Supplemental material

ICARUS: Attacking low Earth orbit satellite networks

Giacomo Giuliari, Tommaso Ciussani, Adrian Perrig, Ankit Singla

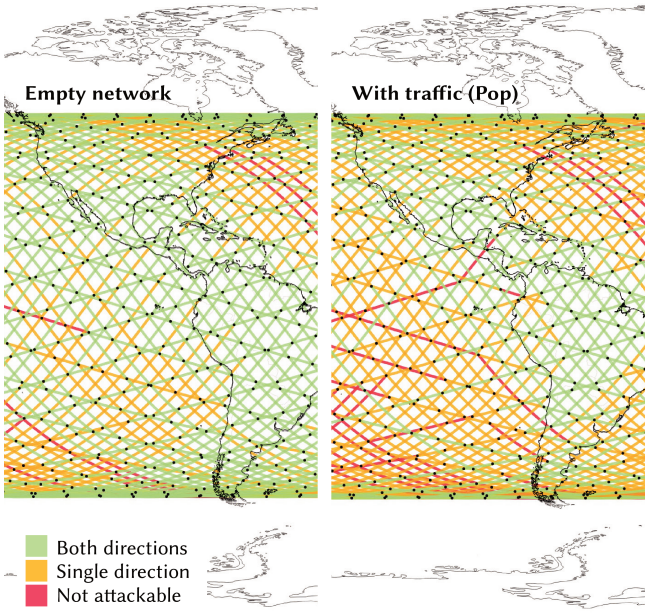

Figure 1: ***Effects of the introduction of traffic on attacks.*** *The map shows the differences between the attacks on an empty network and the attacks with our baseline traffic model. Each link in the map is bi-directional, and can therefore be congested in both directions.*

A Attack location

The following plots concretely show the attacks and their location on the map.

A.1 Effect of baseline traffic on the single-shortest path attack

While adding benign traffic decreases attack cost and lowers maxUp for those target links that are feasible to attack (as seen in §4.3 in the paper), it actually decreases the fraction of feasible target links. However, closer inspection shows that this is not a particularly severe problem for the adversary: the links that become less vulnerable when benign traffic is present are mostly above the oceans, and not the higher-value ones over the more populous regions.

This reduced attack surface is visualized in Fig. 1. In the empty network, most links can be attacked in both directions, and a small fraction of links are not attackable in either direction, while with benign traffic, a large fraction of links can only be attacked in one direction. However, closer inspection shows that this is not a particularly severe problem for the adversary: the links that become less vulnerable when benign traffic is present are mostly above the oceans, and not the higher-value ones over the more populous regions. The links over water are harder to attack, even without benign traffic, because the attack traffic has fewer available routes to traverse them, as there are no up/down-links over water in our model. Adding benign traffic exacerbates this issue because the paths used to reach such difficult-to-attack areas necessarily traverses more congested links, causing attack flows to run into self-congestion before reaching these targets.

A.2 Attack cost for disjoint paths

Fig. 2 shows *where* the cost of the attack is highest across the network's ISLs, with 5-DS. Attacking ISLs that are above land is easier because the adversary can always find a source-destination bot-pair for which there are few, or greatly overlapping paths (reducing uncertainty).

B Multi-target ICARUS is NP-Hard

In §4.4 in the paper we present an attack in which the adversary targets multiple links to congest the communication between two regions. To show that the problem of finding the optimal set of links to congest is NP-hard, we reduce the *minimum set cover problem* [2]—a well-know NP-Hard optimization problem—to it.

Intuitively, the two problems are very similar. The minimum multi-target congestion problem (MMTC) aims at finding the minimum set of attackable links that *collectively congest* all paths. The minimum set cover problem tries to find the minimum number of partially-covering sets that collectively cover a universe of elements. The reduction then consists in mapping elements to paths, and partially-covering sets to attackable links. We now formalize this intuition.

Minimum multi-target congestion problem. Let P be the set of paths the attacker wants to congest. Each path $p \in P$ is composed of (directed) links, $p = \{l_1, \ldots, l_n\}$, which are possibly shared between paths. The union set of all the links in the paths is called L; the adversary can congest—which in this formulation is equal to *removing*—a subset of these, $R \subseteq L$.

The goal of the adversary is to find the minimum set of links A such that removing the links in A disconnects all paths. Equivalently, $\forall p \in P$, $\exists\, l \in R$ s.t. $l \in p$ and $l \in A$, and $|A|$ is minimal.

Minimum set cover problem. Given a set of elements $E = \{e_1, \ldots, e_m\}$, and a collection S of sets of elements (for all $s \in S$ it holds that $s \subseteq E$) for which $\bigcup_{s \in S} s = E$, we want to find a *cover* $C \subseteq S$ such that $\bigcup_{s \in C} s = E$, and $|C|$ is minimal.

Reduction. We show that is always possible to reduce an instance of the minimum set cover problem to an instance of the MMTC problem. The steps are as follows:

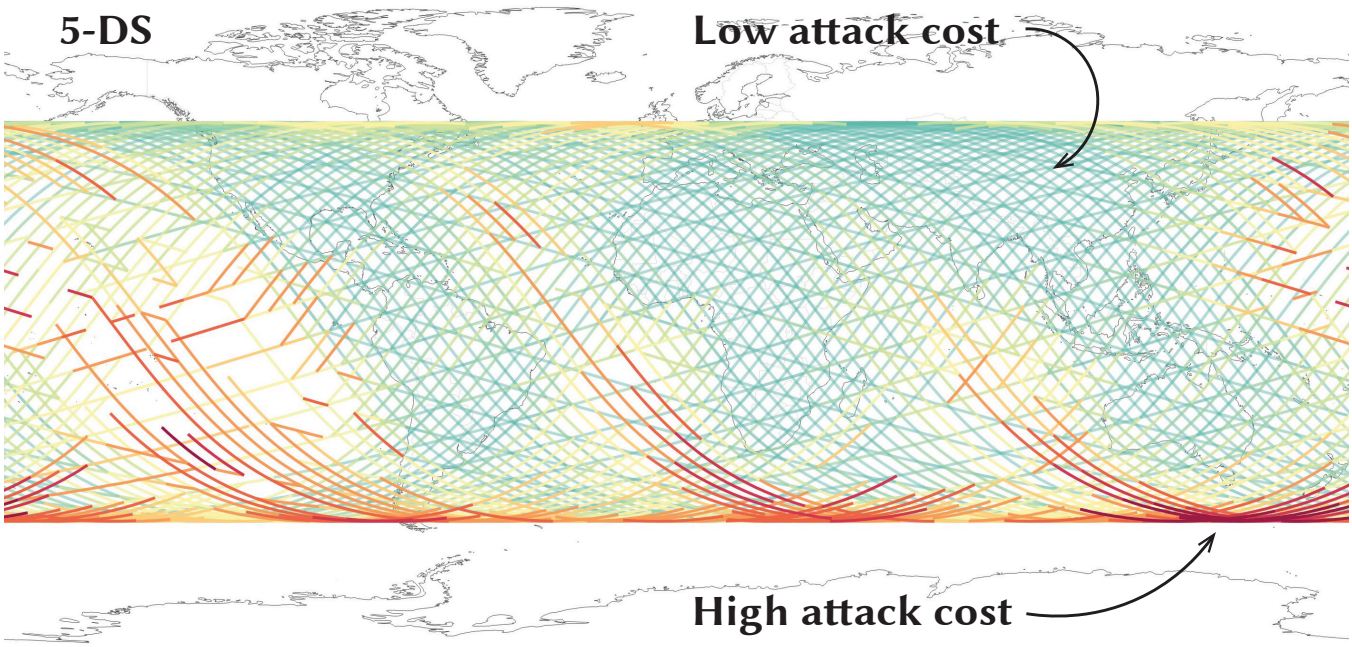

Figure 2: ***Map of attack cost by location.*** *The ISLs with the highest attack cost are the ones far from land—far from where sources are located—as the uncertainty in forwarding is higher.*

- Start with an instance of the minimum set cover problem; S, E are given. We need to create P and L for MMTC.
- For each $s_j \in S$, create a link l_j and add it to L.
- For each $e_i \in E$, create a new path p_i. Add then to p_i all the links $l_j \in L$ such that $e_i \in s_j \subseteq S$. Add the path p_i thus found to P.

We have thus obtained an instance of MMTC in polynomial time (the complexity is at most $|S| \cdot |E|$). If this MMTC instance could be then solved in polynomial time, so would be the related minimum set cover instance, contradicting NP-Hardness. Therefore we conclude that MMTC is NP-Hard.

C Routing schemes

Section 5.1 in the paper introduces the different routing strategies we use in the load-balanced analysis. This appendix describes those schemes in more detail, and underlines the main differences between them. All the metrics in this section refer to paths that share the same source and destination ground points.

Path overlap. We define path overlap as the fraction of common links in a pair of paths that share source and destination ground node. Our four strategies are chosen so that they cover they provide a variety of overlap behaviors, as Fig. 3 shows. A high overlap means less path diversity, which leads to shared links easily becoming a bottleneck and an attack target.

Latency inflation. Path overlap is also an indicator for latency inflation. A low overlap, as for 5-DG and 5-DS, implies a high latency for the longest paths in the source-destination pair (Fig. 4). An intuition on the characteristics of the different load-balancing algorithms is given in Fig. 5.

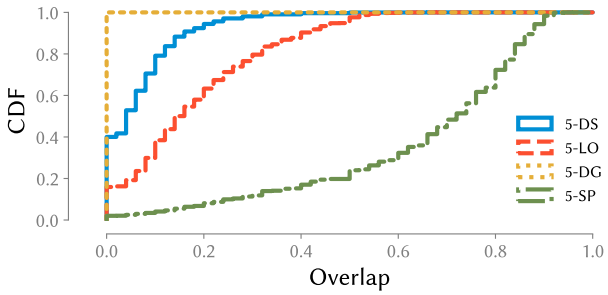

Figure 3: ***Overlap within load-balancing sets.*** *For each pair of paths in a set, we compute the fraction of common links for both paths.*

Number of paths. Fig. 6 shows the number of paths per pair statistics. The low-overlap configurations, with their disjointness constraints, tend to have less source-destination pairs reaching the $k = 5$ mark. This translates into a lower path diversity, which makes attacks easier. The very strong disjointness constraint of 5-DG makes it as an outlier, as the general trend for it is reversed, and the overall number of paths is much lower than its counterparts. This also reflects in Fig. 4, where its median is lower than the others.

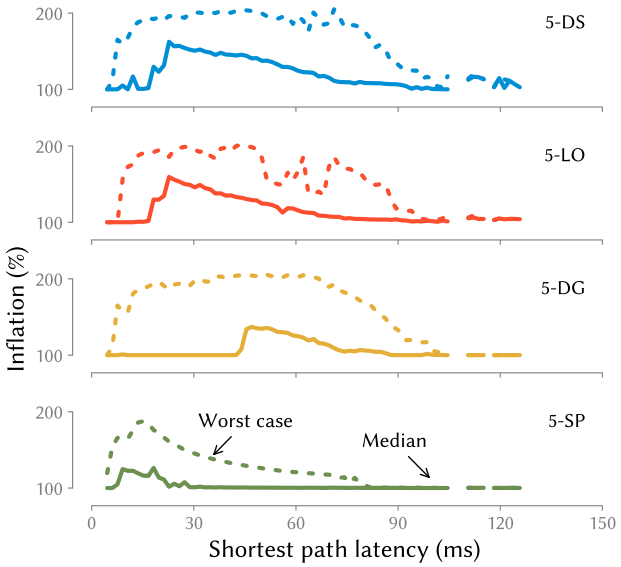

Figure 4: ***Latency inflation induced by non-shortest path routing.*** *For each source-destination pair, we consider the latency ratio between the longest and the shortest path in the load-balancing set. We show here the median and maximum latency, after binning the data in 1ms intervals.*

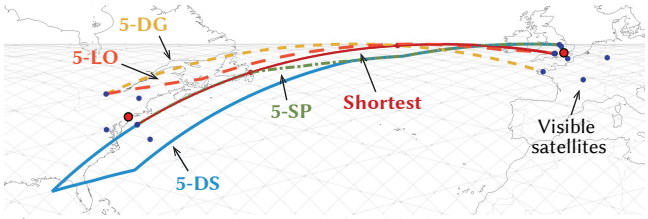

Figure 5: *Longest path for each load-balancing strategy, from New York to London. The longest 5-DG path is shorter than 5-DS. This is due to end-to-end disjointness, and the requirement for each path to be at least as fast as the terrestrial path. A long detour, even if possible, will not be used.*

D Attacking a different constellation

We now present the results for a fictional-but-realistic 40^2 constellation. Our aim is to show that our simulation framework is flexible, and supports running the evaluations in this paper on any constellation configuration. A 40^2 constellation comprises 40 orbits with 40 satellites each, while all other parameters are equal to Starlink: 53° of orbital inclination, 550 km of altitude. This constellation as therefore 16 satellites more than Starlink shell I. It also features a more regular ISL structure, as the lengths of inter- and intra-orbit ISLs are more uniform. The plots in Fig. 7 can be directly mapped to other figures in the corpus of the paper. Generally, we find that results for these parameters are very similar to Starlink shell I. More specific comments can be found in the captions.

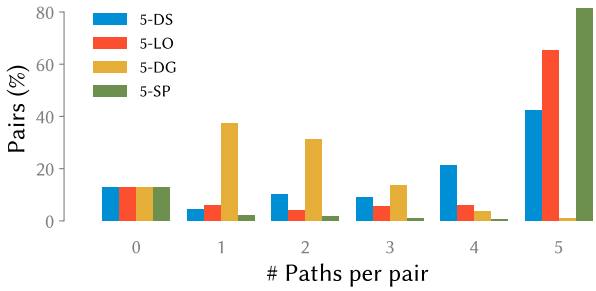

Figure 6: ***Number of paths between pairs.*** *Although our load-balancing set-construction algorithms aim at providing 5 alternative paths between each source and destination pair, other constraints (e.g. disjointness) may reduce the number. For each of the algorithms, 13% of the pairs have 0 available paths. These paths are too close to each other to benefit from the LSN, as the additional latency to reach a satellite and back is higher than the latency of the terrestrial path.*

References

[1] Debopam Bhattacherjee and Ankit Singla. Network topology design at 27,000 km/hour. In *International Conference on Emerging Networking Experiments And Technologies - CoNEXT*, 2019.

[2] Bernhard Korte and Jens Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer-Verlag Berlin Heidelberg, 2012.

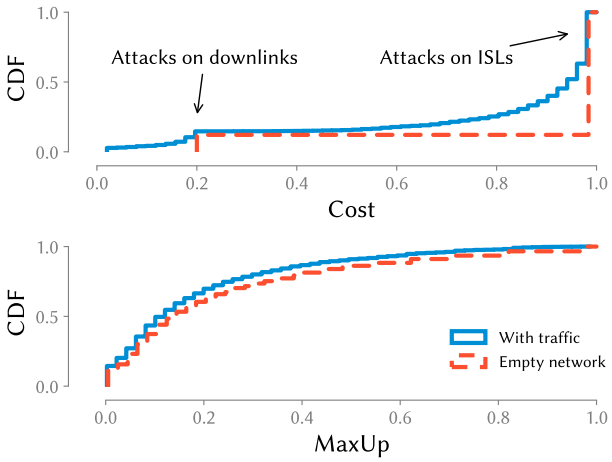

(a) *Cost and maxUp of single-link attacks, GDP traffic model.*

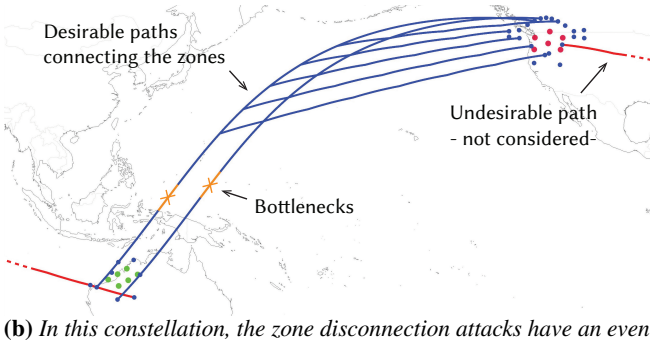

(b) *In this constellation, the zone disconnection attacks have an even more dramatic effect: 2 bottlenecks suffice to congest 24 paths.*

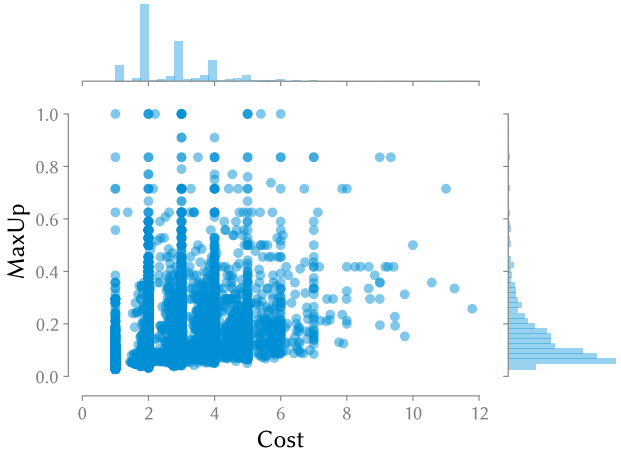

(c) *Zone disconnection attacks cost and maxUp.*

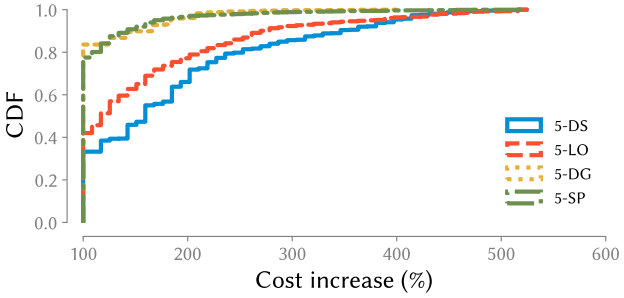

(d) *Cost increase due to load-balancing, with cost optimization. Compared with Fig. 5a in the paper, we see slight increase in median cost.*

Figure 7: ***Attacks on a 40^2 constellation.*** *We run our attack simulations on a fictional constellation [1], with similar characteristics to Starlink shell I. The results are indeed similar to the ones presented in the main body of the paper.*