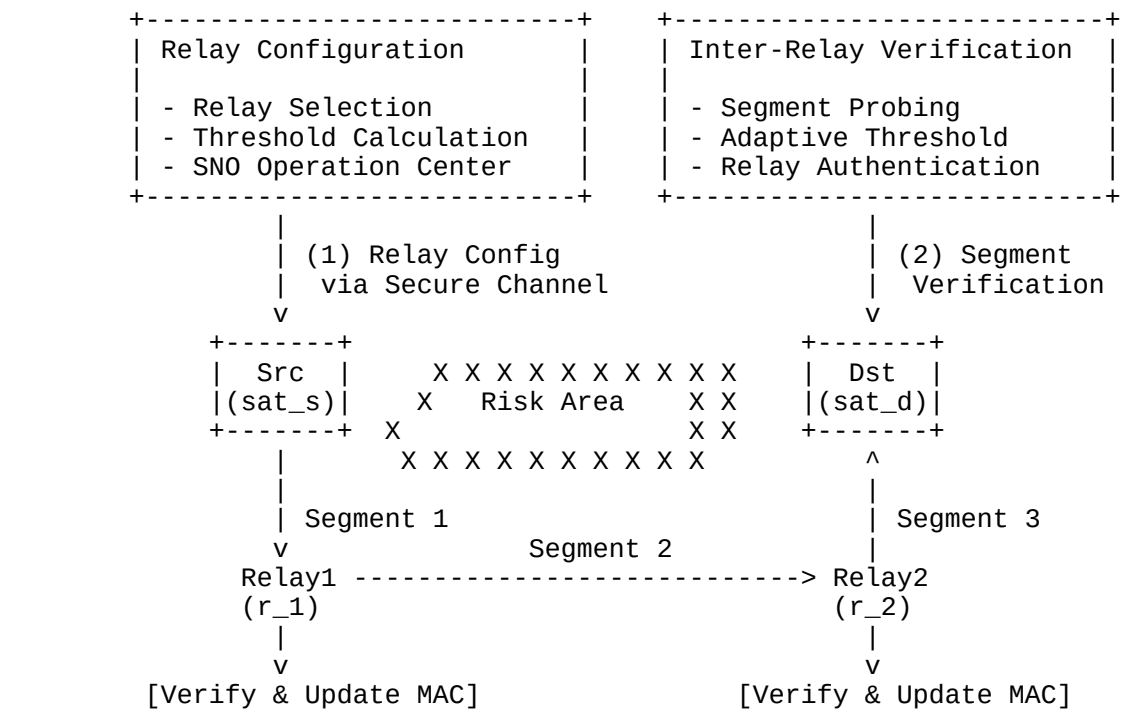

(3) Final Verification: Dst verifies all relay MACs and the last segment

Figure 2: StarVeri Architecture Overview

4.2. Verification Workflow

This section describes the core protocol mechanisms. The verification workflow for each communication pair (Src, Dst) and a given risk area proceeds as follows:

(1) Relay configuration: The SNO operation center calculates relay configurations and estimates segment detour delay thresholds based on predictable satellite trajectories and dynamic LSN topology. Once decided, the operation center delivers the relay configuration to Src via a secure channel.

(2) Intra-segment state probing: Each relay (including sat_d) periodically sends probing packets to its previous relay to measure the current segment delay. Probing packets are authenticated using shared symmetric keys to prevent tampering.

(3) Authentication fields initiation: Before packet departure, Src inserts the corrected sending timestamp, hash value, and relay authentication fields (AUTH) in the packet header.

(4) Relay-based segment verification: When a relay receives a packet, it verifies the segment path based on its probing ground truth and detour threshold. If the packet does not traverse the risk area, the relay computes a MAC value and updates its corresponding AUTH field with the timestamp.

(5) Destination verification: After receiving a packet, Dst verifies the packet's source, AUTH fields updated by relays, and the last segment path to determine whether the packet bypassed the risk area.

4.2.1. Packet header format

Before packet departure, Src constructs the verification header to enable segment-based path verification. This header carries authentication information that will be progressively updated by each relay along the path, allowing Dst to verify that the packet has bypassed the risk area through the designated relay sequence.

The verification header structure is shown in Figure 3.

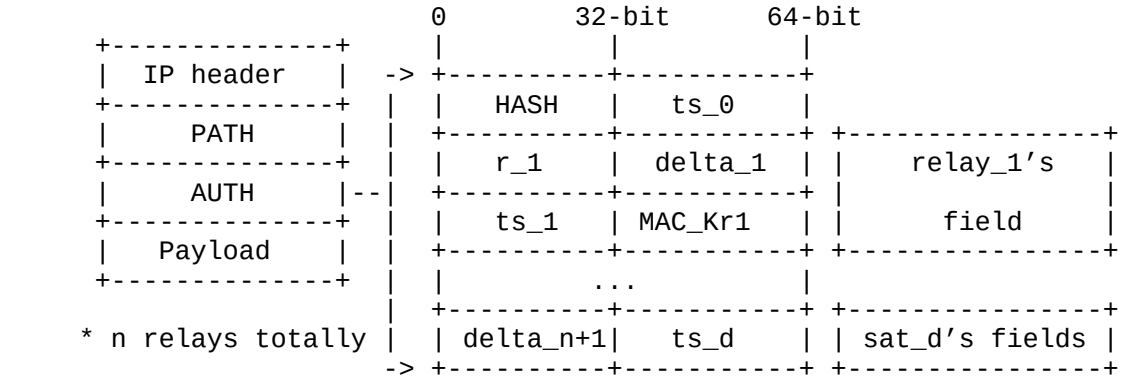

Figure 3: The AUTH header structure of StarVeri

The verification header consists of the following components:

* HASH (4 bytes): Hash value computed as H(Payload||PATH||\Delta||ts_0||RE)[0:4] using SHA-256 and truncated to 4 bytes. Here, PATH is the relay sequence, \Delta is the detour threshold vector, and RE is the risk area identifier.

* Initial Timestamp ts_0 (8 bytes): Corrected sending time ts_0 = ts_{send} + D_{access} + \alpha, where ts_{send} is the actual sending time, D_{access} is the access link delay, and \alpha is a clock synchronization offset.

* AUTH Chain: Reserved fields for relay authentication. The AUTH chain contains n+1 entries: n entries for relays and one entry for destination sat_d.

For each relay r_i (where i = 1, 2, ..., n), the AUTH entry contains: - Relay ID (r_i, 2 bytes): Identifier of the relay - Detour threshold (\delta_i, 2 bytes): Maximum acceptable detour delay from previous node to this relay, i.e., \delta_{r_{i-1},r_i} where r_0 denotes Src - Timestamp (ts_i, 8 bytes): Time when the relay processes the packet - MAC (4 bytes): Message Authentication Code MAC_{K_{r_i}}(HASH||ts_i||r_i)[0:4] using HMAC-SHA256 [RFC2104] [RFC6234], truncated to 4 bytes

For destination sat_d, the AUTH entry contains: - Detour threshold (\delta_{n+1}, 2 bytes): Maximum acceptable detour delay from last relay r_n to destination, i.e., \delta_{r_n,d} - Timestamp (ts_d, 8 bytes): Time when destination receives the packet

4.2.2. Relay verification procedure

When a relay r_i receives a packet, it performs the following verification steps:

1. Check that AUTH_{i-1} has been properly updated by the previous relay. If not, the packet is dropped.

2. Verify segment delay: Check that ts_i - ts_{i-1} \leq \delta_{r_{i-1},r_i} + dt_{r_{i-1},r_i}, where:

    * \delta_{r_{i-1},r_i}: Detour threshold provided by SNO operation center
    * dt_{r_{i-1},r_i}: Probed segment delay from periodic probing

If the delay exceeds this threshold, the packet is dropped as it may have traversed the risk area.

3. If verification passes, the relay updates its AUTH entry and forwards the packet:
    * Record current timestamp ts_i
    * Compute MAC: MAC_{K_{r_i}}(HASH||ts_i||r_i)[0:4] using HMAC
    * Update AUTH_i with: relay ID (r_i), detour threshold (\delta_i), timestamp (ts_i), and MAC
    * Forward packet to the next hop according to the PATH field

4.2.3. Destination verification procedure

Upon receiving a packet, Dst performs the following verifications in order:

1. Source authentication: Verify the packet source using existing authentication mechanisms (out of scope for this document).

2. Integrity check: Recompute HASH using SHA-256 and compare with the value in the packet header. If they do not match, the packet is dropped.

3. Relay chain validation: For each relay r_i in the AUTH chain, authenticate the MAC using HMAC with the shared session key K_{r_i}. If any MAC validation fails, the packet is dropped.

4. Last segment verification: Check that ts_d - ts_n \leq \delta_{r_n,d} + dt_{r_n,d}, where r_n is the last relay. If this check fails, the packet is dropped.

If all verifications pass, Dst accepts the packet as having successfully bypassed the risk area.

4.3. Supporting Mechanisms

This section describes supporting mechanisms that enable the verification protocol. These mechanisms are implementation-specific and out of scope for standardization, but are provided for informational purposes.

4.3.1. Relay configuration distribution

The SNO operation center distributes relay configurations to Src via a secure channel. The configuration contains: - Relay sequence: [r_1, r_2, ..., r_n] - Detour thresholds: [\delta_{s,r_1}, \delta_{r_1,r_2}, ..., \delta_{r_n,d}] - Validity period: Time range during which this configuration is valid

The specific relay selection algorithm is implementation-specific. Different approaches can be used to select relays that minimize detour delay while ensuring verifiability and avoiding risk areas.

4.3.2. Session key negotiation

When a path changes or a new session begins, end users, relays, and SNO negotiate session keys for authentication. Session key establishment is out of scope for this document. Existing key derivation methods such as DRKey or other secure key distribution mechanisms can be used.

4.3.3. Intra-segment probing

Each relay periodically (e.g., every tens of seconds) sends probing packets to its previous relay to measure the current segment delay. The probing mechanism is implementation-specific. A typical implementation includes:

* Probing packet format: Contains a nonce, timestamp, and hop-by-hop MACs for authentication
* Probing reply: Contains the probe ID, measured RTT, and authentication MACs
* Adaptive threshold computation: The measured delay dt_{r_{i-1},r_i} is used together with the detour threshold δ_{r_{i-1},r_i} to determine the segment delay upper bound

Alternative probing approaches or different probing frequencies can be used based on network conditions.

5. Conclusion

This document specifies StarVeri, a path verification framework designed for the unique challenges of LEO satellite networks. StarVeri addresses the limitations of existing verification approaches by introducing a segment-based verification protocol that enables efficient and accurate path verification in highly dynamic satellite environments.

The framework defines a verification protocol based on relay-divided segments. The protocol specifies: (i) a verification header format that carries authentication information progressively updated by each relay, (ii) relay verification procedures that check segment delays and update authentication fields, and (iii) destination verification procedures that validate the complete relay chain and path compliance. By verifying paths through segments rather than hop-by-hop, StarVeri reduces computation overhead on resource-constrained satellites while maintaining verification accuracy.

StarVeri provides satellite network operators with a practical mechanism to enforce and verify risk-avoidance routing policies, contributing to the security and reliability of emerging global satellite Internet services.

6. Security Considerations

This document does not represent a change to any aspect of the TCP/IP protocol suite and therefore does not directly impact Internet security.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. Normative References

[RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, https://www.rfc-editor.org/info/rfc2104 (https://www.rfc-editor.org/info/rfc2104)

[RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, https://www.rfc-editor.org/info/rfc6234 (https://www.rfc-editor.org/info/rfc6234)

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, https://www.rfc-editor.org/info/rfc8402 (https://www.rfc-editor.org/info/rfc8402)

8.2. Informative References

[CoNext11pv] "OPT: On-demand Path Tracing", ACM CoNext, 2011

[SIGCOMM14pv] "ICING: In-band Control for Inter-Domain Routing", ACM SIGCOMM, 2014

[sec20pv] "Secure Path Validation", Security Conference, 2020

[SIGCOMM15pv] Levin, D., Lee, Y., Valenta, L., Li, Z., Lai, V., Lumezanu, C., Spring, N., and B. Bhattacharjee, "Alibi Routing", Proceedings of the 2015 SIGCOMM, pp. 611-624, 2015

[sec17pv] Li, Z., Herwig, S., and D. Levin, "DeTor: Provably Avoiding Geographic Regions in Tor", 26th USENIX Security Symposium, pp. 343-359, 2017

[NDSS19pv] Kohls, K., Jansen, K., Rupprecht, D., Holz, T., and C. P旦pper, "On the Challenges of Geographical Avoidance for Tor", Network and Distributed System Security Symposium, 2019

[INFOCOM22bg] "LEO Satellite Network Background", IEEE INFOCOM, 2022
[CoNext23bg] Tanveer, H., Puchol, M., Singh, R., Bianchi, A., and R. Nithyanand, "Making Sense of Constellations: Methodologies for Understanding Starlink's Scheduling Algorithms", Companion of the 19th International Conference on Emerging Networking EXperiments and Technologies, pp. 37-43, 2023
[HotNets18routing] Handley, M., "Delay is Not an Option: Low Latency Routing in Space", Proceedings of the 17th ACM Workshop on Hot Topics in Networks, pp. 85-91, 2018
[IWQoS18pv] "PPV: Probabilistic Path Verification", IEEE IWQoS, 2018
[ToN23pv] "MASK: Masked Authentication for Scalable Path Verification", IEEE/ACM Transactions on Networking, 2023
[SIGCOMM09linear] "On the Linear Relationship between Delay and Distance", ACM SIGCOMM, 2009
[StarryNet] "StarryNet: LEO Satellite Network Simulator", 2023

9. Acknowledgments

Thanks to all of the contributors.

Authors' Addresses

Qian Wu
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: wuqian@cernet.edu.cn

Yuxuan Weng
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: wengyx25@mails.tsinghua.edu.cn

Zeqi Lai
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: zeqilai@tsinghua.edu.cn

Hewu Li
Tsinghua University
30 ShuangQing Ave
Beijing
100089
China
Email: lihewu@cernet.edu.cn