

Let Me Do It For You: On the Feasibility of Inter-Satellite Friendly Jamming

Ulysse Planta, Julian Rederlechner, Gabriele Marra and Ali Abbasi
CISPA Helmholtz Center for Information Security

Abstract—Unexpected cost reductions in recent years have significantly lowered the hurdle for non-state attackers to gain access to ground stations that can send malicious messages to satellites in low Earth orbit (LEO). In this paper, we investigate the feasibility of on-orbit friendly jamming as a measure to protect LEO satellites from unauthorized and malicious ground station transmissions. We assess the feasibility of integrating an orbital Intrusion Prevention System (IPS) by simulating real-life scenarios. This assessment will lay the groundwork for future experiments involving real satellites. As part of our simulation, we considered a scenario involving a single satellite and an attacker ground station, with a set transmission time. In this scenario, we identified 41 satellites in close proximity that could potentially act as an IPS. Additionally, in a scenario focused on long-term protection, where a single ground station functions as an attacker, we observed that large satellite constellations, such as Starlink, provide substantial protection coverage by acting as an IPS for a specific satellite over the course of a day.

I. INTRODUCTION

With the advent of cheap consumer Software Defined Radios (SDR), Radio Frequency (RF) hardware in general, and the associated open-source projects like GNU Radio [1], the barrier to entry for satellite communications has been lowered considerably. This challenges the assumption that *Attackers* are unable to communicate with satellites in Low Earth Orbit (LEO) and highlights the potential vulnerability of satellites designed with this assumption. It suggests that *Attackers* who were previously unable to contact satellites may now be able to do so. The high reliability requirements caused by the high cost of failure in satellite operations also present a challenge to their implementation. Reliability concerns may also lead to long delays between the discovery of a vulnerability and a patch being ready (assuming the vulnerable system is updatable in the first place). Additionally, at least some satellites are insecure by design and depend on security through obscurity, as indicated by a survey conducted by Willbold et al. [2].

Typically, satellites are equipped with some kind of safe mode. This is a highly restricted operating mode designed to help recovery of the satellite after some incident. Such incidents may include software component corruption or severe hardware failure. Usually, safe mode allows issuing commands via a non-directional antenna at low data rates. This is crucial when a satellite’s attitude control system, which

may not function properly due to a failure, hinders directional communication. The need for absolute reliability makes it hard to prevent *Attackers* from seizing the opportunity to take over control of the satellite. CubeSat missions also usually feature the capability of being controlled via a low data rate unidirectional RF link.

In traditional enterprise networks, firewalls or Intrusion Prevention Systems (IPS) aim to prevent *Attackers* from maliciously interacting with a system. This usually involves a separate IPS/firewall hardware or software node on the network that can intercept, analyze, and selectively forward network traffic. For spacecraft integrating an IPS, operating in this way in safe mode may not be feasible, as the IPS could be a new point of failure for the safe mode functionality.

Furthermore, due to wireless communication, achieving a reliable traffic interception capability for satellites becomes challenging. Traditional methods of intercepting messages over one interface, analyzing them in their entirety, and then forwarding them over another one are rendered impossible. In other words, due to the wireless nature of the connection, inserting a filtering mechanism is difficult.

As demonstrated in other research studies [3], [4], [5], [6], well-intended “friendly” jamming can be used as a means of dropping unauthorized transmissions. To achieve similar functionality as traditional firewalls in a wireless network, Wilhelm et al. [5] propose a wireless firewall that analyzes packets on the fly and then uses jamming to stop packets from being processed by their destination nodes on the wireless network.

The only way to prevent a node from receiving a frame is either to jam all communications to that node and then replay frames that were classified as benign afterward, or to base the decision on a partial frame quickly enough so that it is made before the *Victim* receives the complete frame. In the latter case, the IPS can still jam the signal, preventing the filtered node from receiving the complete frame, and ultimately leading to the corrupted frame being discarded. We propose an IPS based on selective jamming as a way of preventing malicious actors from interacting with vulnerable/insecure-by-design satellites, and make some preliminary evaluations regarding timing limitations.

In this preliminary paper, we describe how an IPS for spacecraft might operate and evaluate whether inter-satellite selective jamming to protect low data rate satellite communications is feasible from a timing perspective, serving as a

preliminary evaluation of the concept for later research.

In summary, we make the following contributions: We propose how selective jamming can be applied to space security, discuss a model to evaluate time constraints, and provide experiments showing the hypothetical selective jamming capabilities of currently active satellites in orbit. To enable others to replicate and build upon our experiments, we will publish the source code, data, and reports of our experiments on GitHub¹.

II. THREAT MODEL

In general, one can consider a ground station or a satellite as an *Attacker*. For the following model description, we will consider an *Attacker* in the form of a malicious ground station that aims to send a malicious frame to a *Victim* satellite. We propose that a *Defender* satellite could function as an IPS in front of the *Victim*, filtering potentially malicious inputs to *Victim* without *Victim* requiring changes to its communications subsystem or software stack. We assume the *Attacker* is targeting a low data rate (around 9600bps), non-directional RF communications system of a *Victim*. This is plausible, as many satellites use a non-directional transceiver for telecommands or as part of their failure recovery scenarios. We argue that it is sufficient to protect this method of communication, as other means of wireless communications may support disabling via commands issued over this system. In case a satellite requires protection, its operators can disable the receivers that would allow *Attackers* to issue commands in a way that the IPS is unaware. The *Defender* would in this case guard the *Victim* until it makes a pass over its legitimate ground station, allowing the legitimate ground station to re-enable other communications systems during the pass and disable them at the end to regain protection from *Defender*.

III. DISTINGUISHING MALICIOUS TRANSMISSIONS

To prevent malicious transmissions from reaching *Victim*, the *Defender* needs a way to distinguish between benign and malicious inputs. As mentioned in Section II, the *Defender* satellite can only interfere with the currently transmitted frame during ongoing transmission. As a result, the decision has to be made based on a partial frame analysis. When there is no Forward Error Correction (FEC) scheme in the physical layer, the jamming decision must be made before a checksum of the complete frame has been transmitted, as they are usually at the end of a transmission. In case FEC is used, the decision has to be made at a time when a sufficient part of the frame can still be jammed for the errors to be uncorrectable. Alternatively, a store-and-forward operation, as suggested by [6], could be implemented. In this approach, not only are all transmissions targeting the *Victim* jammed but also recorded. If they are determined to be benign, they are retransmitted, assuming that full retransmission from the *Defender* is feasible. Detection of malicious transmissions could happen based on the following methods:




Fig. 1. Scenario in which the *Victim* is closer to the *Attacker* than the *Defender*




Fig. 2. Scenario in which the *Defender* is closer to the *Attacker* than the *Victim*

1) *Rule-based detection*: As an initial step, the IPS may filter messages based on rules that describe abnormal data frames that may not be sent during regular operation. The system will decide to jam based on rules similar to the system proposed by Wilhelm et al. [5].

2) *Location-based detection*: An IPS may restrict the geographical access to a satellite by verifying the possibility of a message originating from a known, trusted ground station. The ground station position may be determined through RF direction finding or based on the initial position of the *Defender* at the start of reception.

3) *Timing-based detection*: When it is not possible to inspect the message (as would be the case when *Victim* uses a compromised encryption scheme that the *Attacker* could exploit to create a valid telecommand), the *Defender* can use timing measures to restrict the *Attacker* access to the satellite. This would involve a trusted ground station transmitting an encrypted message to the *Defender*, instructing it not to block all transmissions for a specific time window in which the trusted ground station is transmitting.

4) *Sanitizer-based detection*: In case of checksums being used instead of FEC, the *Defender* may be able to construct an entire message while still being able to jam the checksum at the end of the frame (as is the case for CubeSat Space Protocol (CSP) over AX.25). In this case, the *Defender* can assemble a valid frame before the expiration of the deadline for jamming the *Attacker* signal is over. This is the case because to operate without FEC, the probability of bit errors has to be low. By assuming that the received signal does not contain bit errors, the *Defender* can reassemble the frame while the checksum is still in transmission. This frame can then be forwarded to the

¹<https://github.com/Julian-Rederle/letmedoitforyou>

original network stack running with an Address Sanitizer [7] on the *Defender's* On-Board-Computer. If a memory error is detected, the *Defender* could then jam the signal. This method could provide general protection entirely without requiring active participation of the operator of *Victim* satellite apart from sharing the utilized protocols and network stack. In combination with a store-and-forward strategy [6] approach, this could be adapted to schemes using FEC.

IV. TIMING CONSTRAINTS

Besides considerations about the physical layer, such as a *Defender* satellite's RF hardware limitations and limitations discussed in Section VII, timing constraints play a role in the situations in which the proposed IPS system could function effectively. Wilhelm et al. [8] similarly highlight the reaction time as an important limiting factor in employing jamming as a protection measure. In our scenario, we are considering low bit-rate transmissions with significant distances between nodes. This means that, for a jamming signal to effectively interfere with the reception of a malicious transmission at the *Victim*, it is important that not only the signal reaches the *Victim* with sufficient power, but also that it arrives on time.

The critical timing for this intervention depends on three key distances:

- Distance between the *Attacker* and the *Victim* $D_{A \rightarrow V}$
- Distance between the *Attacker* and the *Defender* $D_{A \rightarrow D}$
- Distance between the *Defender* and the *Victim* $D_{D \rightarrow V}$

For the scheme to work, the *Defender's* signal should reach the *Victim* before the transmission of the *Attacker* has been received, to a point where bit errors will be corrected. Beyond this point, the *Defender* would no longer be able to cause uncorrectable errors that would lead to the frame being discarded.

The time frame available to the *Defender* to act can be calculated based on the time it takes to transmit the remaining part of the message after the last point that could render the message malicious. We will call this part of the message *Tail*. This remaining time is represented by T_{tail} . It is important to also consider the time difference T_{adv} between the arrival of the tail at the *Defender* and the *Attacker*, which can be calculated as follows:

$$T_{adv} = \frac{D_{A \rightarrow V}}{c} - \frac{D_{A \rightarrow D}}{c}$$

That means that the upper bounds of the *Defenders* reaction time is:

$$T_{tail} - \frac{D_{D \rightarrow V}}{c} + T_{adv}$$

The duration of T_{tail} varies based on the protocol used and is influenced by factors such as FEC, postamble, and locations of checksum in the packet. As a simplification, we ignore the relative movement between the satellites during the attacks.

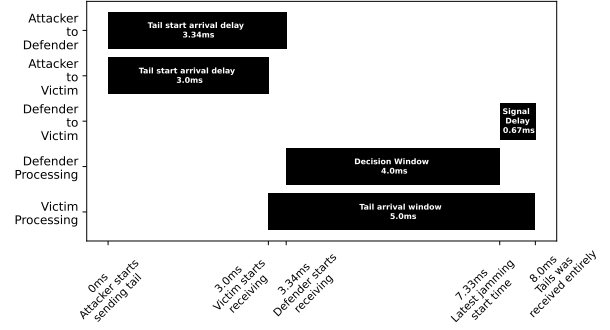


Fig. 3. Timing assuming a bitrate of 9600bps, a *Tail* length of 48 bit and distances as shown in Fig. 1.

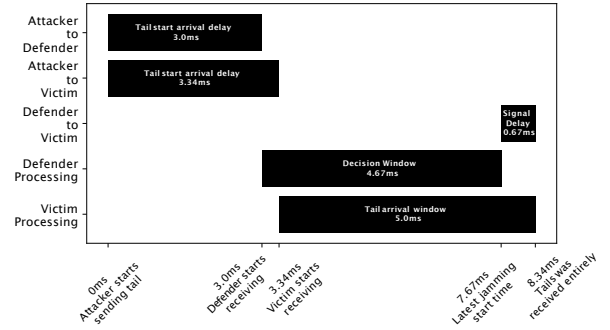


Fig. 4. Timing assuming a bitrate of 9600bps, a *Tail* length of 48 bit and distances as shown in Fig. 2.

V. USE CASES

A. Proximity operation to secure singular satellite

In case of known vulnerabilities in immutable parts of the software of an operational satellite, it may be possible to conduct proximity operations with a CubeSat. This CubeSat could prevent all disallowed interactions with the satellite. The *Defender* could even serve as a proxy for transmitting telecommands. These telecommands would in this case be transmitted to *Defender* in securely encrypted form and then relayed to the *Victim* at low power by the *Defender*. This approach would not only prevent the interception of uplinked data near the ground station, but also limit the *Attacker's* ability to conduct replay attacks by not letting an attacker capture legitimate telecommands. In case the *Defender* experiences a failure in this use case the original satellite could be contacted directly again making the use of *Defender* low risk.

Such capability could be deployed in conjunction with other In-Orbit Services (IOS) e.g., by extending In-Orbit refueling [9] spacecraft (which already need the capability to

approach satellites for their primary missions) with the IPS capabilities. This would enable them to offer single satellite protection as a service on short notice, negating the need for a dedicated IPS satellite.

B. Constellations to cooperatively protect from satellite attacks

With the appearance of modern mega-constellations and the New Space Era [10], there has been an increase in the number of satellites sharing the same software stack. As a result, a single exploit chain targeting one of these constellations could render a significant part of Low Earth Orbit (LEO) satellites compromised allowing for unprecedented attacks. The rise of these constellations also means that for every satellite in LEO there is likely a nearby satellite belonging to one of these constellations. Theoretically, these neighboring satellites could be used to prevent malicious transmissions from reaching other satellites. Utilizing existing constellations may be a way to deploy described protective techniques. Using constellations we could cover a large number of satellites without requiring additional launches. However, this would require collaborative efforts among satellite operators to resolve potential conflicts and ensure optimal protection coverage.

C. On-the-fly honeypot-like behavior

In situations where full coverage of a satellite by IPS satellites cannot be guaranteed, the system may still discourage attacks against targets. If an attack is attempted and prevented by the *Defender* satellite, the *Defender* satellite could respond to a telecommand intended for the protected satellite, effectively assuming the role of a honeypot satellite for the duration of one pass. This strategy could deceive the *Attacker* causing them to disclose their exploits or potential intentions.

VI. EXPERIMENTS

To explore the feasibility of our proposals, we conducted preliminary experiments using Ansys STK [11] in conjunction with the provided Python API. We aimed to identify potential *Defender* candidates by simulating real-life scenarios. The analysis was conducted utilizing a publicly accessible Celestrak dataset [12] of operational satellites in orbit. In all scenarios examined, we assume OPS-SAT (an open research satellite operated by the European Space Agency (ESA)) [13] as the *Victim* satellite and an *Attacker* ground station located near the coast of San Diego. In these preliminary experiments, we investigate feasibility and thus focus only on timing requirements to get a rough estimate of capabilities that can be expected of an inter-satellite IPS. For the experiments, we use a T_{tail} of 5ms which is equivalent to the transmission delay of 48bit (size of checksum of a frame of CubeSat Space Protocol (CSP) over AX.25) at 9600bps. The experiments concentrate solely on a ground station serving as the *Attacker*. However, a similar setup could be used with a satellite acting as the *Attacker*.

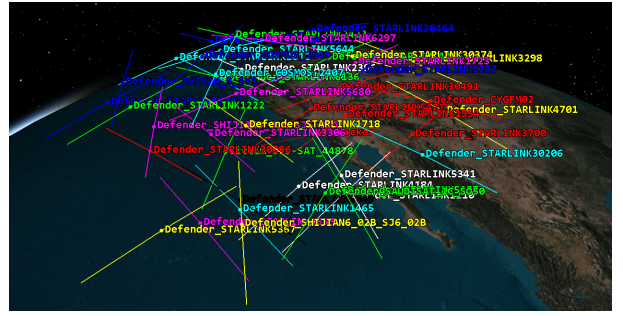


Fig. 5. illustration of potential *Defender* satellites for OPS-SAT and an *Attacker* in San Diego on Jan 1, 2024

A. One-on-one protection

In this experiment, we limit ourselves to finding potential *Defenders* at one specific point in time. With a distance of approximately 1000 kilometers between the *Attacker* and the *Victim*, 03:00:45 a.m. on January 1, 2024 was chosen as the time at which the *Attacker* begins transmitting the *Tail* of the packet.

The procedure is as follows: Each satellite is systematically imported into STK for sequential observation, and access reports are generated for each party involved, calculating the spatial separation between them. These derived values are used in our proposed model to calculate the decision window. The assessment examines whether the *Defender* can make a timely decision to jam or not.

The conducted experiment substantiated the practical applicability of our proposals within a real-world scenario. Notably, it successfully identified 41 satellites nearby, capable of serving as *Defenders* to mitigate potential threats against our designated *Victim*. Figure 5 illustrates the quantity and proximity of capable *Defenders* within the designated context. It is noteworthy to highlight that among the potential *Defender* satellites, 34 of them are part of the Starlink [14] constellation.

B. Protection as a service offered by constellations

In an alternate simulation scenario where the exact initiation time of the *Victim's* transmission of the tail is uncertain, we further explore long-term safeguarding strategies for a designated satellite target. To accomplish this objective, we expand our surveillance scope from a singular defensive satellite at a certain point in time to a succession of *Defenders* ensuring protection over an entire day.

The experiment conducted in this context closely resembles the previous one, with some notable differences. Here, we limit our search for potential *Defenders* to satellites within a specified target constellation. Additionally, we periodically calculate the distance between the involved parties at intervals of 10 seconds for the small and medium constellations and 60 seconds for the large constellations, continuously throughout the entirety of January 1, 2024. To save on computation time, we consider the same fixed *Attacker* ground station position throughout the analysis. We divide our analysis into three subgroups of LEO constellations:

- Small constellations with a satellite count between around 20 to 80, namely Kepler [15], Orbcomm [16], Globalstar [17], Jilin [18], and Iridium [19].
- Medium constellations with a satellite count of around 1000, namely OneWeb [20].
- Large constellations with a satellite count of around 5000, namely Starlink [14].

The results indicate that none of the small constellations can defend our target at any of the observed points in time. For the medium-sized constellation OneWeb, the results are similar, despite its high coverage. The outcome of our analysis is largely influenced by the Walker Star constellation [21] type chosen by OneWeb. In this orbital configuration, the distance between ascending and descending planes increases towards the equator and decreases towards the poles. As a result, near the equator, it is less likely to find a OneWeb satellite close to both a *Victim* and *Attacker*. On the other hand, toward the poles, the conditions are more favorable in terms of distance requirements for a reasonable decision window. Starlink, being the largest constellation, demonstrated a noteworthy presence, with at least one of its satellites within a reasonable defensive range to our designated *Victim* in 79 percent of the 42 time-slots with the *Attacker* in range. In addition to the previous test using a T_{tail} of 5ms, we tested a T_{tail} of 8ms, which increased *Defender* coverage to 95 percent and left only two undefended time slots. These two undefended time slots are caused by especially low T_{adv} , meaning that the propagation delay of the *Attacker's* signal to the *Defender* is significantly higher than the propagation delay to the *Victim*.

VII. LIMITATIONS

In addition to the timing constraints discussed in Section IV there are other non-timing related limitations potentially inhibiting our proposed methods of defense. These can be categorized as follows:

- Computational limitations can render the *Defender* unable to decide to jam in time for it to be effective. In case the differentiation between malicious and benign is based on a copy of the actual protocols stack running with an address sanitizer, the overhead of the address sanitizer (typically 2x computationally and 3x for stack size [7]) has to be considered. If the capability to store-and-forward operation, as suggested by [6] is present, the overhead will just lead to lower overall throughput.
- An attack spread over multiple frames limiting the detection capabilities of the *Defender*. This can be defended against by considering multiple messages to inspect datagrams. Unusually long or improperly fragmented messages might be caught by a rule-based approach.
- The use of directional antennas by *Attacker* may or may not be of significant relevance when determining if the idea is feasible in a certain scenario. When assessing an *Attacker* situated at a considerable distance from the *Defender*, the main lobe of the *Attacker's* antenna grows to a size where it likely encompasses the *Defender*. Conversely, in scenarios involving short-distance attacks

(e.g., Satellite-to-Satellite attacks), the *Defender* may be able to recover signals from a side lobe of the *Attacker's* antenna.

- Doppler shift of the *Attacker's* signal has to be accounted for by the *Defender*. How exactly this is performed is dependent on the exact scenario where the *Defender* is deployed and the RF hardware of the *Defender*. When putting some limitations on possible *Attackers*, the Doppler shift of the *Attacker's* signal can be anticipated.
- False negatives caused by bit errors in the *Defender's* reception for (some methods of distinguishing malicious transmissions). In case the bit error rate is too high, other methods of distinguishing malicious traffic that don't rely on the frame content can be used.

It should be noted that not every satellite we consider as a *Defender* in our experiment has the RF hardware capabilities to jam. Generalizing antenna requirements and limitations in our model poses a significant challenge, particularly due to the dependency on the chosen physical layer, which is why we limit ourselves to investigating time constraints in this preliminary work.

VIII. CONCLUSION

As the barrier for *Attackers* to acquire satellite communication equipment has been significantly lowered with the advance of cheaper RF hardware, there is an increasing need to ensure the security of satellite operations through the implementation of technologies such as selective inter-satellite jamming. By simulating real-life scenarios, the study establishes the theoretical feasibility of our approach. In future work, we will conduct a more practical evaluation of the proposals and determine their real-world applicability. It is also plausible to extend the long-term constellation protection approach to encompass a comprehensive set of all publicly accessible ground stations for a more exhaustive evaluation of constellation protection coverage, which could also uncover potential coverage blind spots. Furthermore, it would be interesting to carry out experiments focusing on the feasibility of the actual antenna requirements and, in particular, concerning the RF hardware onboard active satellites.

ACKNOWLEDGMENT

We thank the Saarbrücken Graduate School of Computer Science for their support.

REFERENCES

- [1] E. Blossom, "Gnu radio: tools for exploring the radio frequency spectrum," *Linux journal*, vol. 2004, no. 122, p. 4, 2004.
- [2] J. Willbold, M. Schloegel, M. Vögele, M. Gerhardt, T. Holz, and A. Abbasi, "Space odyssey: An experimental software security analysis of satellites," in *IEEE Symposium on Security and Privacy*, 2023.
- [3] I. Martinovic, P. Pichota, and J. B. Schmitt, "Jamming for good: A fresh approach to authentic communication in wsns," in *Proceedings of the Second ACM Conference on Wireless Network Security*, ser. WiSec '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 161–168. [Online]. Available: <https://doi.org/10.1145/1514274.1514298>
- [4] A. Bachorek, I. Martinovic, and J. B. Schmitt, "Enabling authentic transmissions in wsns — turning jamming against the attacker," in *2008 4th Workshop on Secure Network Protocols*, 2008, pp. 21–26.

- [5] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Wifire: A firewall for wireless networks," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, p. 456–457, aug 2011. [Online]. Available: <https://doi.org/10.1145/2043164.2018518>
- [6] M. Wilhelm, "Feasibility and applications of a wireless firewall," doctoralthesis, Technische Universität Kaiserslautern, 2016. [Online]. Available: <https://nbn-resolving.de/urn:nbn:de:hbz:386-kluedo-43455>
- [7] "Clang 19.0.0git documentation - addresssanitizer," <https://clang.llvm.org/docs/AddressSanitizer.html>, accessed: 2024-02-08.
- [8] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Short paper: Reactive jamming in wireless networks: How realistic is the threat?" in *Proceedings of the Fourth ACM Conference on Wireless Network Security*, ser. WiSec '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 47–52. [Online]. Available: <https://doi.org/10.1145/1998412.1998422>
- [9] "Ios mission and maturation phase proposal - in-orbit refuelling assessment," <https://nebula.esa.int/content/ios-mission-and-maturation-phase-proposal-%E2%80%93-orbit-refuelling-assessment>, accessed: 2024-01-03.
- [10] O. Kodheli, E. Lagunas, N. Maturo, S. K. Sharma, B. Shankar, J. F. M. Montoya, J. C. M. Duncan, D. Spano, S. Chatzinotas, S. Kisseleff, J. Querol, L. Lei, T. X. Vu, and G. Goussetis, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2021.
- [11] "Ansys stk: Software for digital mission engineering and systems analysis," <https://www.ansys.com/products/missions/ansys-stk>, accessed: 2024-01-03.
- [12] "Celestrak," <https://celestrak.org/NORAD/elements/index.php>, accessed: 2024-01-03.
- [13] "Esa - ops-sat," https://www.esa.int/Enabling_Support/Operations/OPS-SAT, accessed: 2024-02-08.
- [14] "Starlink - world's most advanced broadband satellite internet," <https://www.starlink.com/technology>, accessed: 2024-04-30.
- [15] C. Günther, "Kepler – satellite navigation system description and validation," in *Navitec 2018 Signal Workshop*, 2018. [Online]. Available: <https://elib.dlr.de/126631/>
- [16] "Orbcomm og2," <https://www.orbcomm.com/en/partners/connectivity/satellite/og2>, accessed: 2024-04-30.
- [17] "Satellite technology powered by the globalstar satellite network," <https://www.globalstar.com/en-us/about/our-technology>, accessed: 2024-04-30.
- [18] "Jilin constellation," <https://www.eoportal.org/satellite-missions/jilin-con>, accessed: 2024-04-30.
- [19] K. Maine, C. Devieux, and P. Swan, "Overview of iridium satellite network," in *Proceedings of WESCON'95*, 1995, pp. 483–.
- [20] "Constellations connecting people all over the globe," <https://www.airbus.com/en/space/telecom/constellations>, accessed: 2024-04-30.
- [21] E. Lagunas, S. Chatzinotas, K. An, and B. F. Beidas, *Non-Geostationary Satellite Communications Systems*. Institution of Engineering and Technology, Dec. 2022. [Online]. Available: <http://dx.doi.org/10.1049/PBTE105E>