

# Poster: Probabilistic Chunk-Dispersed Routing for Mitigating Link-Flooding Attack in LSN

Hyeon-Min Choi  
Incheon National University  
chm031313@inu.ac.kr

Jae-Hyeon Park  
Incheon National University  
jhp@inu.ac.kr

Eun-Kyu Lee  
Incheon National University  
eklee@inu.ac.kr

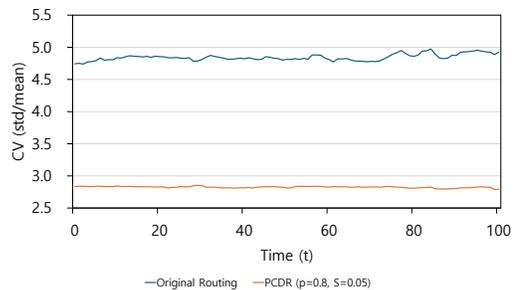
**Abstract**—Low Earth Orbit (LEO) satellite networks (LSNs) offer global low-latency backhaul, but are vulnerable to link flooding attacks (LFAs) because ISL/GSL capacities are limited and shortest-path routing is predictable. We propose Probabilistic Chunk-Dispersed Routing (PCDR), in which the source ground station probabilistically fragments each packet payload into  $n$  chunks and source-routes them in parallel over distinct paths sampled from the ECMP shortest-path set. In our SKYFALL-based simulations on a Starlink shell-1 topology under the SKYFALL and ICARUS attack models, PCDR reduces the coefficient of variation (CV) of ISL utilization and lowers the total loss over 100 s by about 17% (vs. shortest-path routing). Under the SKYFALL attack model, reaching a throughput degradation target of 0.4 requires about 55% more bots and about 57% more compromised regional blocks.

## I. INTRODUCTION

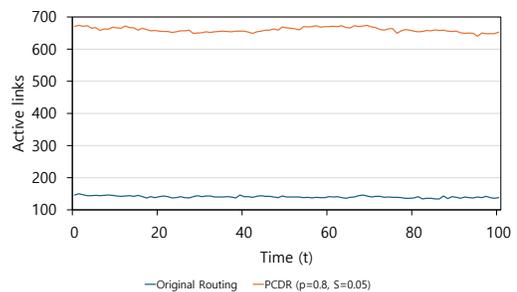
LSNs form a dynamic mesh via inter-satellite laser links (ISLs) and ground-satellite links (GSLs). However, uneven ground-station distribution and lower GSL capacity create time-varying bottlenecks vulnerable to LFAs [1]. The fundamental problem is that deterministic shortest-path routing makes traffic patterns predictable, enabling adversaries to steer malicious flows toward shared bottlenecks using inferable topology information. While ICARUS shows LFA feasibility [2], SKYFALL proves that adaptive tracking of these moving bottlenecks can cripple throughput [1]. This motivates the need for a defense that introduces routing uncertainty to disrupt the attacker’s targeting precision. To be practical, such mitigations must operate in real time with minimal on-board satellite processing.

## II. PROBABILISTIC CHUNK-DISPERSED ROUTING (PCDR)

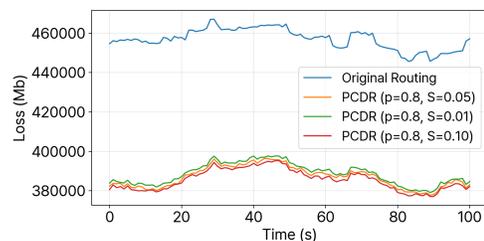
PCDR combines (i) probabilistic chunking (payload fragmentation), (ii) randomized ECMP multi-path assignment, and (iii) randomized downlink selection among the closest satellite and its 1-hop ISL neighbors near the destination ground station (GS). Each chunk is a payload fragment; chunks are forwarded independently and reassembled at the destination GS.



(a) CV of ISL utilization



(b) Active ISLs over time



(c) Loss time series

Fig. 1. PCDR load balancing and loss reduction over 100 s. (a–b) use  $p = 0.8$ ,  $S = 0.05$ . (a) PCDR reduces the CV of ISL utilization, indicating less concentrated load. (b) PCDR increases the number of active (nonzero-traffic) ISLs, implying broader use of underutilized links. (c) PCDR reduces loss over time (varying  $S$ ), lowering the total loss aggregated over 100 s.

The source GS computes the number of available ECMP shortest paths to the destination, denoted as  $P_{\text{avail}}$  (per-slot snapshot). Given payload size  $F$ , minimum chunk size  $S$ , weight parameter  $p$ , and  $u \sim \mathcal{U}(0, 1)$ , the chunk count  $n(F)$  is selected as:

$$n(F) = \min \left( P_{\text{avail}}, \left\lceil \frac{F}{S} (1 + pu) \right\rceil \right), \quad u \sim \mathcal{U}(0, 1) \quad (1)$$

### III. EXPERIMENTS AND RESULTS

We evaluated PCDR in the SKYFALL simulator by comparing it with baseline shortest-path routing [1], [2]. The topology follows the Starlink shell-1 configuration with 165 ground stations. Link capacities are set to 4 Gbps for both GSL uplink/downlink and 20 Gbps for ISLs, and the network state is updated in 1-second time slots. UDP transport is assumed. Each bot transmits 20 Mbps, with up to 20 bots per geographic block. SKYFALL performs adaptive flooding by targeting the identified time-varying bottleneck GSLs in each time slot [1], whereas ICARUS launches concurrent flooding from multiple regions [2]. Metrics include ISL-utilization CV, active ISLs, loss (Mb), and attacker cost (bots/blocks) for a target throughput degradation  $D = 1 - \frac{T_{\text{legal}}}{T_{\text{base}}}$  (no-attack  $T_{\text{base}}$ ).

Fig. 1 summarizes PCDR’s load-dispersion effect. PCDR lowers the CV of ISL utilization and increases the number of active ISLs (Fig. 1(a–b)), indicating broader participation of previously underutilized links. Fig. 1(c) shows the loss (Mb) time series over 100 s; PCDR reduces the total loss aggregated over 100 s by about 17% relative to the baseline.

Fig. 2 summarizes attacker cost. Under SKYFALL, reaching a throughput degradation target of 0.4 requires about 55% more bots and about 57% more compromised regional blocks (in our evaluation). Under ICARUS, the increase is smaller because the attack is inherently more dispersed, but required resources still tend to rise.

### IV. CONCLUSION

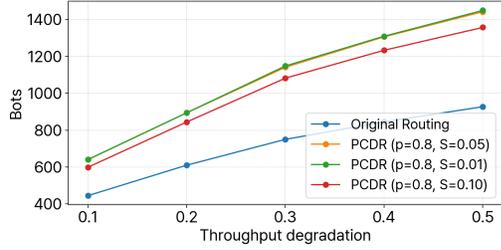
PCDR is a routing-based mitigation that aims to dilute bottleneck concentration rather than explicitly detecting and filtering LFA traffic. It relies on lightweight randomization and ECMP path sampling at the ground station (payload fragmentation + source routing), while satellites perform header-guided forwarding, aligning with on-board resource constraints. The benefits are more limited against attacks that do not focus on specific bottlenecks (e.g., broadly distributed flooding) [2], and without congestion awareness, some chunks may traverse already congested links. Chunking/reassembly introduces overhead, and under TCP, path delay diversity may increase reordering and trigger unnecessary retransmissions. Future work includes QoS-aware probabilistic control using congestion/delay feedback and integrating loss-recovery mechanisms (e.g., FEC or retransmission policies). Overall, PCDR increases attacker cost for bottleneck-tracking LFAs such as SKYFALL [1] while improving load balancing and reducing loss in the evaluated scenarios.

### ACKNOWLEDGMENT

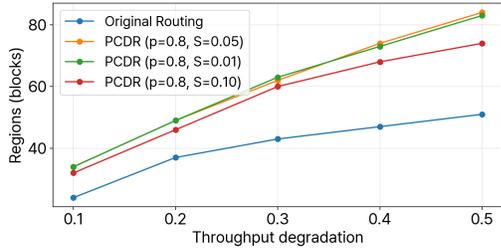
This research was supported by the MSIT, Korea, under the ITRC support program (IITP-2026-RS-2023-00259061) supervised by the IITP.

### REFERENCES

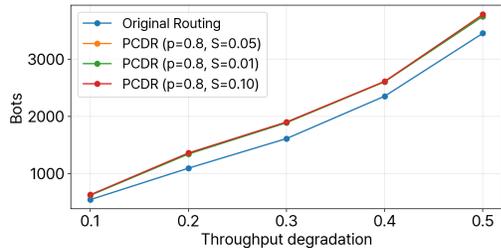
- [1] Y. Deng *et al.*, “Time-varying bottleneck links in LEO satellite networks: Identification, exploits, and countermeasures,” in *Proc. NDSS*, 2025.
- [2] G. Giuliani, T. Ciussani, A. Perrig, and A. Singla, “ICARUS: Attacking low Earth orbit satellite networks,” in *Proc. USENIX Annu. Tech. Conf. (USENIX ATC)*, 2021.



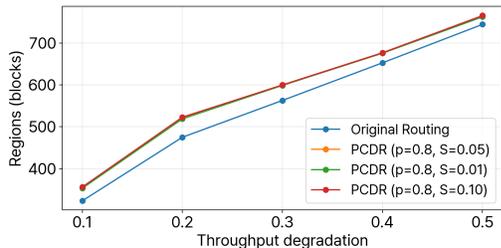
(a) Number of bots required for target degradation (SKYFALL)



(b) Number of regional blocks required for target degradation (SKYFALL)



(c) Number of bots required for target degradation (ICARUS)



(d) Number of regional blocks required for target degradation (ICARUS)

Fig. 2. Comparison of the required attack cost for different routing methods.

The source samples  $n$  distinct paths from the ECMP set (without replacement), attaches a source-route header, and forwards each chunk on a different path in parallel. In the downlink phase, PCDR randomly selects the final downlink satellite among the closest satellite to the destination GS and its 1-hop ISL neighbors (four candidates). This increases route uncertainty and spatial dispersion, reducing the chance of saturating a single bottleneck link.

# PCDR: Probabilistic Chunk-Dispersed Routing for Mitigating Link-Flooding Attack in LSN

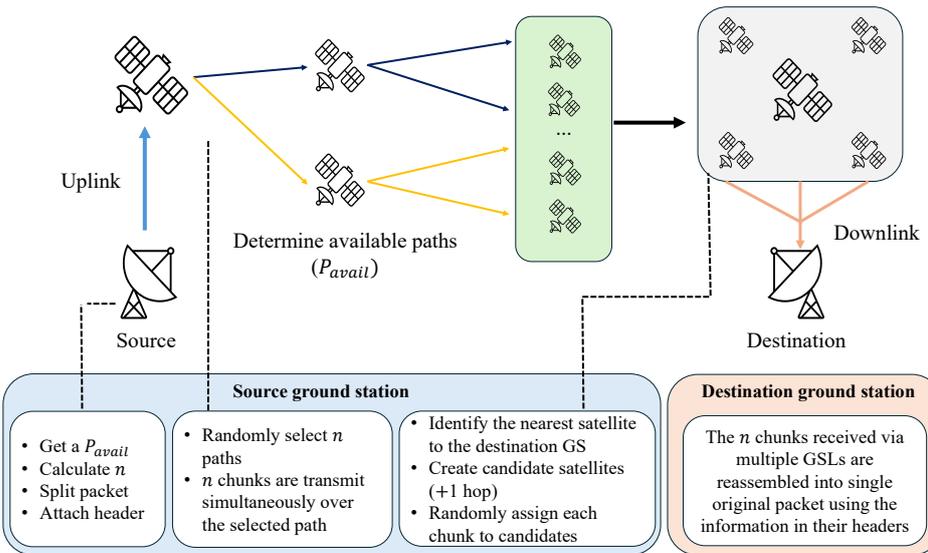
Hyeon-Min Choi, Jae-Hyeon Park, and Eun-Kyu Lee

Dept. of Information and Telecommunication Engineering, Incheon National University, Incheon, Korea

## Abstract

Low Earth Orbit (LEO) satellite networks (LSNs) offer global low-latency backhaul, but are vulnerable to link flooding attacks (LFAs) because ISL/GSL capacities are limited and shortest-path routing is predictable. We propose Probabilistic Chunk-Dispersed Routing (PCDR), in which the source ground station probabilistically fragments each packet payload into  $n$  chunks and source-routes them in parallel over distinct paths sampled from the ECMP shortest-path set. In our SKYFALL-based simulations on a Starlink shell-1 topology under the SKYFALL and ICARUS attack models, PCDR reduces the coefficient of variation (CV) of ISL utilization and lowers the total loss over 100~s by about 17% (vs. shortest-path routing). Under the SKYFALL attack model, reaching a throughput degradation target of 0.4 requires about 55% more bots and about 57% more compromised regional blocks.

## Overview of PCDR



## Key idea

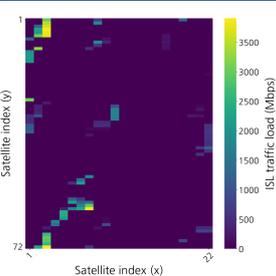
**PCDR spreads traffic over multiple shortest paths using lightweight source-side control, reducing bottleneck concentration and thereby weakening LFA impact.**

- 1. Probabilistic chunking**  
split each packet payload into  $n$  chunks (randomized).
- 2. Randomized ECMP selection**  
source-routes chunks over distinct ECMP shortest paths.
- 3. Randomized downlink**  
choose among the closest satellite and its 1-hop ISL neighbors near the destination GS.

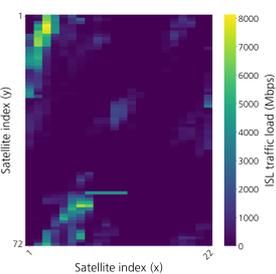
**Probabilistic chunk count equation:**

$$n(F) = \min(P_{avail}, \left\lceil \left\lfloor \frac{F}{S} \right\rfloor \cdot (1 + pu) \right\rceil \right), u \sim U(0,1)$$

$P_{avail}$ : number of available ECMP shortest path  
 $F$ : payload size,  $S$ : minimum chunk size,  
 $p$ : randomness weight,

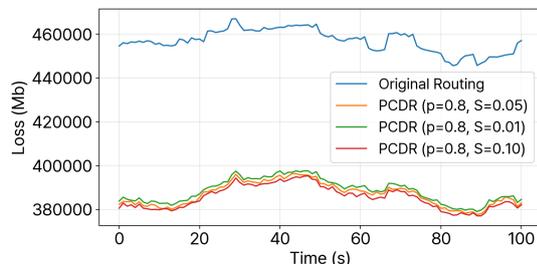


Original routing

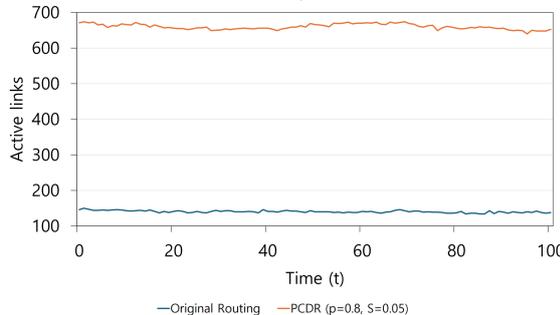


PCDR ( $p=0.8, S=0.05$ )

**Heatmap of traffic loads over all the ISLs at  $t=100s$ .**



Traffic loss timeseries for 100s, PCDR reduces loss over time.



PCDR increases the number of active (nonzero-traffic) ISLs, implying broader use of underutilized links.

**Quantitative analysis of PCDR's load-balancing effect.**

## Result and Conclusion

### Key results

- More balanced ISL load: lower utilization CV, more active ISLs.
- Reliability: 17% lower total loss over 100 s vs. shortest-path.
- SKYFALL (bottleneck-tracking LFA): to reach  $D=0.4$ , needs +55% bots and +57% regions.
- ICARUS (globally dispersed LFA): cost increase is smaller (up to ~20%).

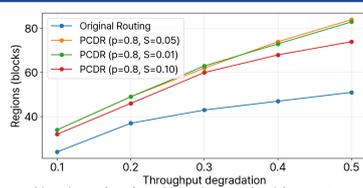
### Why it is practical

- Computation at source GS: probabilistic chunking + ECMP path sampling + source routing.
- Satellites only do header-guided forwarding  $\rightarrow$  fits on-board constraints.
- Uses ECMP shortest paths, so latency inflation is limited.

### Limitations & next steps

- Without congestion awareness, some chunks may still traverse busy links.
- Chunking/reassembly overhead; TCP may suffer from reordering.
- Future: congestion-aware probability control + loss recovery (FEC/retransmission).

**Eval. setup:** Starlink shell-1 (1,584 sats, 165 GS), GSL 4 Gbps, ISL 20 Gbps, 1-s slots, UDP.



**Comparison of the required attack cost for different routing methods.**

