Time-varying Bottleneck Links in LEO Satellite Networks: Identification, Exploits, and Countermeasures

Yangtao Deng[†], Qian Wu[†‡], Zeqi Lai[†‡], Chenwei Gu[†], Hewu Li[†‡], Yuanjie Li[†‡], Jun Liu[†‡]
[†]Tsinghua University, [‡]Zhongguancun Laboratory
{dengyt21, gcw22}@mails.tsinghua.edu.cn, {wuqian, lihewu}@cernet.edu.cn,
{zeqilai, yuanjiel}@tsinghua.edu.cn, juneliu@mail.tsinghua.edu.cn

***Abstract*—In this paper, we perform a multifaceted study on the security risk involved by the unique time-varying bottleneck links in emerging Low-Earth Orbit (LEO) satellite networks (LSNs). We carry out our study in three steps. First, we profile the spatial and temporal characteristics of bottleneck links and how they might be exploited for bottleneck identification. Thus, the bottleneck links imposes a new risk of link flooding attack (LFA) on LSNs. Second, we propose SKYFALL, a new LFA risk analyzer that enables satellite network operators to simulate various LFA behaviors and comprehensively analyze the consequences on LSN services. Concretely, SKYFALL's analysis based on real-world information of operational LSNs demonstrates that the throughput of legal background traffic could be reduced by a factor of 3.4 if an attacker can manipulate a number of compromised user terminals to continuously congest the bottleneck links. Based on our analysis, we finally discuss the limitations of traditional LFA countermeasures and propose new mitigation strategies for LSNs.**

I. INTRODUCTION

Low-Earth Orbit (LEO) satellite networks (LSNs) are undergoing rapid development, drawing increasing attention for their potential to offer global Internet services. Major players in this space, such as SpaceX's Starlink [1], Amazon's Kuiper [2], and Boeing [3], are deploying mega-constellations consisting of thousands of satellites. With the use of high-speed laser or radio links [4], these LSNs are poised to offer worldwide, low-latency, and high-throughput Internet services. On the ground segment, a satellite could be connected to a ground station (GS). Numerous GSes [5], [6] have been constructed and put into use. In terms of user statistics, SpaceX announced reaching 3 million subscribers in May 2024 [7].

While LSNs hold immense potential for global access, security concerns remain a potential problem. Researchers and engineers are closely examining underlying risks in terminal systems and satellite communications, such as link jamming, electronic attacks [8], [9], energy-draining [10], and spoofing [11]. Besides, a security researcher at the Black Hat security conference demonstrated the ability to launch a fault injection attack on a Starlink user terminal (UT) and bypass its security protections by a homemade printed circuit board (PCB) [12]. Furthermore, on November 18, 2022, an organization claimed responsibility for three deliberate Distributed Denial-of-Service (DDoS) attacks against Starlink, leading to user reports of difficulty logging in for several hours [13]. Understanding the impact and scope of various potential security risks in LSNs, and exploring corresponding countermeasures is crucial for LSN operators.

One of the security risks is link-flooding attack (LFA) based DDoS, which represents a significant and growing concern in LSN security. Public Starlink data reveals that DDoS attack traffic accounted for 25.6% of the network and application layer attack traffic in the previous year up to September 2024 [14]. In [15] and [16], compromised UTs are manipulated to serve as bots and launch malicious traffic at each other, effectively overwhelming target links with LFA.

Due to the relatively even structure of today's mega-constellation and the uneven distribution of ground facilities [1], global traffic, and user distributions [14], there must be *bottleneck links* in LSNs. The bottleneck links refer to the downlink GSLs that transmit more background traffic from a wider range of places and populations, and these bottleneck links change frequently due to the dynamism. Because the disclosure of constellation design [17], satellite trajectories [18], [19] and ground station distribution [6], the network structure of an LSN as well as its bottleneck links are identifiable. If attackers identify these bottleneck links and launch an LFA, it can significantly impact the quality of network services in LSNs. Therefore, analyzing the adverse impact and potential risks of LFAs is crucial for LSNs.

Existing analysis methods of LFAs can be categorized into two main types. The first kind is analysis focusing on LFAs of terrestrial Internet, such as Coremelt [20] and Crossfire [21]. However, these studies do not quantify the impact of attacks on user throughput and network quality at a global scale under the circumstances of certain compromised malicious botnets. Other works target LFAs in satellite networks, such as [15] and [16]. Nevertheless, these works only focus on how to target certain inter-satellite links (ISLs) or ground-satellite links (GSLs) at a static time slot under various routing schemes. They do not provide insights into how the legal traffic in LSNs will be affected. The dynamism of satellites is not considered for a long-term analysis. The analysis results of the bottleneck links being flooded are also not shown.

Network and Distributed System Security (NDSS) Symposium 2025
24-28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-8-3
https://dx.doi.org/10.14722/ndss.2025.230109
www.ndss-symposium.org

To overcome the limitations of existing LFA analysis methods and comprehensively understand the risks and consequences of LFAs that target the time-varying bottleneck links in LSNs, we carry out our study in the following steps.

First, we analyze the temporal and spatial characteristics of the bottleneck links in LSNs. These links are distributed in multiple places and also time-varying. Recent measurements show that the Starlink satellites even have a handover or rescheduling interval of 15 seconds with UTs [22], [23]. Knowing that LSNs have the inherent dynamic movements [24], the bottleneck links might also change from time to time. With more traffic volume from a wide range of regions, there are natural risks if they are flooded by malicious traffic. The legal traffic will be severely degraded or cut off the connectivity from the network. Unlike indiscriminate flooding, targeting these bottleneck links inflicts more severe damage and disrupts a greater number of areas and user traffic.

Second, we propose a new LSN risk analyzer for LSNs called SKYFALL. It is capable of performing comprehensive and in-depth analysis of the time-varying bottleneck links in LSN. It also presents how to analyze the risks of LFAs if compromised UTs are given. Its analytical methodology consists of mainly two stages. In *Data Gathering Stage*, SKYFALL collects the basic information of LSN topology, routing, and traffic distribution. In the following *Analysis Stage*, SKYFALL consequently analyzes the risks on the legal traffic and GSLs and some significant factors affecting the adverse impact (*e.g.,* the number of compromised UTs and their regional blocks). With SKYFALL, we quantitatively analyze the multifaceted consequences of LFAs on LSNs under various attack scenarios. We assess the validation of the bottleneck links, the impact of LFAs on legal user traffic and GSLs, factors influencing the risk impact, and the stealthiness of the compromised UTs. The results demonstrate the global damage influences. The throughput of legal traffic could be reduced by a factor of 3.4 in the worst case.

Lastly, analysis of traditional countermeasures explains why previous methods will not work for LSNs in practice. We then propose possible mitigations against such LFA with fundamental evaluation results. Effective approaches we suggested for operators include traffic scheduling, traffic throttling, differential charging, *etc.*

The contributions of this paper are:

- We explain the definition of bottleneck links in LSNs, and illustrate their characteristics and pervasiveness in a case study with public information of real mega-constellations and GSes in use.
- We propose SKYFALL, a mechanism to analyze the impact on the legal traffic and users when compromise UTs in certain regions are provided to exploit the time-varying bottleneck links. Stages of SKYFALL's assessment cover the data gathering process and analysis process.
- We present comprehensive risk analysis results, with respect to the congested links, actual throughput and degradation, and stealthiness. The results illustrate the effect of utilizing the UTs. We also quantify the possible factors that will influence the consequences brought by the UTs from the perspective of throughput degradation.

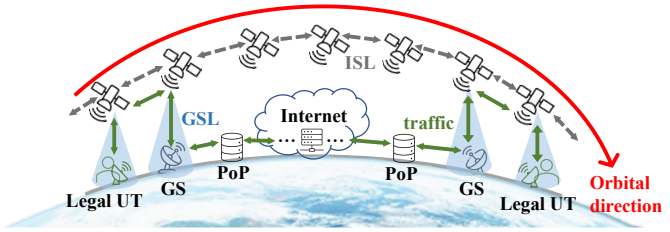

Fig. 1: LEO satellite networks (LSNs).

- We discuss the traditional countermeasures and their limitations in LSNs. Possible defenses are examined along with evaluation results to minimize the reduced user traffic.

The rest of this paper is organized as follows. §II explains the basic structure of the LSNs and LFA risks. The intuition and introduction of the bottleneck links are discussed in §III. The detailed mechanism of the risk analysis is provided in §IV. §V illustrates the analysis results on real-world LEO mega-constellations. The results demonstrate that bottleneck links could be effectively exploited for flooding by compromised UTs to degrade the legal traffic. §VI discusses the traditional countermeasures and their limitations, possible proposals for LSNs. We conclude the paper in §VII. The experiment settings and artifact descriptions are in the appendix.

II. Background and Related Works

A. Emerging LEO Satellite Networks (LSNs)

As shown in Figure 1, LSNs consist of two key components, a space segment containing thousands of LEO satellites, and the rest named the ground segment [25]. The space segment facilitating global connectivity through high-speed ISLs [23]. On the other hand, individual user terminals connected to the satellites may first transmit data [26], [27] by multiple hops of ISLs and then to a GS through a GSL, followed by a Point-of-Presence (PoP) where Carrier Grade NAT is implemented to translate the IP addresses before reaching terrestrial network servers. Terrestrial users access the Internet via an LSN. The data flow must initially be transmitted to a satellite then descend to the ground via a GSL. If the accessed satellite is not connected with a GS, the data will be transmitted through several ISL hops to a satellite with a GSL. For individual users, direct communication through the ISLs in between, without any ground segment for transfer, is infeasible. So it is inevitable for two end users who aim to communicate mutually to connect to the terrestrial Internet first.

One of the salient characteristics of LSNs is the spatial disparity of ground-satellite connectivity. While the space segment is relatively evenly distributed, the GSes and UTs are dispersed unevenly due to various factors, such as geographical reasons (*e.g.,* weather conditions, terrain features) and economic reasons (*e.g.,* population distribution). As a result, the number of accessible satellites per GS varies [25]. The discrepancy of such ground-satellite connections makes the GSes differ in terms of the background traffic bandwidth. The GSL capacity is much smaller than that of ISLs, creating the possibility that some GSLs are bottleneck links.

LSNs also exhibit a remarkable characteristic of high dynamism. Typically, a ground-satellite link in an LSN has a

relatively short duration, typically lasting only a few minutes [24]. As a result of the high-speed movements, the satellites establish connections with new GSes as they move out of range.

B. Link-flooding Attacks (LFA)

In an LFA, a botnet operates by generating data flows between pairs of bots or toward public servers, with the intention of congesting or blocking target network links or nodes. Notably, Coremelt [20] and Crossfire [21] are the most relative attacks that are undistinguishable from legitimate traffic.

Coremelt. A Coremelt attack involves pairs of bots that send continuous traffic to each other, with the traffic paths crossing the same target link(s). Since the traffic can appear to be legitimate, traditional traffic classification methods are ineffective. In Coremelt, attackers must have knowledge of the network topology and routing to select the pairs intentionally.

Crossfire. Crossfire is a similar attack to Coremelt, but instead of compromised botnet nodes sending traffic to each other, it targets public servers. By flooding bottleneck links upstream from these servers, users lose access to them. This technique effectively isolates targets or even an entire region from services without communicating with them. Similarly, traditional traffic classification methods are also insufficient in detecting the malicious traffic generated in Crossfire attacks.

C. Link-flooding Risks in LSNs

Link-flooding risks also exist in emerging LSNs, due to the following unique characteristics. First, an LSN's operations, including the topology and routing, are public knowledge or can be inferred from open sources and measurements. Note that the satellite trajectories including positions are publicly disclosed by NORAD Celestrak [18]. The error in the latitude, longitude, and altitude (LLA) information estimated from the Two-Line-Element (TLE) and ephemeris data is anticipated to be no more than a few kilometers [28]. Thus, full knowledge of the LSN topology is available. For routing, given massive measurements of latency between a large set of bots or ground servers, the attacker is able to make a rational inference about the potential routes or likely connectivity [15]. Apart from the knowledge of LSN topology and routing, the global bot distribution is also available with global coverage provided by current LSN operators. These LSN characteristics make it possible for an attacker to launch an LFA.

Some recent works have begun to explore the possibility of LFAs in LSNs and analyze their risks, such as [15], [16], [29]. All of them employ Coremelt-like flooding in LSNs, where compromised UTs initiate traffic to each other, flooding target links to hinder communications between users or regions. However, methodologies of how to assess the influences under current LSN structures with compromised UTs are not illustrated, such as how the LSN topology and routing could be collected and how the legal traffic volume change could be monitored. Besides, their analysis results do not show the risks on legal traffic if some time-varying bottleneck links are flooded (comparison between flooding bottleneck links and randomly flooding the same number of GSLs in §V).

Takeaways: Until now, there has not been an LFA analysis that properly adapts to the latest LSNs, such that the time-varying bottleneck links are identified for flooding. Previous studies above do not show the methodologies to analyze the potential risks towards legal traffic and GSLs. Quantified results of throughput degradation are not shown for operators to assess the risks. Thus, this motivates us to explore the time-varying bottleneck links and analyze the non-negligible flooding performances with a number of compromised UTs considering the current time-varying LSN structures.

III. Time-varying Bottleneck Links in LSNs

A. Bottleneck Links: The Potential Risks in LSNs

In a communication network, typically a bottleneck link is a link that is fully utilized (saturated) by all flows of traffic sharing this link. Bottleneck links are significant not only because they determine the performance (*e.g.,* the maximum throughput) and scalability of the network, but also because they are prone to be the targets of malicious attacks and involve potential risks (*e.g.,* by LFAs). Understanding the characteristics of bottleneck links, analyzing the consequences of being attacked, and preparing corresponding countermeasures to cope with the potential risks are important for any communication network, including LSNs.

As mentioned in Figure 1, in an LSN, all network traffic between satellites and terrestrial Internet aggregates at GSes and exchanges via GSLs. Since ISLs typically have much higher capacity than GSLs [15], GSLs are likely to be bottleneck links in an LSN. In addition, the LEO mobility involves a unique feature on bottleneck links in an LSN: because the network structure and routes frequently and endlessly change over time, the bottleneck links are time-varying, *i.e.,* the locations of bottleneck links in the LSN frequently change over time. Therefore, in this paper we define bottleneck links as a set of GSLs with high utilization at different times.

B. Characteristics of Time-varying Bottleneck Links

To quantitatively characterize the time-varying bottleneck links in LSNs, we combine real-world LSN traffic distribution and simulation to analyze the following characteristics related to GSLs.

- **Uneven Service Time of GSes.** For GSes, their individual accumulated service time varies. The discrepancy originates from the uneven satellite density and various GS-satellite connection time.
- **Time-varying GS Occurrence.** The uneven distribution of GSes across different locations causes differences in their occurrences on the routes from all the geographical blocks. Temporally, the occurrences of a GS change over time, due to the varying number of satellites connected to it at different times.
- **Time-varying GSL Throughput.** The spatially uneven distributions of populations and GSes make the GSL throughput vary. Temporally, its background legal throughput changes over time, due to the varying satellites connected to a UT or the GS at different times.

Methodology. We carry out experiments by combining LSN simulation, real constellation information, and traffic distribution of operational LSNs. Specifically, our simulation is based on the constellation design of Starlink's first shell including 1584 satellites orbiting at the latitude of 550 km with an inclination of 53 degrees. The satellites are simulated

based on their operational trajectories provided by the open Celestrak database [18]. We simulate the well-known +Grid LSN topology [30], [25], [31], [32], [4] in which each satellite connects to two satellites in the same orbit (*i.e.,* front and back) and two satellites in the adjacent orbits (*i.e.,* left and right). The deployment of ISLs is also substantiated by Starlink, with every transoceanic transmission being accomplished exclusively via ISLs [33]. We simulate 165 Starlink geo-distributed GSes based on their real-world locations [34], [17]. The simulation of satellite routing is based on a recent statement made by SpaceX [33] and a recent measurement study on Starlink [35] which reveals that in over 70% of the situations, user data routes to the nearest GS. Concretely, in our simulation, if a user terminal and a nearby GS are in the transmission range of a certain satellite, then traffic from the user is transferred to the GS via the satellite through transparent forwarding (*i.e.,* the well-known "bent-pipe" routing) without the assistance of ISLs. Only for those remote users far away from available GSes, user traffic is forwarded to an available GS via ISLs.

For legal traffic, due to the lack of an LSN background traffic model, we model user traffic based on the real Starlink traffic distribution in more than 50 countries provided by Cloudflare [14]. Table I summarizes the geographical distribution. Near 50% of the traffic originates from North America, followed by Europe and the Pacific. We discretize the geo-locations on the Earth's surface into geographical grid blocks that are $1°(longitude) \times 1°(latitude)$, with each block representing an area of as large as 12,000 km^2 in the equator. To mimic real-world traffic distribution, we generate background legal traffic based on the distribution from Table I and the traffic generation method introduced in [15]. The probability of adding a background flow of 0.5Mbps from a geographical block is proportional to the percentage of the country it belongs to. We sample this traffic generation process 2,000,000 times. Flows that exceed the link's capacity are discarded. Following the settings in [15], the uplink/downlink capacity of a GSL is set to 4Gbps, while an ISL has a capacity of 20Gbps.

For simulation, it is technically difficult to conduct real analysis on operational LSNs. Therefore, the data-driven simulation is a mainstream methodology widely used in LEO research (*e.g.,* [36], [37], [38], [39]). Similarly, we implement a simulator for the following bottleneck link identification as well as SKYFALL's analysis in around 3000 lines of Python codes. The simulator could be adaptable to analyze the risks under customized constellation settings (*e.g.,* number of orbits, per-orbit satellite number, inclination, and height). To enhance realism, SKYFALL considers the following LSN-related information into simulation: i) the satellites, GSes, UTs, and UTs as network nodes; ii) the constant mobility and movements of satellites, including their geodetic LLA coordinates at the second-level granularity; iii) ground-satellite connections about when and where a GS, compromised UT, and legal terminals ought to be connected to which satellite; iv) the traffic transmission path from legal UTs or compromised UTs to terrestrial Internet; and v) the throughput and capacity per link.

Characterizing Uneven GS Service Time. For a GS, its service time is defined as the total accumulated time with all connectable satellites throughout the day so as to offer services. By prioritizing global GSes in accordance with this service time, we can pinpoint the GSes that may suffer from a long-term flooding attack. The orbital configurations and geographical aspects influence satellite directional shifts in mid to high latitude areas, resulting in a larger satellite density at these locations. Consequently, GSes situated within these zones establish more consistent satellite connections. Figure 2 plots the GS distribution classified according to their service time. The results exhibit substantial geographical variations. A GS having a longer service time suggests a higher activity level, and therefore, takes precedence over other GSes, which could be targeted by the attacker for a long-duration flooding.

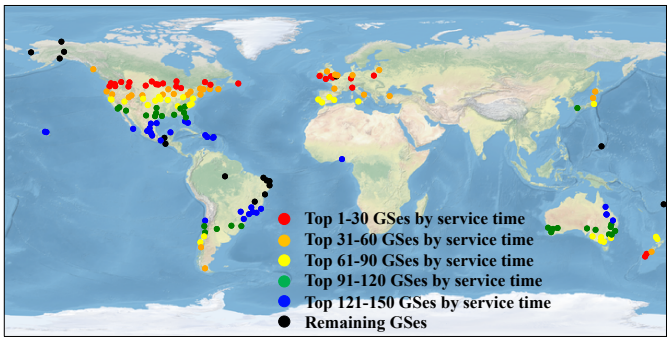

Fig. 2: GS distribution with respect to service time.

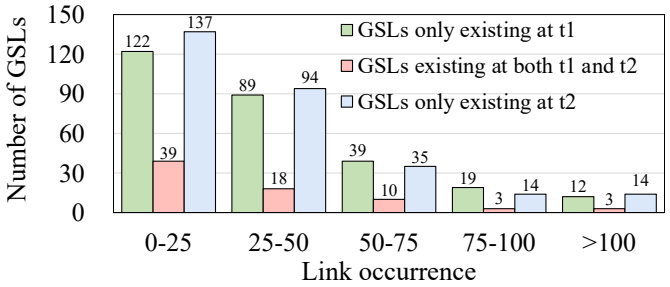

Fig. 3: GSL occurrence distribution in routes from different geographical blocks to the Internet.

Characterizing Time-varying GS Occurrence. For a GS, its occurrence [40] at a time slot is defined as the occurrence by the times of its presence on the routes from all the geographical blocks to the Internet via LSNs. A higher occurrence indicates the GS is responsible for serving more nearby blocks, creating security risks of being attacked. Therefore, by prioritizing GS occurrences for each time slot, we can effectively identify the geographically most influential ones.

Figure 3 illustrates the results by comparing the downlink GSL occurrence at time slot $t1$ and $t2$ respectively (separated by 120 seconds since a GSL usually lasts two to three minutes [24]). The GSLs are categorized based on their link occurrences. Firstly, the link occurrence for each GSL varies significantly at a certain time slot. The majority of them have fewer than 100 occurrences, but some can reach as high as 250. A higher occurrence represents a higher accessibility by geographical blocks and users. Apart from spatial disparity, temporal dynamism is also obvious. Only 73 satellites in total remain connected with GSLs after 120 seconds (the bars colored in red). By comparing the numbers of GSLs between the two time slots within each occurrence range, more GSLs at $t1$ own an occurrence greater than 50, while at $t2$, there are more links

TABLE I: Geographical traffic distribution of Starlink [14].

Country/Region	United States	Canada	Australia	Brazil	Mexico	Germany	Ukraine	Philippines	United Kingdom	France	Others (> 40 countries)
Percentage	36.5%	10.9%	7.4%	5.0%	4.6%	3.6%	3.5%	3.2%	3.0%	2.6%	19.7%

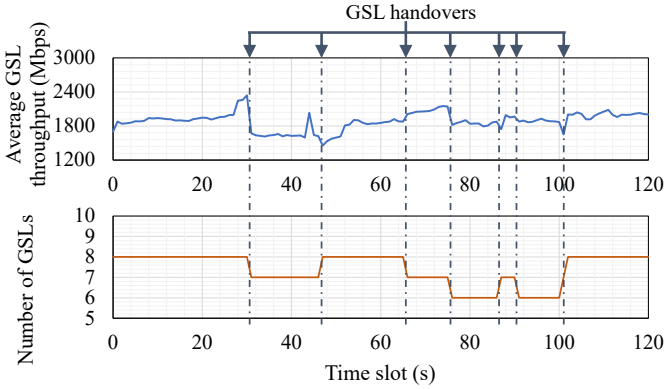

Fig. 4: GSLs and their throughputs are time-varying due to the uneven ground-satellite connectivity and dynamic handovers.

with a lower occurrence. This indicates that a small number of links remain unchanged, while the occurrence distribution undergoes significant changes over time.

Characterizing Time-varying GSL Throughput. As shown in Table I, near 50% of the traffic originates from North America, followed by Europe and the Pacific. The spatially uneven traffic distribution offers us an opportunity to depict the throughput variations among GSLs. Meanwhile, GSL throughput experiences significant temporal variations in an LSN due to frequent ground-satellite handovers. For example, Figure 4 represents the number of GSL connections and the average GSL throughput at the GS located in Ajigaura, Japan, over a period of 120 seconds. The number of GSLs connected to the GS alters within tens of seconds when a satellite orbits away from the GS receiving coverage or it establishes a new connection with an alternative GS. Concurrently, traffic flows initiated from UTs might be relayed by satellites orbiting in different orbital directions owing to the handovers. As a result, the average throughput of the connected GSLs also exhibits a temporal fluctuation. Thus, to identify the GSLs with more traffic, we sort the GSL throughput for each time slot, considering the spatial disparities and temporal variations.

C. Identifying Bottleneck Links in LSNs

We leverage the characteristics introduced above to identify the time-varying bottleneck links in an LSN. Firstly, we are able to identify some vital GSes for each time slot based on the GS service time and GS occurrence. In each time slot, we identify the overlapping GSes from their rankings, for instance, those ranked among the top n GSes in their service time, and concurrently ranked among the top n GSes according to their occurrences at this time slot. We use G_t to mark these overlapping vital GSes at time slot t. Once G_t has been

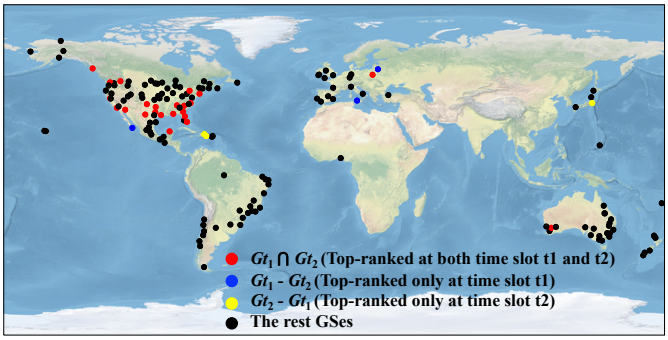

(a) Top-ranked GSes by service time and occurrence.

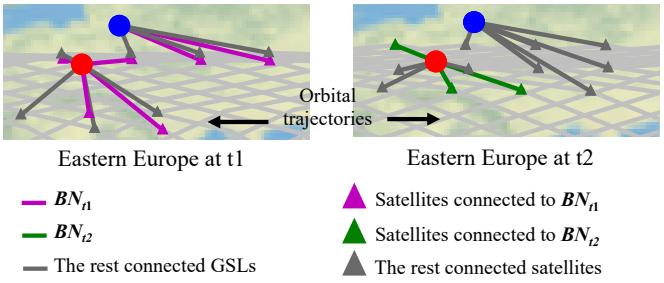

(b) Zooming in on the bottleneck links BN_{t1} and BN_{t2} in Eastern European: focusing on a GS in Poland and another in Lithuania as an example.

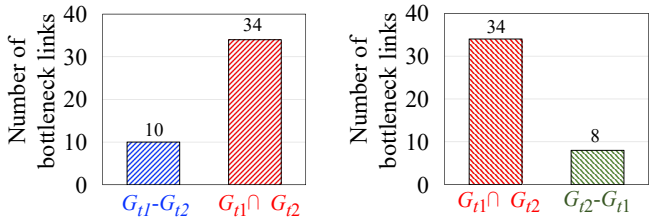

(c) BN_{t1} number and distribution. (d) BN_{t2} number and distribution.

Fig. 5: Bottleneck link identification results.

identified, we prioritize their connected GSLs based on their link throughputs. We are able to prioritize the connected GSLs with the background traffic volume. For instance, we identify GSLs with background throughput that exceeds half the GSL capacity, which we denote as BN_t, representing the bottlenecks at time slot t. Then we define *bottleneck links* $= \cup_{t \in T} BN_t$, where T is the period we want to analyze, consisting of multiple time slots. It is important to note that the *bottleneck links* form a set of vulnerable GSLs for each time slot and involve the dynamics of the LSN. Since the bottleneck links vary over different time slots and are not static, we consider them to be time-varying.

Case Analysis - Bottleneck Link Identification Result. As described above, we are able to locate the overlapping vital G_t,

those who are among the top n GSes from their GS service time and also among the top n GSes from their GS occurrences (*e.g.,* $n = 40$). Figure 5a shows the resulting vital G_{t1}, G_{t2} and the rest GSes at time slot $t1$ and $t2$ (*e.g.,* $t1 = 0$ and $t2 = 120$ with an interval of 120 seconds). The highlighted GSes in red, blue, and yellow are $G_{t1} \cap G_{t2}$, $G_{t1} - G_{t2}$, and $G_{t2} - G_{t1}$. These GSes are more important and are distributed globally in Europe, North America, Asia, Oceania, and South America. However, there is a distinction between G_{t1} and G_{t2}, which confirms that the bottleneck is time-varying. After estimating the link throughput that has more than half of the GSL capacity occupied by background traffic, we are able to identify BN_{t1} and BN_{t2}. To have a better understanding of the time-varying characteristic, we zoom in on the distributions of BN_{t1} and BN_{t2} in Eastern Europe as an example in Figure 5b, categorized by which GS they are connected to at the two time slots respectively. Due to the dynamism, the number of GSLs fluctuates from time to time. Besides, the identified BN_{t1} and BN_{t2} also change at different time slots. The global number of BN_{t1} and BN_{t2} is shown in Figure 5c and Figure 5d, where 44 and 42 links are identified as bottleneck links at $t1$ and $t2$ respectively.

Takeaways: Through this case, we analyze and show the identification results of the bottleneck links, which confirms our intuition that the bottlenecks change due to the frequent handovers between GSes and satellites. By identification of *bottleneck links* $= \cup_{t \in T} BN_t$, where T is the period of multiple time slots, we define bottleneck links as the time-varying target for flooding security analysis by SKYFALL. These findings emphasize not only the importance of the existence of bottleneck links but also the temporally dynamic nature of GSLs.

IV. SKYFALL: AN LFA RISK ANALYZER FOR LSNS

Given the presence of bottleneck links in LSNs, it is intuitive to consider the potential risks of these links being flooded. In this section, we propose SKYFALL, a new LFA risk analyzer for comprehensively analyzing various consequences of LFAs on the time-varying bottleneck links in an LSN. We describe the analysis objective, metrics, and the analysis methodology used by SKYFALL.

A. SKYFALL's Analysis Objective: Impact on the Legal Traffic, GSLs, and Users.

As discussed in previous studies [41], [42], [43], we operate under the assumption that a number of compromising UTs are given in specific regions. The compromised UTs could be used to continuously generate malicious traffic to their connected satellites in the same manner as legal terminals. As some GSLs will carry extra traffic owing to the malicious flooding, they will become congested first. Thus, the analysis objective of SKYFALL is to assess the extent to which the capacity of the links, when flooded by compromised UTs, will congest the legitimate traffic, the GSLs, and impact users in specific geographical blocks. SKYFALL also evaluates the factors influencing the potential risks, such as the number of compromised UTs and their regional blocks.

B. Metrics for Quantifying Consequences

SKYFALL leverages the following metrics to quantify various consequences of LFA risks on time-varying bottleneck links.

- **Throughput Degradation.** The flooding effectiveness could be measured by the throughput degradation [44] or the reduced throughput. Throughput degradation is calculated by normalizing the reduced aggregate throughput of legal background traffic after the flooding. A higher degradation indicates a stronger performance. Considering the time-varying bottleneck links, the worst case is to have a stable and constant throughput degradation using the compromised UTs for each time slot. In other words, for a period T, the overall throughput degradation of BN_t should be promised for $t \in T$. Note that T represents the analysis period during which we evaluate the impact of flooding the bottleneck links, and this period may vary based on our analytical requirements.
- **Ratio of Congested GSLs.** Since the compromised UTs will generate malicious traffic consistently, some transmission links will be congested. The ratio of congested GSLs refers to those that are congested due to the loaded malicious traffic. If the bottleneck links are targeted for the compromised UTs, such a relatively small number of congested downlink GSLs could cause high throughput degradation.
- **Ratio of Affected Blocks.** By flooding the bottleneck links, legal traffic from multiple geographical blocks will be congested. The ratio of affected blocks refers to those whose legal traffic will be congested during its transmission at the downlink bottleneck GSLs. A larger range of affected blocks represents a stronger performance.
- **Stealthiness.** The traffic volume increase of an access satellite caused by the compromised UTs shows how bursty the total traffic is. A higher bandwidth usage brings a higher risk of discovery by the operator.

C. Analysis Methodology

With the above information, SKYFALL has two main stages for conducting an analysis. As shown in Figure 6, in the first stage, SKYFALL acquires essential data about the targeted LSN, including its topology, routing, and legal traffic distribution to have a precise analysis. The information of the targeted LSN could be obtained through available public information, automatic crawlers, reconnaissance, or speculation. A detailed explanation of how to obtain the information is in §IV-C1. With the basic information, SKYFALL then measures network information about the legal traffic distribution. Subsequently, in the second stage, SKYFALL scrutinizes the bottleneck links as described above. SKYFALL then assesses the negative impacts introduced by the given compromised UTs in certain regions, and evaluates how the performance will be affected by factors, such as the number of given compromised UTs, and the malicious traffic volume of each UT.

1) Data Gathering Stage: SKYFALL collects the basic information of LSN topology, routing, and traffic distribution.

Topology and Routing Gathering. For detailed security analysis, SKYFALL must collect basic information about the LSN's topology and routing. Given that the ground-satellite connections might change due to satellite movements, this is a periodic process that can be executed offline.

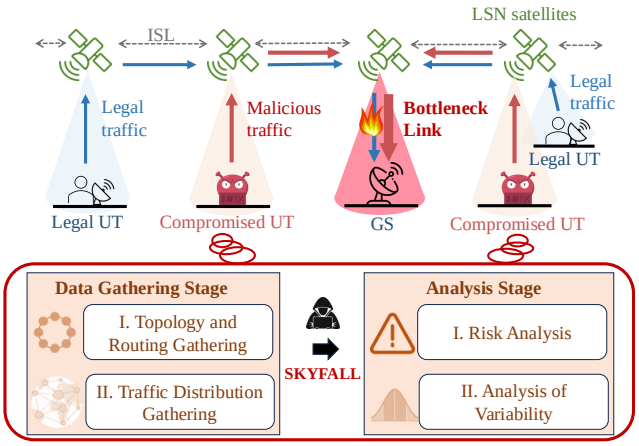

Fig. 6: Data Gathering Stage: SKYFALL obtains information on LSN topology, routing, and legal traffic. Analysis Stage: SKYFALL analyzes the potential risks posed by compromised UTs to legal traffic and users, along with evaluating factors that influence the risks.

Leveraging publicly available constellation data such as Two-line Element sets (TLEs) and Conjunction reports [19], SKYFALL is able to have a knowledge of the LSN topology and speculate routing scheme, including the number of satellites, their orbital altitude, and trajectories. TLEs, a standard format devised by NASA for expressing the trajectories of space objects, are monitored via ground-based telescopes, radars, or onboard satellite sensors. Utilizing TLE data, one can predict a satellite's position and velocity with orbit propagators like the SGP4 package [45]. Conjunction reports [19], on the other hand, are standardized notifications detailing upcoming close conjunction between space objects. TLEs are updated every 3.0 to 34.7 hours, while conjunction reports are updated every 8 hours. Utilizing these sources, researchers in [46] have accurately predicted satellite trajectories and self-driving behaviors. Further, studies such as [22] have effectively hypothesized how a UT might select a satellite for access through extensive simulation and analysis. Thus, SKYFALL is able to obtain the LSN topology and also speculate ground-satellite connections.

Traffic Distribution Gathering. Apart from the basic topology and routing information, the traffic distribution of the LSNs is also essential. To independently estimate the legal traffic volume of a GSL link, SKYFALL could employ a method called *throughput estimation* akin to `iperf`. Specifically, we establish connections between multiple UTs and a ground cloud server, transmitting UDP packets with a linearly increasing throughput. Simultaneously, we employ `pings` to document the latency. Once the latency increases and the accumulated throughput from the multiple UTs fails to increase, the landing GSL encounters congestion. Thus, we record the transmitting throughput at this point. By subtracting the estimated throughput at this point from the GSL capacity (*e.g.,* 4Gbps [15]), the background traffic flow at the GSL can be identified. We do this for each time slot and document the time-varying throughput change. To reduce bandwidth consumption and avoid detection, we only need to generate intermittent burst traffic for throughput estimation, instead of creating continuous high-volume traffic.

2) Analysis Stage: In this stage, SKYFALL analyzes the potential risks based on the information gathered in §IV-C1.

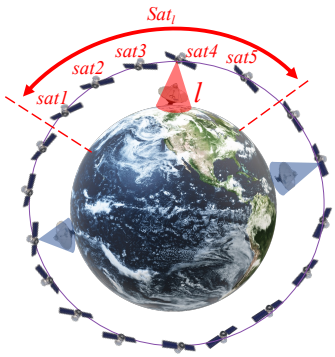

Fig. 7: GSL l and corresponding Sat_l whose accessed malicious traffic flowing into l.

SKYFALL also assesses several significant factors affecting the adverse impact.

Risk Analysis. The first step for SKYFALL is to identify the bottleneck links described in §III. Targeting these links results in more legal traffic being congested by the UTs, as well as more related geographical blocks. Bottleneck links are defined as a set of GSL links in both temporal and spatial dimensions and are identified based on the stepd described in §III. In short, for a time slot, GSes can be ranked by GS service time and GS occurrence. Among the top GSes in both metrics, we identify their connected GSLs with larger throughputs at this slot. Applying this to each slot, all the identified GSLs collected from each time slot of the period are aggregated into the set.

Given the availability of compromised UTs, SKYFALL can analyze congested links and legal flows from a network-wide perspective, as well as affected geographical blocks from the users' view. Note that these compromised UTs at certain regions, are controllable geo-distributed hosts with global access to LSNs. Firstly, since the IP address ranges and corresponding geographical areas of LSNs (*e.g.,* Starlink) are publicly known [47], SKYFALL can verify whether a compromised UT is connected to the target LSN using its IP and geolocation. Furthermore, UTs can be coordinated using a control channel, similar to those in terrestrial networks. With the aid of the Global Navigation Satellite System (GNSS), UTs can be synchronized [48], [49] and then generate malicious traffic towards LSNs. Subsequently, SKYFALL can infer the routes of malicious traffic based on the LSN topology and the ground-satellite connection. For instance, in a topology like that depicted in Figure 7 where traffic typically flows to the nearest GS, the malicious traffic from compromised UTs connected to a satellite in Sat_l will be transmitted by the downlink GSL l. Lastly, based on the estimated throughput (§IV-C1) of link l and the malicious traffic, SKYFALL can determine if l is experiencing congestion and can identify the congested legal flows and their originating blocks. This multi-layered analysis will provide a comprehensive evaluation of the LSNs under potential risks from compromised UTs.

The worst-case scenario could be the compromised UTs being positioned near the bottleneck links identified by SKYFALL. In such situations, these bottleneck links are likely to be the first to experience congestion, leading to a more severe reduction of legal traffic and a greater impact on users.

Analysis of Variability Influencing the Risks. A number of

factors may affect the risk analysis results, such as the number of compromised UTs, the UTs' regions, the unit traffic each compromised UT transmits, and so on. To test this, we change the number of compromised UTs or the region coverage of the UTs and compare the legal throughput degradation and the stealthiness (§V).

To conclude, through such analysis, we can have an in-depth understanding of the reduced throughput of the legal traffic and affected geographical blocks, and determine if the bottleneck links (referenced in §III) can be targeted for flooding to achieve significant effects. LSN operators can thus simulate various LFAs and analyze the potential risks.

V. Analyzing the Consequences

A. Experiment Setup

Constellation and GS Configuration. As previously described in §III-B, the SKYFALL analysis uses the same setting of GS distribution and constellation structure. The LSN comprises 165 GSes and several LEO satellites. Apart from Starlink, we also choose Kuiper for evaluation, compromising 34 orbital planes and 34 satellites for each, with an inclination of 51.9 degrees at the height of 630 km [2], [50].

Topology Setting. To model the network from a broader view and to show SKYFALL's flexibility for arbitrary LSN topologies, we consider the following topology schemes.

- **+Grid Topology.** As used in §III-A, a satellite is connected to the nearest two neighbors within the same orbit, with another two links connected to the adjacent orbits. The topology is shown in Figure 8.. We use **+Grid** for the following evaluations by default.
- **Circular Topology.** Depicted in Figure 8, the satellites within the same orbital plane form a circular structure, where each one communicates with two neighbors. We propose this topology based on the following evidence. Starlink's v1.5 satellites have two intra-orbit inter-satellite links [51]. Since intra-orbit distances are more consistent while inter-orbit connections are more dynamic, it is reasonable to assume that intra-orbit connections between neighbor satellites. Besides, SpaceX has begun to use ISLs for traffic transmission. Experiments are conducted with 10 polar satellites to establish connectivity between the Arctic and the Antarctic in January 2021 [52]. In 2022 and 2023, Starlink effectively provided services in polar regions where no GSes are accessible [53], [54], and the only way for transmission is by ISLs [55], [56]. In conclusion, it is reasonable to assume that the satellites are distributed in orbits, and each satellite only communicates with two intra-orbit neighbors to form a *Circular Structure*.

Threshold Setting. The parameters utilized in the preceding analysis are established on actual measurements and logical assumptions. We set the unit traffic U that a compromised UT can transmit as 20Mbps, based on the uplink measurement outcomes from [26], [27]. The maximum threshold B_u is defined as 400Mbps, implying that the number of UTs connected to a single satellite cannot exceed 20.

The background legal traffic, simulation, and other experiment settings are the same as that in §III-B by default. For the legal traffic of the Kuiper constellation, we follow the traffic generation model in [15] for the comparative experiments below. Apart from the simulator in §III-A, we also conduct experiments based on the container-based emulation tool offered by [37] to demonstrate the feasibility of *throughput estimation* in §IV-C1. The interval between two adjacent time slots is one second, but this can be adjusted for different granularity. A more comprehensive summary of the experiment settings is in the Appendix.

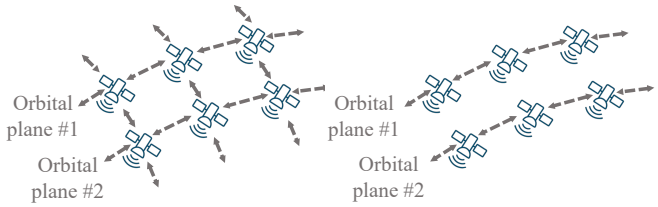

(a) +Grid Topology. (b) Circular Topology.

Fig. 8: +Grid and Circular topologies.

B. Analysis Results

In this subsection, we first evaluate the identification of the bottleneck links. Subsequently, we showcase the effects on the traffic volume and GSLs of LSNs under the influence of these compromised UTs, and how the number of compromised UTs will affect the risk. Overall, our experiments cover the following aspects:

- **Validation of the Bottleneck Link Identification (§V-B1).** We demonstrate the coverage of the geographically affected blocks and the volume of the congested legal traffic by flooding the bottleneck links. We compare the consequences when the bottleneck links are not targeted.
- **Risk Analysis Results: the Adverse Impacts on Traffic Volume and GSLs (§V-B2).** We demonstrate the continuous throughput degradation of legal traffic and congested GSLs with a number of available compromised compromised UTs in a comparative evaluation.
- **Variability Analysis Results by Different Numbers of Compromised UTs and Regional Blocks (§V-B3).** We showcase how the throughput degradation will be influenced by factors such as the number of compromised UTs and their regional blocks.
- **Stealthiness Analysis under Various Topologies (§V-B4).** We evaluate the stealthiness of the UTs by comparing the malicious traffic with the legal traffic transmitted by each satellite.
- **Feasibility Analysis of Throughput Estimation (§V-B5).** We use a container-based emulation tool [37] to evaluate the feasibility of *throughput estimation* technique.

1) Validation of the Bottleneck Link Identification: Utilizing the compromised UTs targeting these meticulously identified bottleneck links would yield considerably larger impacts on the legal users. We demonstrate the feasibility of identified bottleneck links (§III-C) through comparative analysis. We assume SKYFALL congests the identified bottleneck links. We compare SKYFALL by randomly congesting the same number of GSLs (approximately 8% of the global GSLs). We then contrast the congested legal traffic of these GSLs once they are targeted

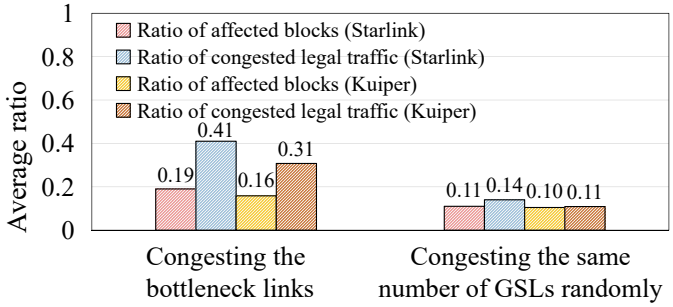

Fig. 9: Validation of the bottleneck links - in comparison with congesting the same number of GSLs randomly. The affected blocks refer to those whose legal traffic will be congested.

by the compromised UTs, and the affected geographical blocks where the background traffic originates. We conduct the same experiments with Kuiper to show SKYFALL's generality.

Figure 9 exhibits both the ratio of affected blocks across global blocks, along with the ratio of congested legal traffic across the total traffic. By congesting merely 8% of global GSLs, over 41% of legal traffic will be congested, originating from approximately 19% of the blocks. By dividing the ratio of congested traffic (affected blocks) and congested GSLs, the gains or payoffs from the perspectives of affected blocks and traffic are as high as 2 × and 5 × respectively. In contrast, merely between 11% and 14% of the traffic would be congested from around 10% blocks if randomly congesting the GSLs. The same extends to other constellations, such as Kuiper. These results imply that bottleneck links indeed yield more substantial adverse impacts and more effective choices for potential attackers. Meanwhile, the bottleneck link identification methods and better performance are not only limited to one specific constellation.

2) Risk Analysis Results: the Adverse Impacts on Traffic Volume and GSLs: After the identification of the bottleneck links (§III), SKYFALL leverages the available compromised compromised UTs (§IV) to evaluate their effects on the LSNs and users and assess the exploitation results of the bottleneck links. To have a more obvious understanding of the risks, we assume the compromised UTs are given at regions near GSes with the identified time-varying bottleneck links. Thus, in this worst-case scenario, since the bottleneck links are supposed to be congested first, more legal traffic will be reduced. Meanwhile, we conduct a comparative experiment with ICARUS [15]. The comparison tries to demonstrate that targeting the time-varying bottleneck links with nearby compromised UTs can result in a continuous and influential adverse effect on legal traffic and users, which is non-negligible.

In this experiment, we compare SKYFALL with ICARUS [15]. ICARUS is an attack scheme of flooding links by mutual malicious traffic between pairs of compromised UTs, but bottleneck identification, the utilization of compromised UT, and performance analysis are not discussed in it. We assume both SKYFALL and ICARUS are given the same number of compromised UTs for flooding. The UTs for SKYFALL are given near the bottleneck links and can intentionally congest the identified bottleneck links while the same number of terminals for ICARUS is put in proportion to the traffic distribution of

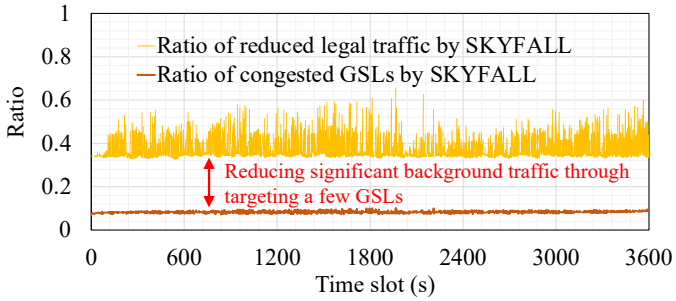

(a) Risks on Starlink's legal traffic and GSLs by SKYFALL.

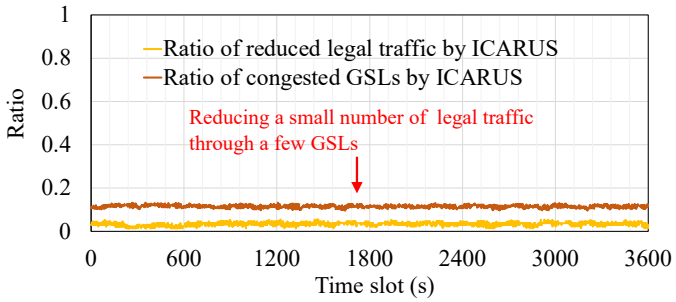

(b) Risks on Starlink's legal traffic and GSLs by ICARUS.

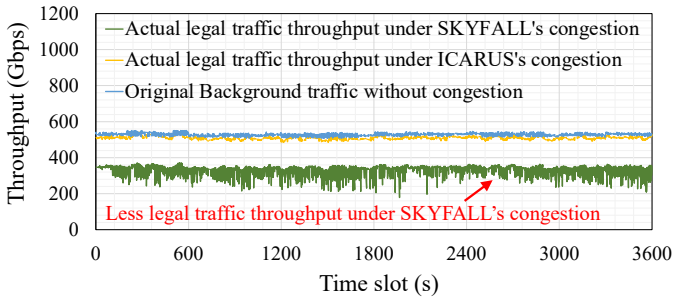

(c) Throughput change on Starlink.

Fig. 10: Risk analysis results.

Starlink [14]. In other words, in places where the traffic is denser, more UTs are available. As no other LFAs in LSNs are currently available, we only compare SKYFALL with ICARUS, using the latter as a representative generalized LFA to show the risks of targeting the bottleneck links. We then assess the reduced and actual throughput of legal traffic over time when the GSLs become congested under both mechanisms. To maintain consistency with our work, we set the unit traffic of the UTs in ICARUS to 20Mbps, and both topologies are +Grid by default.

The ratio of congested GSLs is depicted in Figure 10a and Figure 10b for SKYFALL and ICARUS respectively. Notably, by congesting only 8% of the GSLs on average, SKYFALL successfully reduces the total legal traffic by 37% on average, while 11% for ICARUS throughout the one-hour period. The traffic volume reduced under SKYFALL's scenario is 3.4 × that reduced by ICARUS. The throughput change of the legal traffic is shown in Figure 10c, compared with the original background traffic. Under SKYFALL's application of compromised terminals, the volume of legal traffic is reduced more, compared with ICARUS. Such significant long-term

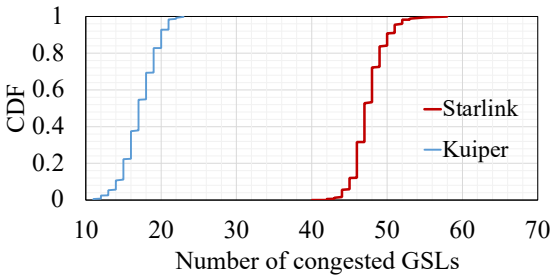

(a) CDF of the number of congested GSLs by SKYFALL.

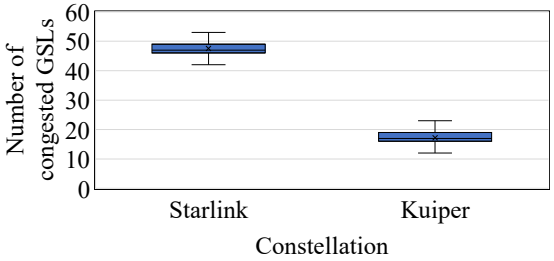

(b) Box-plot of the number of congested GSLs by SKYFALL.

Fig. 11: SKYFALL's adaptable analysis results across various constellations.

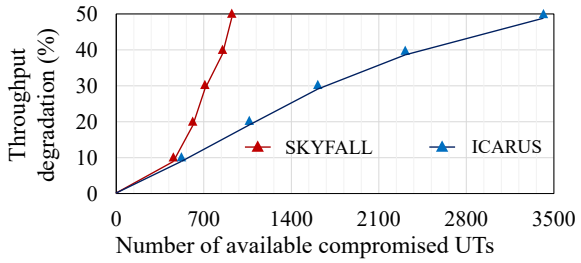

(a) Throughput degradation under the +Grid topology.

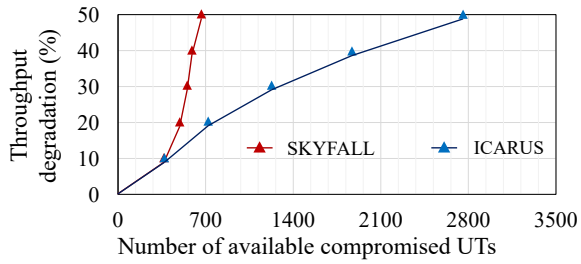

(b) Throughput degradation under the Circular topology.

Fig. 12: Throughput degradation with varying numbers of compromised UTs.

reduced throughput indicates that the small number of congested GSLs brings a great impact on the legal traffic under the proper usage of the UTs. It shows that the UTs are able to effectively congest the bottleneck links and bring a non-negligible throughput reduction throughout the entire period. On the contrary, ICARUS does not consider the variability of the GSLs both temporally and spatially.

To demonstrate the SKYFALL's adaptation across various constellations, Figure 11a and Figure 11b provide a statistical perspective of the number of congested GSLs for each time slot during the entire one-hour analysis period. It is satisfactory that congested link numbers are stable throughout the period, ranging from 10 to 22 for Kuiper, and 40 to 58 for Starlink. Thus, in the worst-case scenario, compromised UTs can have a continuous flooding performance and are also adaptive across different constellation settings.

3) Variability Analysis Results by Different Numbers of Compromised UTs and Regional Blocks: The risk analysis results depend primarily on factors, such as the number and regions of available compromised UTs, which determines the malicious traffic volume.

More compromised UTs correlate with a more pronounced degradation of the legal traffic. To evaluate the extent to which the UTs are able to congest the legal traffic, we vary the number of compromised UTs given to SKYFALL and analyze the according throughput degradation of the legal traffic. We also compare ICARUS when the same number of additional UTs are given. As depicted in Figure 12a and Figure 12b, the throughput degradation for various numbers of available compromised UTs is illustrated. The number of compromised UTs for SKYFALL remains within a small range for both +Grid and Circular topologies. Allocating more UTs for SKYFALL, we can have greater degradation of the congested legitimate traffic [44]. With 1200 compromised UTs available, SKYFALL is able to achieve up to 50% degradation in throughput. In contrast, the achievable degradations for ICARUS are also shown in Figure 12. More than double or triple the number of compromised UTs is required to reach a throughput degradation exceeding 30%. Note that there are currently millions of subscribers owning terminals of Starlink [57]. Given such a small portion of malicious terminals, our analysis shows that with a limited number of compromised UTs, we are able to achieve non-negligible and high-payoff performances.

On the other hand, the geographical regions of available compromised UTs will also significantly influence the feasibility of an exploit. Intuitively, the presence of UTs in a larger coverage of regions near the GSes connected with the bottleneck links leads to greater throughput degradation. To quantify the relationship between the regional coverage and the resulting adverse consequences, we vary the number of regional blocks of the given compromised UTs to study the corresponding changes in throughput degradation. We continue to focus on the worst-case scenario where the regional blocks are near the GSes connected with the identified bottleneck links under both topologies. For comparison, we also vary the number of regional blocks of the compromised UTs for ICARUS and compare their effect on the throughput degradation. Figure 13 illustrates the throughput degradation for various region coverage of the available compromised UTs. It is noticeable that there is a substantial reduction in the number of regional blocks for SKYFALL and ICARUS. Compared with ICARUS, the most significant reduction exceeds 15 $\times$ in the Circular topology with a 50% traffic degradation. In short, if 24 regional blocks are provided with compromised UTs, SKYFALL's analysis shows they achieve a 10% degradation, while this number climbs to no more than 50 for a 50% degradation under the +Grid topology. Overall, more regional blocks are required for +Grid

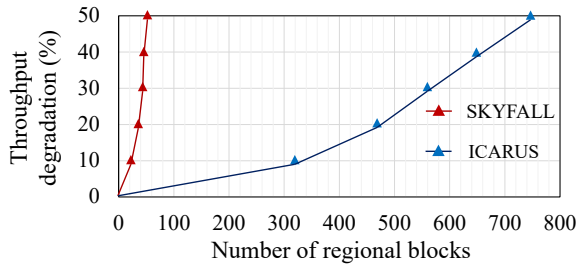

(a) Throughput degradation under the +Grid topology.

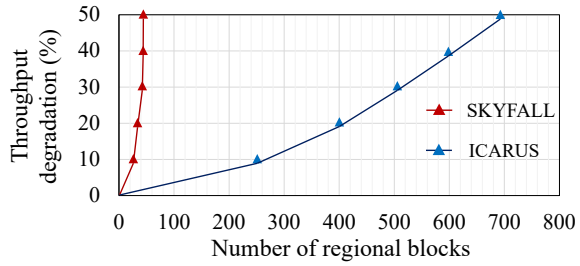

(b) Throughput degradation under the Circular topology.

Fig. 13: Throughput degradation with varying numbers of regional blocks.

than for Circular, although this trend diminishes for SKYFALL. Note that there are currently 7000 or so land blocks that have access to Starlink satellites where UTs are allowed to be put into use. Our analysis results show that compromised UTs at a small regional coverage are capable of reducing the legal traffic volume when targeting the bottleneck links.

In sum, if SKYFALL is given a number of compromised UTs in certain regions, both the number of UTs and the regions will dominate their effect on the LSNs. Admittedly, there are additional factors that need to be put into consideration, such as the malicious traffic volume each UT sends. Yet, through our analysis, the entire compromised UTs could be effectively used to generate malicious traffic persistently towards the bottleneck links.

4) Stealthiness Analysis under Various Topologies: The detectability or stealthiness of flooding is a critical factor in identifying the UTs's success in avoiding detection and mitigating any countermeasures. This factor is measured by the total malicious traffic of a satellite from all accessed compromised UTs. The larger the traffic is, the higher the risk of detection is. Under the same worst-case scenario settings, we present the detectability of each satellite under various throughput degradations in Figure 14. We also include the throughput of legal (background) traffic transmitted by each satellite in the figures (colored in green). Remarkably, the highest detectability is 400Mbps, which takes up only one-eighth of the link capacity. Additionally, more than 90% of the satellites do not have any malicious traffic, and no more than 30 satellites have a malicious traffic volume of more than 200Mbps. Thus, the malicious traffic is insignificant compared to the background legal traffic for a satellite. These findings suggest that if each compromised UT sends a traffic flow of 20Mbps, they have the potential to discretely malicious traffic, making it challenging for operators to detect and mitigate the flooding risks.

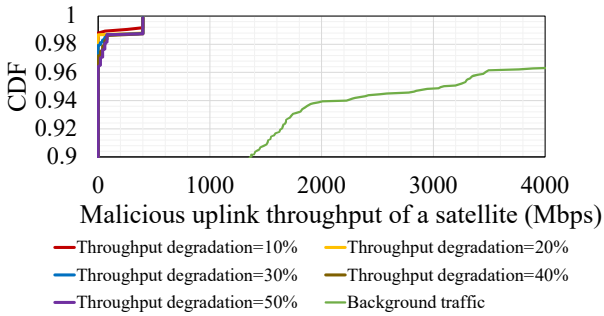

(a) Detectability under the +Grid Topology.

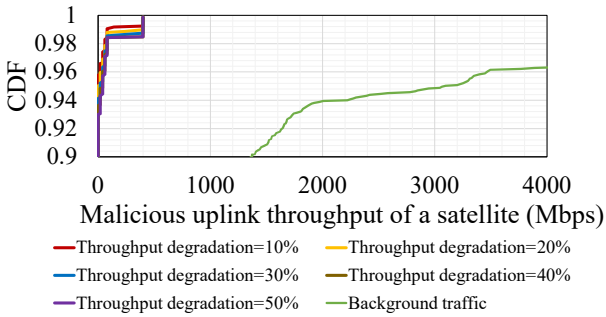

(b) Detectability under the Circular Topology.

Fig. 14: Malicious uplink traffic volume and legal traffic per satellite under various topologies.

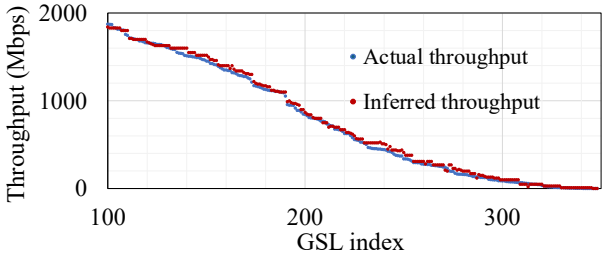

Fig. 15: Comparison between inferred throughput and actual throughput.

5) Feasibility Analysis of Throughput Estimation: To demonstrate the practicality of the proposed *throughput estimation* technique in §IV-C1, we utilize an open-source tool [37] to assess the accuracy of the approach in an emulated environment. The environment is comprised of nodes and GSLs, and runs OSPF protocol on container-based nodes while transferring data traffic between nodes. The generated background traffic (§V-A) for each GSL is also emulated. The *throughput estimation* method is then used to probe the throughput. The actual traffic volume of each GSL, along with the inferred throughput results, are depicted in Figure 15. The results show that the actual and inferred throughputs are well-matched, with a maximum deviation of 192Mbps. On average, the inferred throughput represents 1.1 times the actual throughput. The reason why the inferred throughputs are a little higher than the actual ones is that the critical value of probing bandwidth when the bandwidth no longer increases or the RTT starts to soar is always smaller than the free link bandwidth, so the inferred background traffic is larger, though the difference is ignorable overall. This validates the effectiveness of estimating the link background

throughput, thus creating opportunities for subsequent data gathering and risk analysis of SKYFALL.

VI. COUNTERMEASURES

In this section, we compare various methods for mitigating the GSL congestion based on previous analysis and outline possible approaches.

A. Why Existing Countermeasures are Insufficient?

Filtering. Packet filtering techniques, such as per-packet filtering, are commonly used to create customized blacklists for distinguishing malicious traffic [58], [59]. However, in the case of SKYFALL, Coremelt, and Crossfire, the data is encapsulated by network protocols that cannot be easily differentiated from benign traffic by setting simple filtering rules.

Learning for Classification. Machine learning and deep learning techniques employ models to classify traffic [60], [61]. Similarly, such methods are impractical for the context of our study. Moreover, the classification process requires computational resources, while satellites are typically equipped with low-power processors [36].

Capability. The utilization of cryptographic tokens for sender identification [62], [63] is based on the receiver indicating the desired sender. Thus the receiver only receives data from whom he or she wants. However, this technique is not suitable for our attack objective of disrupting data transmission to public servers, since servers are open to all users. Moreover, in the LSN context, cryptographic tokens would be vulnerable to leakage to the attacker, as compromised UTs are compromised end-users or completely owned by the attacker.

CDN and Cloud. Cloud-based mitigation strategies, such as those proposed in [64], utilize the plentiful memory and network resources of clouds to perform traffic offloading and filtering. However, this may introduce significant latencies, which are non-negligible, compared with the low-latency LSNs. Furthermore, since satellites have limited onboard resources, they cannot serve the same function as clouds.

Routing Protocol Characteristics. Previous studies [65], [66] have utilized loop detection mechanisms in Border Gateway Protocol (BGP) to effectively reroute traffic to an uncongested path nearby. However, since LFA in SKYFALL does not necessarily rely on any specific routing protocol, these protocol vulnerabilities are not relevant to our proposed system.

Topology Obfuscation. Obfuscation can be used to hide the topology from attackers [67]. If the attacker lacks knowledge of the network topology or routing, they are unable to launch successful attacks. However, the satellite topology is publicly available through FCC filings and NORAD [18], [17], making it difficult to conceal geographical information from the public completely. Users are at least able to speculate the satellite geo-locations and their landing PoPs, even measure the network topology using tracing tools.

Packet Marking. Packet marking [68], [69] is a proactive defense strategy against DDoS attacks, especially when spoofed source IP addresses are used and a packet's true origin is hard to detect. In this method, a path fingerprint is embedded within each packet. This mechanism operates deterministically on a per-packet level. Consequently, each packet following the same path carries the same identifier, allowing the victim to utilize the path fingerprint to filter out packets that match the attackers' identifiers on an individual packet basis, irrespective of source IP address spoofing. However, packet marking has high computation overhead and suffers from poor real-time performance in tracing all packets in real-time. The problems become more severe in highly dynamic LSNs.

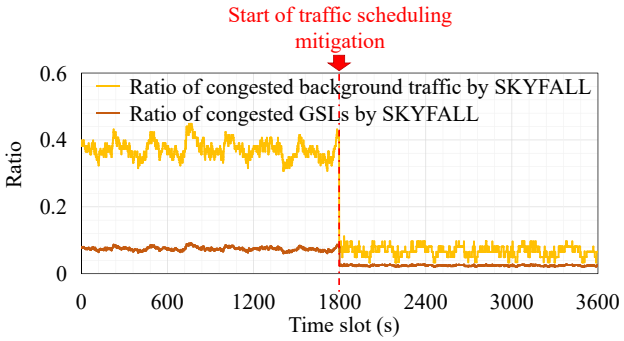

(a) Risks on legal traffic and GSLs with/without the customized traffic scheduling mitigation.

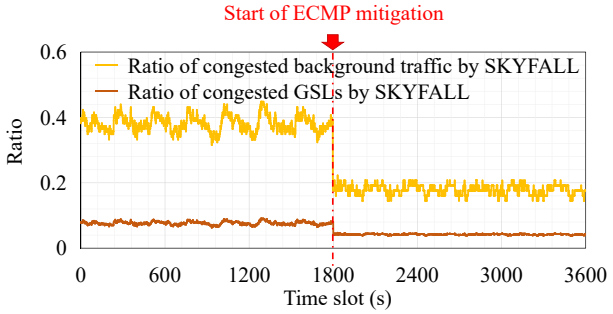

(b) Risks on legal traffic and GSLs with/without the ECMP mitigation.

Fig. 16: Mitigation performance.

B. Possible Countermeasures

Thus, defenses against traditional LFA cannot be used to alleviate SKYFALL' impact. We discuss three possible countermeasures.

Customized Traffic Scheduling for Routing Hiding and Obfuscation. Lack of knowledge of routing decreases the effectiveness of attacks and makes it more expensive for attackers to congest the bottleneck links. Network operators can use traffic scheduling, changing the ISL connections, or other methods to prevent attackers from knowing the network structure. It is detected that Starlink is using a controller for traffic scheduling this year [22]. Thus, a traffic scheduling might be leveraged for mitigation.

Traffic scheduling defenses are mainly classified into two kinds. Distributed approaches perform traffic scheduling based on local link state information in a greedy manner to adjust traffic paths [70]. Centralized traffic scheduling mechanisms compute optimal solutions for the network-wide routing in the controller to determine the routing paths for each node. Routing updates are then sent to each node. In countermeasures against backbone network DDoS [71], [72], border routers

will notify the source AS to perform rerouting based on the Border Gateway Protocol (BGP) to bypass the congested autonomous system (AS). The SDN-based traffic scheduling of ground networks [44], [73], [74] depends on controllers to detect network congestion links and perform centralized routing or flow table updates to prevent Crossfire attacks. Though traffic scheduling might bring extra overhead in LSNs and its feasibility should be verified, it is still necessary to study its defense abilities.

Figure 16a illustrates the results after a traffic scheduling mitigation under the same compromised UT availability in §V-B2. In this scenario, a centralized traffic scheduling approach [75], [44] is employed, wherein an SDN-based ground controller collects real-time global GSL background traffic and recalculates the relay paths for each satellite. Satellites with GSL throughput exceeding a threshold (*e.g.,* 95% of the link capacity) redirect surplus traffic from the congested GSL to the closest satellite with an available GSL. Simultaneously, other satellites continue transmitting their traffic through their connected GSLs. The results show that both the ratio of affected background traffic and congested GSLs drop to one-fourth of those when no mitigation is adopted. This helps the network avoid being congested at certain bottlenecks but disperse the traffic to multiple positions instead. For generality, this proposed countermeasure could potentially be incorporated into relevant Internet traffic engineering standards (*e.g.,* RFC 3272 [76]) or other onboard engineering techniques.

However, this approach still has limitations. The computational expenses are significant since recalculating real-world topologies with more than ten nodes consumes over half an hour [77]. This extended computation time makes it impractical to frequently reschedule traffic. Moreover, due to the dynamic GSL handovers, by the time the recalculation process ends, the network topology may have already changed. Additionally, signaling overhead remains essential as the controller continuously collects GSL states and sends commands to each satellite. The satellites also need to negotiate their routing using source routing (SR) or other mechanisms.

Equal Cost Multiple Path. Equal Cost Multiple Path (ECMP) presents a possible strategy to distribute network traffic across various routes, hopefully preventing congestion on any single targeted link while maintaining a balanced load distribution. In LSNs, while some traffic is transmitted in the 'bent-pipe' approach, where satellites simply forward the received user data directly to the connected GS, others are transmitted through more hops in ISLs. Considering these ISLs in LSNs, we propose the adaptation of ECMP when traffic routes involve ISLs. Through such equal partitioning of traffic over multiple ISL paths, not only is malicious traffic dispersed, but the throughput of legal traffic will experience a significant increase.

Figure 16b shows the mitigation results under the same setting as that in §V-B2. Approximately half of the attacked traffic and GSLs are recovered, as the dispersed malicious traffic no longer congests the bottleneck GSLs. Nonetheless, a number of traffic streams are transmitted in bent-pipe mode, where no ECMP is used, so the mitigation performance of ECMP is not as satisfactory as the traffic scheduling approach in Figure 16a. Additionally, whether ECMP could be practically implemented for each data stream is unknown, since the number of transmission hops and distances are constantly changing.

Traffic Throttling and Differential Charging. It is also possible to address the issue of continuous and high-volume traffic by imposing traffic limitations or differential pricing. It is likely that traffic congestion will occur due to high user density, and operators may limit user traffic to manage this issue. Besides, if the traffic charging is high-cost, the low payoff will force the attacker to give up sending malicious traffic from the root. The growing number of subscribers may necessitate a shift towards charging based on traffic volume usage, so differential pricing is a recommended approach.

C. Discussion with Satellite Network Operators

To the best of our ability, we have sent our analysis results together with a brief explanation to the major operators of today's LSNs [78], through the contact information provided by their official websites. Since these operators have sufficient user terminals (now or in the future) to be potentially impacted by the identified risks, we hope that disclosing our findings can help them become aware of such new risks and take appropriate countermeasures in advance. In addition, we have discussed the identified risks in more detail with our partners from two related satellite network operators in our country. We reached a consensus that the adverse impact of congesting time-varying bottleneck links and degrading the service quality of LSNs is not only a network security risk (*e.g.,* if an attacker can manipulate a large number of compromised UTs in certain regions to generate malicious traffic congesting the bottlenecks), but also a performance/scalability issue that all operators have to face (*e.g.,* if the number of satellite users continues to grow rapidly in the future, even the traffic of legitimate users will cause congestion on the bottleneck link and cause performance degradation). They agree that possible countermeasures like traffic scheduling and re-routing mentioned in §VI-B make sense in principle. Since their constellations are still under heavy construction with insufficient UTs at this moment, countermeasure evaluation in their real experimental network environment will be scheduled as a future work.

VII. Conclusion

This paper introduces the definition of the bottleneck links and how to identify them in real LSNs. We demonstrate the dynamism of the bottleneck links and find that they are susceptible to link-flooding attacks. To study how the time-varying bottleneck links could be exploited for LFA and how to simulate various LFA behaviors, we propose an analyzer SKYFALL. It analyzes the consequences of having compromised UTs to flood the LSNs and helps LSN operators to have an in-depth understanding of the risks. Through comprehensive analysis, our findings shed light on the importance of identifying the bottleneck links. The analysis also demonstrates the impact of the UTs on legal traffic and GSLs. We finally discuss the shortcomings of traditional countermeasures. Possible solutions are proposed to mitigate the potential risks.

Acknowledgments

We thank our anonymous shepherd and all the NDSS reviewers for their helpful comments and feedback. This work was supported by the National Key R&D Program of China (No. 2022YFB3105202) and National Natural Science Foundation of China (NSFC No. 62372259). Zeqi Lai is the corresponding author.

REFERENCES

[1] Starlink: high-speed, low latency broadband Internet. https://www.starlink.com/.

[2] FCC Authorizes Kuiper Satellite Constellation. https://docs.fcc.gov/public/attachments/FCC-20-102A1.pdf.

[3] FCC authorizes boeing broadband satellite constellation. https://www.fcc.gov/document/fcc-authorizes-boeing-broadband-satellite-constellation.

[4] Inigo del Portillo, Bruce G Cameron, and Edward F Crawley. A Technical Comparison of Three Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband. *Acta Astronautica*, 159:123–135, 2019.

[5] Amazon. AWS Ground Station. https://aws.amazon.com/ground-station/.

[6] Starlink Ground Station. https://starlinkinstallationpros.com/starlink-ground-station-backbone-of-satellite-internet/.

[7] 3M Starlink users. https://x.com/Starlink/status/1792678386353213567.

[8] Prime: Cybersecurity risk management strategies for satcom networks. http://www.milsatmagazine.com/story.php?number=1142237172/.

[9] Unmanned x-47b aircraft completes sea trial. http://news.cnet.com/8301-11386_3-57560226-76/unmanned-x-47b-aircraft-completes-sea-trial/.

[10] Yaoying Zhang, Qian Wu, Zeqi Lai, Yangtao Deng, Hewu Li, Yuanjie Li, and Jun Liu. Energy Drain Attack in Satellite Internet Constellations. In *2023 IEEE/ACM 31st International Symposium on Quality of Service (IWQoS)*, pages 1–10. IEEE, 2023.

[11] Threats to united states space capabilities. https://spp.fas.org/eprint/article05.html/.

[12] The hacking of starlink terminals has begun. https://www.wired.co.uk/article/starlink-internet-dish-hack/.

[13] Killnet claims attacks against starlink, whitehouse.gov, and united kingdom websites. https://www.trustwave.com/en-us/resources/blogs/spiderlabs-blog/killnet-claims-attacks-against-starlink-whitehousegov-and-united-kingdom-websites/.

[14] Starlink traffic data. Accessed on: September 9, 2024, https://radar.cloudflare.com/as14593?dateRange=52w.

[15] Giacomo Giuliari, Tommaso Ciussani, Adrian Perrig, and Ankit Singla. ICARUS: Attacking Low Earth Orbit Satellite Networks. In *USENIX Annual Technical Conference (ATC)*, pages 317–331. USENIX, 2021.

[16] Patrick Tser Jern Kon, Diogo Barradas, and Ang Chen. Stargaze: a LEO Constellation Emulator for Security Experimentation. In *Proceedings of the 4th Workshop on CPS & IoT Security and Privacy*, pages 47–53, 2022.

[17] Starlink Services. PETITION OF STARLINK SERVICES, LLC FOR DESIGNATION AS AN ELIGIBLE TELECOMMUNICATIONS CARRIER. https://www.mass.gov/doc/dtc-21-1-starlink-final-order/download, 2021.

[18] CelesTrak. NORAD two-line element sets current data. https://celestrak.com/NORAD/elements/.

[19] CCSDS Recommendation for Space Data System Standards (508.0-B-1): Conjunction Data Messages. https://public.ccsds.org/Pubs/508x0b1e2c2.pdf.

[20] Ahren Studer and Adrian Perrig. The Coremelt Attack. In *14th European Symposium on Research in Computer Security (ESORICS)*, pages 37–52. Springer, 2009.

[21] Min Suk Kang, Soo Bum Lee, and Virgil D Gligor. The Crossfire Attack. In *Symposium on Security and Privacy*, pages 127–141. IEEE, 2013.

[22] Hammas Bin Tanveer, Mike Puchol, Rachee Singh, Antonio Bianchi, and Rishab Nithyanand. Making Sense of Constellations: Methodologies for Understanding Starlink's Scheduling Algorithms. *arXiv preprint arXiv:2307.00402*, 2023.

[23] Nitinder Mohan, Andrew Ferguson, Hendrik Cech, Prakita Rayyan Renatin, Rohan Bose, Mahesh Marina, and Jörg Ott. A Multifaceted Look at Starlink Performance. *arXiv preprint arXiv:2310.09242*, 2023.

[24] Yuanjie Li, Hewu Li, Lixin Liu, Wei Liu, Jiayi Liu, Jianping Wu, Qian Wu, Jun Liu, and Zeqi Lai. "Internet in Space" for Terrestrial Users via Cyber-Physical Convergence. In *Proceedings of the 20th ACM Workshop on Hot Topics in Networks (HotNets)*, pages 163–170. ACM, 2021.

[25] Yaoying Zhang, Qian Wu, Zeqi Lai, and Hewu Li. Enabling Low-latency-capable Satellite-Ground Topology for Emerging LEO Satellite Networks. In *Proceedings of International Conference on Computer Communications (INFOCOM)*, pages 1329–1338. IEEE, 2022.

[26] François Michel, Martino Trevisan, Danilo Giordano, and Olivier Bonaventure. A First Look at Starlink Performance. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC)*, pages 130–136. ACM, 2022.

[27] Daniel Perdices, Gianluca Perna, Martino Trevisan, Danilo Giordano, and Marco Mellia. When satellite is All You Have: Watching the Internet from 550 ms. In *Proceedings of the 22nd ACM Internet Measurement Conference (IMC)*, pages 137–150. ACM, 2022.

[28] TS Kelso et al. Validation of SGP4 and IS-GPS-200D against GPS Precision Ephemerides. In *AAS/AIAA Space Flight Mechanics Conference*, 2007.

[29] Yan Zhang, Yong Wang, Yihua Hu, Zhi Lin, Yadi Zhai, Lei Wang, Qingsong Zhao, Kang Wen, and Linshuang Kang. Security Performance Analysis of LEO Satellite Constellation Networks under DDoS Attack. *Sensors*, 22(19):7286, 2022.

[30] Simon Kassing, Debopam Bhattacherjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. Exploring the "Internet from Space" with Hypatia. In *Proceedings of the 20th ACM Internet Measurement Conference (IMC)*, page 214–229. ACM, 2020.

[31] Debopam Bhattacherjee and Ankit Singla. Network Topology Design at 27,000 km/hour. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies (CoNEXT)*, pages 341–354. ACM, 2019.

[32] Zeqi Lai, Qian Wu, Hewu Li, Mingyang Lv, and Jianping Wu. OrbitCast: Exploiting Mega-Constellations for Low-Latency Earth Observation. In *29th International Conference on Network Protocols (ICNP)*. IEEE, 2021.

[33] Inter-satellite lasers. https://twitter.com/elonmusk/status/1776985584596287622.

[34] Federal Communications Commission. Spacex gen2 non-geostationary satellite system. attachment a. technical information to supplement schedule s. https://fcc.report/IBFS/SAT-LOA-20200526-00055/2378671, 2020.

[35] Liz Izhikevich, Manda Tran, Katherine Izhikevich, Gautam Akiwate, and Zakir Durumeric. Democratizing leo satellite network measurement. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 8(1):1–26, 2024.

[36] Bradley Denby and Brandon Lucia. Orbital Edge Computing: Nanosatellite Constellations as a New Class of Computer System. In *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, page 939–954. ACM, 2020.

[37] Zeqi Lai, Hewu Li, Yangtao Deng, Qian Wu, Jun Liu, Yuanjie Li, Jihao Li, Lixin Liu, Weisen Liu, and Jianping Wu. StarryNet: Empowering Researchers to Evaluate Futuristic Integrated Space and Terrestrial Networks. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 1309–1324, Boston, MA, April 2023. USENIX Association.

[38] David Koisser, Richard Mitev, Marco Chilese, and Ahmad-Reza Sadeghi. Don't shoot the messenger: Localization prevention of satellite internet users. *arXiv preprint arXiv:2307.14879*, 2023.

[39] Eric Jedermann, Martin Strohmeier, Vincent Lenders, and Jens Schmitt. Record: A reception-only region determination attack on leo satellite users. In *USENIX Security Symposium*, 2024.

[40] Min Suk Kang and Virgil D Gligor. Routing Bottlenecks in the Internet: Causes, Exploits, and Countermeasures. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 321–333, 2014.

[41] Tohid Shekari, Alvaro A Cardenas, and Raheem Beyah. {MaDIoT} 2.0: Modern {High-Wattage}{IoT} botnet attacks and defenses. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 3539–3556, 2022.

[42] Saleh Soltan, Prateek Mittal, and H Vincent Poor. {BlackIoT}:{IoT} botnet of high wattage devices can disrupt the power grid. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 15–32, 2018.

[43] Tohid Shekari, Celine Irvene, Alvaro A Cardenas, and Raheem Beyah. Mamiot: Manipulation of energy market leveraging high wattage iot

botnets. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 1338–1356, 2021.

[44] Jiarong Xing, Wenqing Wu, and Ang Chen. Ripple: A Programmable, Decentralized Link-flooding Defense Against Adaptive Adversaries. In *USENIX Security Symposium*. USENIX, 2021.

[45] sgp4 package. https://pypi.org/project/sgp4/.

[46] Yuanjie Li, Hewu Li, Wei Liu, Lixin Liu, Wei Zhao, Yimei Chen, Jianping Wu, Qian Wu, Jun Liu, Zeqi Lai, et al. A networking perspective on starlink's self-driving leo mega-constellation. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, pages 1–16, 2023.

[47] IP addresses and corresponding areas of Starlink. https://pan.uvic.ca/~clarkzjw/starlink/.

[48] Aleksandar Kuzmanovic and Edward W Knightly. Low-rate TCP-targeted Denial of Service Attacks and Counter Strategies. *IEEE/ACM Transactions on Networking*, 14(4):683–696, 2006.

[49] Ryan Rasti, Mukul Murthy, Nicholas Weaver, and Vern Paxson. Temporal Lensing and Its Application in Pulsing Denial-of-service Attacks. In *Symposium on Security and Privacy*, pages 187–198. IEEE, 2015.

[50] Application of Kuiper Systems LLC for Authority to Launch and Operate a Non-Geostationary Satellite Orbit System in Ka-band Frequencies. https://licensing.fcc.gov/myibfs/download.do?attachment_key=1773885.

[51] Two intra-orbit inter-satellite links of Starlink v1.5 satellites. https://mikepuchol.com/modeling-starlink-capacity-843b2387f501.

[52] SpaceX adds laser crosslinks to polar Starlink satellites. https://spacenews.com/spacex-adds-laser-crosslinks-to-polar-starlink-satellites/.

[53] Musk's Polar Starlink Satellites Win Raves at Pentagon While Twitter Flails. https://www.bnnbloomberg.ca/musk-s-polar-starlink-satellites-win-raves-at-pentagon-while-twitter-flails-1.1846012.

[54] Update on Antarctic Communications Improvements. https://future.usap.gov/communications-improvements/.

[55] All future Starlink satellites will have laser crosslinks. https://spacenews.com/all-future-starlink-satellites-will-have-laser-crosslinks/.

[56] Wikipedia. Starlink. https://en.wikipedia.org/wiki/Starlink, 2021.

[57] Spacex starlink internet service surpasses 1m subscribers. https://www.satellitetoday.com/broadband/2022/12/19/spacex-starlink-internet-service-surpasses-1m-subscribers/.

[58] Muthusrinivasan Muthuprasanna and Govindarasu Manimaran. Distributed Divide-and-conquer Techniques for Effective DDoS Attack Defenses. In *International Conference on Distributed Computing Systems (ICDCS)*, pages 93–102. IEEE, 2008.

[59] Yang Xiang, Wanlei Zhou, and Minyi Guo. Flexible Deterministic Packet Marking: An IP Traceback System to Find the Real Source of Attacks. *IEEE Transactions on Parallel and Distributed Systems*, 20(4):567–580, 2008.

[60] Battista Biggio. Machine Learning Under Attack: Vulnerability Exploitation and Security Measures. In *Workshop on Information Hiding and Multimedia Security*, pages 1–2. ACM, 2016.

[61] Nicolas Papernot, Patrick McDaniel, Ian Goodfellow, Somesh Jha, Z Berkay Celik, and Ananthram Swami. Practical Black-box Attacks against Machine Learning. In *Asia Conference on Computer and Communications Security (ASIACCS)*, pages 506–519. ACM, 2017.

[62] Barath Raghavan and Alex C Snoeren. A System for Authenticated Policy-compliant Routing. In *Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 167–178. ACM, 2004.

[63] Xin Liu, Xiaowei Yang, and Yong Xia. Netfence: Preventing Internet Denial of Service from Inside Out. *SIGCOMM Computer Communication Review*, 40(4):255–266, 2010.

[64] Seyed K Fayaz, Yoshiaki Tobioka, Vyas Sekar, and Michael Bailey. Bohatei: Flexible and Elastic DDoS Defense. In *USENIX Security Symposium*, pages 817–832. USENIX, 2015.

[65] Jared M Smith and Max Schuchard. Routing Around Congestion: Defeating DDoS Attacks and Adverse Network Conditions via Reactive BGP Routing. In *Symposium on Security and Privacy*, pages 599–617. IEEE, 2018.

[66] Muoi Tran, Min Suk Kang, Hsu-Chun Hsiao, Wei-Hsuan Chiang, Shu-Po Tung, and Yu-Su Wang. On the Feasibility of Rerouting-based DDoS Defenses. In *Symposium on Security and Privacy (SP)*, pages 1169–1184. IEEE, 2019.

[67] Roland Meier, Petar Tsankov, Vincent Lenders, Laurent Vanbever, and Martin T Vechev. NetHide: Secure and Practical Network Topology Obfuscation. In *USENIX Security Symposium*, pages 693–709. USENIX, 2018.

[68] Abraham Yaar, Adrian Perrig, and Dawn Song. Pi: A path identification mechanism to defend against ddos attacks. In *2003 Symposium on Security and Privacy, 2003.*, pages 93–107. IEEE, 2003.

[69] Minho Sung and Jun Xu. Ip traceback-based intelligent packet filtering: A novel technique for defending against internet ddos attacks. *IEEE Transactions on parallel and Distributed Systems*, 14(9):861–872, 2003.

[70] Antoine B Bagula, Marlene Botha, and Anthony E Krzesinski. Online traffic engineering: the least interference optimization algorithm. In *2004 IEEE International Conference on Communications (IEEE Cat. No. 04CH37577)*, volume 2, pages 1232–1236. IEEE, 2004.

[71] Dimitrios Gkounis, Vasileios Kotronis, Christos Liaskos, and Xenofontas Dimitropoulos. On the interplay of link-flooding attacks and traffic engineering. *ACM SIGCOMM Computer Communication Review*, 46(2):5–11, 2016.

[72] Soo Bum Lee, Min Suk Kang, and Virgil D Gligor. Codef: Collaborative defense against large-scale link-flooding attacks. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*, pages 417–428, 2013.

[73] Albert Gran Alcoz, Martin Strohmeier, Vincent Lenders, and Laurent Vanbever. Aggregate-based congestion control for pulse-wave ddos defense. In *Proceedings of the ACM SIGCOMM 2022 Conference*, pages 693–706, 2022.

[74] Zaoxing Liu, Hun Namkung, Georgios Nikolaidis, Jeongkeun Lee, Changhoon Kim, Xin Jin, Vladimir Braverman, Minlan Yu, and Vyas Sekar. Jaqen: A {High-Performance}{Switch-Native} approach for detecting and mitigating volumetric {DDoS} attacks with programmable switches. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 3829–3846, 2021.

[75] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, et al. B4: Experience with a Globally-deployed Software Defined WAN. *ACM SIGCOMM Computer Communication Review*, 43(4):3–14, 2013.

[76] RFC 3272. https://www.rfc-editor.org/rfc/rfc3272.

[77] Yarin Perry, Felipe Vieira Frujeri, Chaim Hoch, Srikanth Kandula, Ishai Menache, Michael Schapira, and Aviv Tamar. {DOTE}: Rethinking (Predictive){WAN} Traffic Engineering. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 1557–1581, 2023.

[78] Operational satellite Internet constellations. https://en.wikipedia.org/wiki/Satellite_internet_constellation#Operational.

APPENDIX A
EXPERIMENT SETTINGS

The appendix describes the experiment settings in §III and §V.

A. Geo-location Settings

The geo-locations on Earth's surface are discretized into geographical grid blocks measuring 1° longitude by 1° latitude. A block in the equator represents an area of 12,000 km^2.

B. Constellation and GS Configurations

Experiments are conducted using the Starlink Shell One constellation with 1584 satellites at a latitude of 550 km, inclined at 53 degrees [1], [17], as per NORAD Celestrak data [18]. For broader evaluation in §V, the Kuiper constellation is also modeled, consisting of 34 orbital planes each with 34 satellites at an inclination of 51.9 degrees and an altitude of 630 km [2], [50]. We include 165 operational Starlink GSes in the network [34], [17].

C. Topology

Two proposals of topologies are considered. Primarily, the +Grid topology connects satellites within the same orbit and adjacent orbits, which is widely recognized and studied [30], [25], [31], [32], [4]. Additionally, the Circular topology allows communication between each satellite and its two intra-orbit neighbors as shown in Figure 8. By default, we use +Grid topology for evaluation.

D. Routing

Transmission is facilitated via single-hop bent-pipe relays or ISLs [25], [10], [32]. It involves paths to the nearest GS [35] and connections without or with ISLs [23]. This creates a robust framework for routing simulations consistent with current configurations.

E. Background Legal Traffic

Traffic modeling leverages open Starlink data from Cloudflare [14] to simulate real-world traffic conditions. Traffic is generated based on the distribution patterns across geographical blocks. Traffic intensity aligns with the public data proportions, offering a realistic backdrop to assess network performance under typical loads.

F. Other Settings

The experimental parameters are determined based on existing real-world measurements and logical assumptions. Following the settings in [15], the uplink/downlink capacity of a GSL is typically 4Gbps, while an ISL has a capacity of 20Gbps. The traffic unit U that a compromised UT can transmit is set to 20Mbps. The maximum allowable throughput B_u per satellite is capped at 400Mbps to prevent exceeding the capacity of 20 UTs per satellite. The time slot duration for analysis is set at one second to ensure timely processing and responsiveness in dynamic network conditions. These settings are employed across various experiments for consistency and validation of the results.

APPENDIX B
ARTIFACT APPENDIX

The artifact appendix describes the steps to run SKYFALL in an emulated environment.

A. Description & Requirements

This section provides all the necessary details to set up the requirements needed to run the experiments.

1) How to access: The artifact will be available in GitHub repository at https://github.com/SpaceNetLab/SKYFALL and the reproduced dataset can be found at Google Cloud at the following URLs: https://drive.google.com/file/d/1rTuCinLNDnB9q8lyPyZgIaXxpHscX5my/view?usp=drive_link and https://drive.google.com/file/d/1eNZg-OF8xsjjjJNGbJ_8kE_j0x-MtSFR/view?usp=drive_link. To clone, use `git clone https://github.com/SpaceNetLab/SKYFALL`. The README.md provides all the details on running the experiments. The DOI is 10.5281/zenodo.13978898. The target URL is https://doi.org/10.5281/zenodo.13723143.

2) Hardware dependencies: Our primary performance results were obtained and achieved on a DELL R740 server with two Intel Xeon 5222 Processors and 8*32G DDR4 RAM. A server with multiple logical processors is preferred, but a commodity desktop (*e.g.,* an x86-64 CPU with 8 cores and 16 GB of RAM) can also run a demo shown in README.

3) Software dependencies: Our code depends on Python 3.6, and a variety of Python libraries from pip including astronomy packages. A requirements file and installation step have been provided in the README.md file. The operating system could be CentOS 7.9.2009 or Ubuntu 18.04 or above.

B. Artifact Installation & Configuration

Clone the repository using `git clone https://github.com/SpaceNetLab/SKYFALL`. Build packages following the installation instructions provided in the README.md.

C. Experiment Workflow

The workflow includes identifying the vital ground stations (GSes) and bottleneck links, risk analysis for considering flooding circumstances and the analysis of variability influencing the risks. Specifically, our artifacts first build the satellite geo-locations and legal background traffic under the Low-Earth Orbit (LEO) satellite network (LSN) model (steps 1, 2, and 3 in README.md). Then the vital GSes can be identified in step 4. In steps 5 and 6, the usage of the compromised UTs will be analyzed. Step 7 will generate evaluation results based on the intermediate outputs from previous steps.

D. Major Claims

The major claims made by the paper are:

- (C1): Targeting the time-varying bottleneck links with nearby compromised UTs can result in a continuous and adverse effect. This is proven by the results in Figure 10.

- (C2): The number and regions of available compromised UTs will determine the throughput degradation, as shown in Figure 12 and Figure 13.
- (C3): Compromised UTs can have good stealthiness by comparing the total malicious traffic of each satellite from all UTs to the legal traffic. The results in Figure 14 demonstrate the detectability.

E. Evaluation

We offer an example to demonstrate the functionality, configurability, and ease of use of our attribution codes. The results obtained from the example have been used to create the figures in our paper. The structure of our paper results follows that of our example. To make the example applicable to commodity desktops and reduce its size, a demo has also been included in the README.md file. Similar results will be obtained. The artifact README.md provides a more comprehensive set of instructions, with specific commands and command-line arguments for each workflow.

1) Experiment (E1): [30 human minutes + up to 6 compute hours]: The experiment includes building the LSN topology and its legal traffic, identifying the vital GSes and bottleneck links. It also shows how to analyze the risks and detectability given various number or regions of compromised UTs. (Claims C1, C2, and C3).

[Preparation] Please see the sections ”Preparation” and ”Installation” in README.md.

[Execution] Please see the sections ”Getting started”, ”How to reproduce the results in the paper?”, and ”How to run a small demo?” in README.md.

[Results] Please see the section ”Results” in README.md.

F. Customization

In addition to the LEO constellations run in our experiments, the other constellations with various numbers of orbits and per-orbit satellites could also be configured to run all the experiments again. The results may vary, but the overall performance still holds.

The length of the attack period could also be customized. For a full description of all command-line arguments, please see the artifact README.md.