Death By A Thousand COTS: Disrupting Satellite Communications using Low Earth Orbit Constellations

Frederick Rawlins
University of Oxford

Richard Baker
University of Oxford

Ivan Martinovic
University of Oxford

***Abstract*—Satellites in Geostationary Orbit (GEO) provide a number of commercial, government, and military services around the world, offering everything from surveillance and monitoring to video calls and internet access. However a dramatic lowering of the cost-per-kilogram to space has led to a recent explosion in real and planned constellations in Low Earth Orbit (LEO) of smaller satellites.**

These constellations are managed remotely and it is important to consider a scenario in which an attacker gains control over the constituent satellites. In this paper we aim to understand what damage this attacker could cause, using the satellites to generate interference.

To ground our analysis, we simulate a number of existing and planned LEO constellations against an example GEO constellation, and evaluate the relative effectiveness of each. Our model shows that with conservative power estimates, both current and planned constellations could disrupt GEO satellite services at every groundstation considered, albeit with effectiveness varying considerably between locations.

We analyse different patterns of interference, how they reflect the structures of the constellations creating them, and how effective they might be against a number of legitimate services. We find that real-time usage (e.g. calls, streaming) would be most affected, with 3 constellation designs able to generate thousands of outages of 30 seconds or longer over the course of the day across all groundstations.

I. Introduction

Space infrastructure underpins an array of services that the world relies upon, including GNSS for navigation and timing, broadcast media and internet access in remote areas. Historically, the cost and complexity of developing and deploying space infrastructure was so vast that it was accessible only to the most well-funded government and commercial entities. However, a 'New Space' revolution is underway, bringing orbital operations within the reach of a far larger number of entities, enabled by the development of flexible, commodity hardware; multi-tenanted launches; and an ecosystem of support services that have come about together [20].

Workshop on Security of Space and Satellite Systems (SpaceSec) 2023
27 February 2023, San Diego, CA, USA
ISBN 1-891562-89-4
https://dx.doi.org/10.14722/spacesec.2023.233980
www.ndss-symposium.org

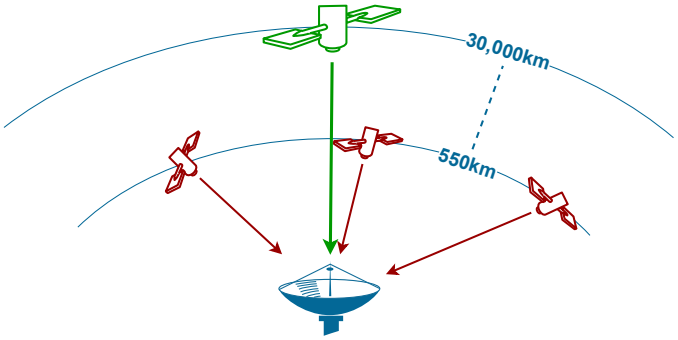

Fig. 1: An illustration of the attack setup. Three LEO satellites with signals reaching the victim receiver, interfering with the legitimate signal (centre).

Noting this vast expansion in the quantity of space infrastructure, emergence of new services (presenting more attack surface), and the newfound ease with which more participants can launch their own satellites, there is a novel potential for Downlink Interference Attacks (DIAs) by current constellations and those planned for the future. Despite the substantial disparity in transmission power between the attacker and victim satellites, the difference in orbital distance and the far greater number of potential attacking satellites make an attack such as this important to consider.

For years, terrestrial Distributed Denial of Service (DDoS) attacks have mostly used 'botnets' composed of thousands of compromised computers to disrupt services across the internet [7]. An attack of this scale may seem disproportionate for space, but there are a number of multi-thousand satellite constellations planned. It is therefore worth examining whether these constellations might represent the same threat. They would operate on slightly different principles, relying on the fact that with a sufficiently dense constellation, there will always be a satellite close to the line between the victim satellite and groundstation, where interference would be strongest. However, the same attack vector exists: thousands of near identical systems that through direct (transmitted to the satellite, targeting remote management interfaces) or indirect (via groundstation attack or insider threat) means can potentially be compromised from almost anywhere on Earth.

Even with no malicious activity, it is recognised that interference between satellites operating at the same frequen-

cies can lead to degraded communications [14]. As such, modern LEO constellations are required by regulators like the Federal Communications Commission (FCC) to address how they plan to avoid interference with existing infrastructure. As evidenced from the filings themselves, they take a good-faith, proactive approach to this by "ensuring that there is the necessary amount of angular separation" between themselves and potential victim satellites when targeting groundstations [23, p. 27]. This works because the receiving antenna will have high directionality at the angle directly towards its intended satellite. This interference-avoidance therefore relies on each component of the constellation continually monitoring its own position relative to a number of other satellites to avoid broadcasting in case of insufficient angular separation.

In our case, the attacker either does not implement this measure, or overrides it, such that satellites transmit irrespective of their angular separation, as shown in Fig. 1.

Concretely, we make the following contributions:

- We design and implement a simulation framework for Downlink Interference Attacks.
- We evaluate current and planned constellations' effectiveness at jamming the Inmarsat Global Xpress system in the Ku/Ka bands.
- We analyse the effects of constellation choice, geometry and receiver location upon the jamming effectiveness.
- We discuss the impact of an attack on realistic GEO use-cases.
- We discuss the limitations of CubeSats in the communications space, and their potential for disrupting GNSS services.

II. Background

A. New Space

'Old Space' satellites were almost all built by governments or large defence/communication companies. Every satellite was custom-built for its mission over a span of years, and excessively tested on the ground to ensure its success and reliability in space.

'New Space' satellites on the other hand make use of commercial-off-the-shelf (COTS) hardware for the most part, drastically reducing costs and development time. This is largely in part due to the CubeSat Design Specification, which has presented a modular standard for others to build upon. Even non-CubeSats, like the satellites in the Starlink constellation, are using COTS hardware alongside their own proprietary hardware.

In general these CubeSats are significantly less powerful than traditional 'Old Space' satellites, since they focus on being small, light, and easy to build. However this cheap approach is reflected not only in their physical hardware and software, but also in their orbit. These satellites are almost exclusively in Low Earth Orbit (LEO), and what they lose in coverage, they dramatically gain in signal strength. This is understandably important when using COTS SDRs and antennas, since much more powerful communications systems would be required in Geostationary Orbit (GEO), and on Earth to compensate. The final advantage, and the most crucial for satellites like Starlink, is the lower latency. In order to achieve the responsiveness required for real time video calls, gaming, and streaming, the latency must be at least of the same order of magnitude as terrestrial broadband (20 ms). It's clear then that GEO, where the round trip for light is 200 ms not including any processing at either end, would never be feasible. Starlink on the other hand, in its orbit of roughly 550 km, has been measured to have a latency of 45 ms.

Given the fact that each satellite in LEO can see much less of the Earth at any given point, and it is only above a particular for a matter of minutes before its rapid orbit takes it over the horizon. If global coverage is desired, then hundreds or thousands of satellites are required, which is the basis for constellations.

B. SDRs in Space

Software-defined radios (SDRs) have dramatically lowered the barrier to entry for satellite operators, and increased flexibility. This represents an evolution from earlier 'bent-pipe' repeater systems towards fully regenerative designs. The resulting ability to change frequency, modulation, encoding schemes, and processing on the fly enables a much more dynamic and adaptable system.

Licenses for space operation give strict frequency bands and operating parameters for satellites. As identified in [26], a satellite developer must provide a transmitter survey and potentially verification of emission characteristics, however these mostly rely on a benign satellite operator providing trustworthy data [6], [21]. The adversarial analysis done by Pavur et al. shows that physical changes are limited by size and power, but software is extremely difficult to completely externally verify without a detailed code review, which is currently not required for all commercial launches.

This means satellites verified for one mode of operation on the ground can change once in orbit. They have the physical capability through the use of SDRs and wide-band antennas to broadcast at frequencies beyond their license, though this may be limited by other hardware such as RF filters. Despite this, due to the congested nature of frequency band allocation, even small excursions from the assigned band of operation will almost certainly overlap with other satellites.

C. Related Work

The main body of intersatellite interference research, small though it is, focuses on ensuring that the good faith solutions proposed work to keep interference at acceptable levels, and the side effects on both systems. The effectiveness of angular separation and beam control in reducing interference have been studied with a view to understanding the cost to LEO coverage that it imposes [19], [30], and steps have been taken to find analytical alternatives to orbital simulation for efficient interference calculation [31].

These studies however do not take an adversarial perspective or consider a threat model, and so are not able to provide an answer for a number of our questions surrounding the potential capabilities of different actors. They also lack a comparative component considering often a single constellation at a time, making it much more difficult to generalise the

Attack	Military	Intelligence Agency	Corporate Insider	Hardware Supplier	APT	Organised Crime	Terrorist Group	Individual Hacker	Activist Group
Constellation Launch	✓	✓	x	x	x	*i*	*i*	x	*i*
Pre-Launch Tampering	✓	✓	✓	✓	*i*	*i*	*i*	*i*	*i*
In-Space Tampering	✓	✓	✓	x	✓	x	*i*	*i*	*i*
Distributed Jamming	✓	✓	x	x	✓	x	x	*i*	*i*

Key: ✓ - Attacker is likely both capable of executing the attack and motivated to do so. *i* - Attacker is likely interested in the attack, but has limited capacity to execute it. **x** - Attacker is likely neither interested in nor capable of executing the attack. Note: There may be crossover between categories, such as an insider threat sponsored by an intelligence agency. This matrix is intended as a demonstrative summary of the four major stages of the attack, with one of the first three required for execution of the fourth.

TABLE I: Threat Matrix

results or find which constellations might be more appealing for takeover/creation.

"Should We Worry About Interference in Emerging Dense NGSO Satellite Constellations?" [9] studies a number of constellations and their interference effects, but does so with a Monte Carlo approach to satellite positioning. This Monte Carlo approach provides a more even distribution of relative positions, but does not allow us to see how interference changes over time. This is crucial when considering how different interference patters affect the user's ability to interact with the system, as explored in Section VI-A.

III. Threat Model

We consider a threat model in which the victim is a legitimate user of the GEO satellite(s), for communication or navigation purposes. The attacker seeks to disrupt the availability of a GEO signal for a period of time, by injecting noise or interfering signals.

An attacker in this scenario has either the capability to take control of a substantial portion of an existing constellation, be it through malicious cyber attack or state-level take-over, or funding to place their own constellation into orbit. While widespread attacks on satellite constellations are not yet commonplace, we nevertheless believe it to be a realistic scenario in the near future. A range of security mechanisms do exist in contemporary constellations, such as separation of data- and control-plane traffic and standards for link encryption and constellation protection [11], [22]. However, similar measures have not fully prevented attacks in other fields and it is reasonable to assume that vulnerabilities could still exist here. Based on the framework proposed by Pavur et. al [25], we present in Table I a threat matrix covering major threat actors and their capabilities to engage in such attacks. We assess the attack to be feasible for nations (*Military*, *Intelligence Agency*), insiders at satellite operators or manufacturers (*Corporate Insider*, *Hardware Supplier*) and the most skilled cybercriminal groups (*APT*).

The number of entities with the resources to produce, launch, and maintain a constellation of their own, is smaller. While a small number of CubeSats have been launched by university consortia or private enterprises [2], [3], at present only a handful of providers (e.g. Starlink, OneWeb, or Iridium) have launched commercial constellations with more than 100 satellites [1].

IV. Simulation

The simulation is composed of two stages. The first stage is an orbital propagation model that accepts a description of the victim and attacker constellations and calculates their progression over time; yielding the relative positions and geometry at given time steps. The results of this are passed to a second stage, which models satellite downlink radio transmissions; computing their relative power and the resultant signal integrity. This two-stage architecture decouples the models, allowing one orbital scenario to be analysed under different radio-propagation conditions, without repeating the orbital calculations.

A. Orbital Propagation

Satellite simulations were performed in FreeFlyer, an orbital propagation program designed to simulate multiple groundstations, satellites, and celestial bodies over a period of time. The simulation accepts both existing Two-Line Element sets (TLEs), or constellations specified via the GUI or scripting language. TLEs are a format of storing satellite positions and orbits, allowing tracking. The simulation then takes these known satellite positions and orbits, and steps through to calculate their position over time. A number of techniques are available to do this efficiently, though we use FreeFlyer's default for simplicity [1]. Atmospheric drag is included in the simulation, though not relevant for the timescales involved.

B. Radio Communication

To estimate interference we need to understand how the signals interact with each other at the ground. The signal strength depends on the frequency of signal, the distance between the transmitter and the receiver, the power and gain of transmitter and receiver, the angle between them, and signal loss to atmosphere as seen in Eq. 1.

[1]Runge Kutta 8(9) Integrator, with a fixed step size of 300s, and a Relative Error Tolerance of 1e-9, with Norad SGP4 for TLEs.

Value	Name	Unit
$EIRP$	Effective Isotropic Radiated Power	dBW
A	Atmospheric Attenuation	dB
f	Transmission Frequency	Hz
b	Transmission Bandwidth	Hz
c	Speed of Light	ms^{-1}

TABLE II: Constants

We simulate the satellite's orbital propagation (Section V-A) and at each time step, with the distances and angles recorded, we can use Eq. 1 to determine the received power r_p from each satellite. However, the effect of transmissions that are occluded by the Earth is negligible, so these are excluded for reasons of computational efficiency. At every simulation step, a set of the visible LEO satellites is generated (those more 15 degrees above the horizon) and only those are used in further calculations. Signal powers are computed for each victim GEO satellite and then the strongest is subsequently used, representing the most difficult signal for the attacker to overpower. For each of the selected LEO attackers and the strongest GEO victim, the simulation stores distance, azimuth, and elevation.

Given constants as shown in Tab. II and values for satellite distance d and incident angle at the receiver φ, we calculate r_p as follows:

$$r_p = EIRP + r_g - L - A \tag{1}$$

$$L = 20 \cdot \log\left(\frac{4\pi df}{c}\right) \tag{2}$$

Due to the highly-directional antennas employed, the receiver antenna gain r_g varies substantially with angle φ. In our simulation it is modeled according to the European Radiocommunications Committee's (ERC) report on analysis of Inmarsat receivers [12]. Details of the calculation are provided in App. A.

For efficiency, all time-invariant factors in a given scenario (e.g. EIRP, frequency-dependent path loss components) are precomputed, leaving only time-variant factors (distance, azimuth) to be calculated in the main simulation loop.

The results of the received power calculations allow us to compute the Signal to Interference plus Noise Ratio (SINR) experienced at the receiver. SINR provides a measure of signal clarity, and is used to determine whether a groundstation is likely to be able to correctly decode the legitimate signal.

Both the orbital propagation and RF simulation code will be made publicly available at https://github.com/ssloxford to assist additional security research into satellite interference in future. A visualisation by the authors for the constellations used can be found at https://cotstellation.netlify.app.

V. Evaluation

A. Experiment Configuration

As demonstration groundstation locations we chose those listed as AWS Ground Stations [2]; illustrated in Fig. 2. These

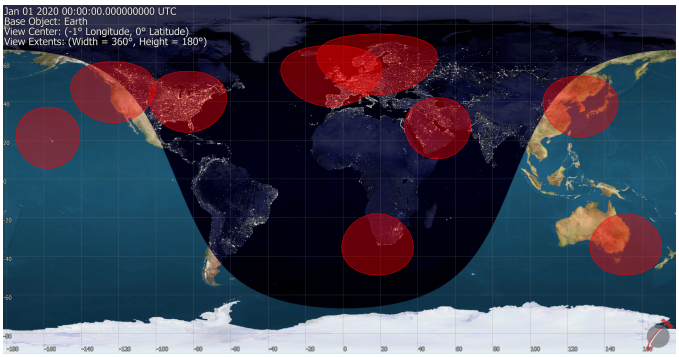

Fig. 2: AWS Groundstation locations used as examples for receivers.

are 10 stations around the world, and as the Ground-Station-As-A-Service industry grows, may represent a substantial part of satellite-ground communications. While fairly evenly distributed longitudinally, they are focused mostly in the Northern Hemisphere.

While selecting the locations as above, we consider the receiver hardware to be that of an Inmarsat receiver, consistent with that service being considered the victim (this includes an Inmarset-specific antenna pattern, as described in Sec. IV-B). Consistent with this, we set f = 19.2 GHz and b = 250 MHz, as used by the Inmarsat downlink (and still covered by the Ku-Band Starlink user downlink transmitters) [27]. We also set A = 0.35 dBW [18] and EIRP according to values for each constellation under consideration (see Tab. III).

The victim group of GEO satellites chosen belong to the Inmarsat Global Xpress Series. These were chosen as they represent an existing communications network, exemplify well the 'Old Space' approach of few, powerful satellites, and operate in the same frequency bands as Starlink and OneWeb (Ku/Ka band).

Inmarsat do not release much information about the power of their satellites, however the limited information on the 2, 3, 4 series as well as the then upcoming Alphasat (which became Inmarsat-4A F4) give the maximum EIRP link of 70 dBW [17], which we have imcreased by an order of magnitude to account for technical development, and ensure we do not overestimate interference.

The orbits for the 5 satellites comprising the Global Xpress series were taken from Celestrak TLEs [1]. A total of six different constellations were modelled as the attacking party. These consisted of various deployments by Starlink [29], OneWeb [24] and PlanetLabs [4], along with a hypothetical CubeSat constellation. Details of each can be found in Tab. III.

The simulation was run with a propagation step size of 10 seconds, and a total elapsed time of 24 hours. We consider an attack to constitute successful jamming when it reduces the SINR at the receiver below 10 dB [3]. This is a coarse initial value used to standardise the results across constellations. In practice, this value will vary based on the modulation and

[2] https://aws.amazon.com/ground-station/locations/

[3] Most services do not publish their receivers' required SINR, but one service lists 10 dB as required to show correct alignment [8] and that figure is used here.

Constellation	Status	Number	Satellite	EIRP (dBW)	Orbit (km)
Inmarsat Global Xpress **(Victim)**	Exists	5	Global Xpress	80	35,000
PlanetLabs Dove Constellation	Exists	150	Dove	8	530
CubeSat Walker Constellation	Hypothetical	396	CubeSat	6	550
Starlink First Group	Exists	1,584	Starlink	39	550
Starlink Phase 1	Planned	4,408	Starlink	39	550
One Web Phase 1	Planned	716	OneWeb	45	1,200
One Web Phase 2	Planned	6,372	OneWeb	45	1,200

Note: Due to limited public information, the numbers here come from a variety of sources, or represent best estimations from available public filings. See Appendix B for more details.

TABLE III: Constellation Details

coding schemes, which may be updated dynamically by the victim system.

B. Results

With no interference, the average received signal strength from Inmarsat was -96.02 dBW, with an SINR of 23.98 dBW. Aggregated interference figures from each constellation, considered independently, are given in Tab. IV. While these averages provide a small insight into the relative attack strength possible using each constellation, they obscure important details that become clear when examining the behaviour over time.

TABLE IV: Mean Interference Results

Name	Mean SINR (dBW)	Mean Time Jamming (%)	Mean period of jamming (s)
Dove	23.94	0	0
CubeSat	23.96	0	0
Starlink First Group	13.27	12.41	31.25
Starlink Phase 1	9.80	27.35	38.88
OneWeb	15.56	7.36	47.52
OneWeb Phase 2	3.81	77.67	150.67

1) Feasibility of Denial of Service: Fig. 3 shows the percentage of time that a receiver is jammed (SINR $<$ 10 dB) as a function of the EIRP of the attacking satellites, when assuming two different attacker constellations. Fig. 3a plots this function for the Starlink First Group constellation (as currently deployed), while Fig. 3b plots it for the full Starlink Phase 1, expected in the future. It can be seen that attacks are affected by the size of the constellation (1,584 satellites in Fig 3a vs. 4,408 in Fig 3b), with the disruption curves rising earlier and faster when the larger constellation is used. It is also clear that the effectiveness is heavily affected by the EIRP, with almost no jamming possible if EIRPs are below 30 dBW and near-total effectiveness for EIRPs over 55 dBW.

The plots in Fig. 3 show the jamming effectiveness of each constellation when emitting at a range of powers. However we can use the constellations' known power values (or expected values). For Starlink First Group, with an actual EIRP of 39 dBW, it means that for all groundstations barring Stockholm there would be disruption $5-15\%$ of the time (40% at Stockholm). The greater number of satellites in Starlink Phase 1 benefits the attack equivalent to roughly 5 dBW of extra power. In other words, an attacker could use ≈ 5 dBW less power per satellite (3.16 times less) if using the larger Starlink Phase 1 constellation.

2) Effect of Groundstation Location: From the groundstation curves, in Fig. 3, it can be seen that the non-uniform layout of the constellations leads to different levels of effectiveness at different groundstations. Even if the curves between stations are similar, their slight offset and rapid growth means that for the same power level they might experience very different jamming results. For example, comparing the SINR directly over the course of the 24 hours between Manama and Stockholm it can be seen that the sporadic coverage given to Manama in Fig. 4a compared to the denser coverage over Stockholm (Fig. 4b), an effect noted in [31]. Due to its inclination of $53.2°$, Starlink has a much higher density near groundstations like Dublin and Stockholm ($53.3°$ and $59.3°$ latitude respectively), than the likes of Honolulu and Manama ($21.3°$ and $26.2°$ latitude respectively). This shows how important specific orbital parameters, as well as the locations an attacker wishes to jam, are in determining the potential jamming effectiveness of a constellation.

3) Effect of Constellation Choice: However, not all constellations exhibit the patterns shown above. OneWeb, for example, has a much more cyclic and predictable nature, as shown in Fig. 4c. While the receiver is jammed for less time overall, it is more consistent during jammed periods. Looking at time jamming across all ground stations for OneWeb as seen in Fig. 3c the lower number of satellites, higher orbit, and strict periodic nature of the jamming leaves it much less effective than Starlink, despite the higher power of each satellite.

Contrasting constellations based on density, such as between OneWeb (Fig. 4c) and OneWeb Phase 2 (Fig. 4d), we can also see how much of a difference the 'secondary' satellites make (those not directly between the groundstation and victim satellite). The greatest difference between the single strongest interfering satellite, and the sum of all interfering satellites is ≈ 5 dBW for One Web Phase 2.

4) CubeSat vs. Commercial New-Space: So far we have still been considering constellations with hundreds or thousands of relatively powerful satellites however. Looking at the other end of the spectrum, at a much more achievable CubeSat constellation, the results are very different.

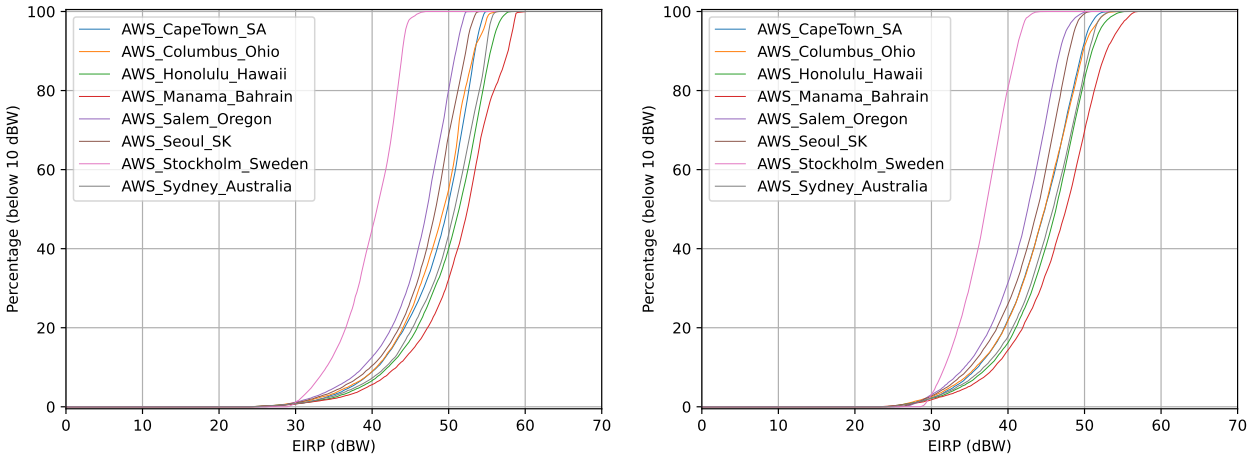

(a) Starlink First Group

(b) Starlink Phase 1

(c) OneWeb

(d) CubeSat Walker Constellation

Fig. 3: Time spent jamming over 24 hours as a function of the power of each satellite in the constellation. Jamming threshold taken as < 10 dB SINR.

The CubeSat constellation (Fig. 3d) itself has better performance compared to OneWeb (Fig. 3c), with the notable improvement of Stockholm (again explained by the lower inclination), until a power of 60 dBW. This is currently far beyond the power achievable at such a scale, and so matters much less in comparison. The crucial difference however the actual EIRP of each CubeSat sits at 6 dBW, where no groundstation is being jammed at all. Taking an example groundstation, in this case Salem, we can see that the constellation never even comes close to appreciably reducing the SINR.

VI. Discussion

The results give a number of interesting considerations. For one, the latitude makes a large difference, with groundstations closer to the equator suffering less interference than those higher up. This is chiefly because constellations with a non-polar inclination (such as Starlink) are much denser at the latitude of their inclination (53.2 degrees). This means that there are more secondary satellites contributing, and the primary satellites spend more time with low angular separation.

Next we can see that the nonlinearity of the SNIR curves (such as Fig. 3a) means a relatively small change in power can result in a dramatic increase in jamming time. Satellite manufacturers are often very restrictive in the information they provide about their satellites. As such, lots of the information used is taken from legal filings. This has the unfortunate consequence that all of the information provided is what the manufacturers *intend* to do, not their full capabilities (which is what we care about in an adversarial scenario). There are numerous references to lowering satellites' EIRP to acceptable levels, which implies that they are more capable than the given figures suggest [29]. More concretely, RF transmitters are typically operated with a backoff of 3-6dBs to ensure their linearity. This is of much less concern while jamming, and so moderate increases of transmit power in the order of 3dB might be achievable once compromised by raising the amplifier gain to power saturation levels, subject to the satellite's DC power

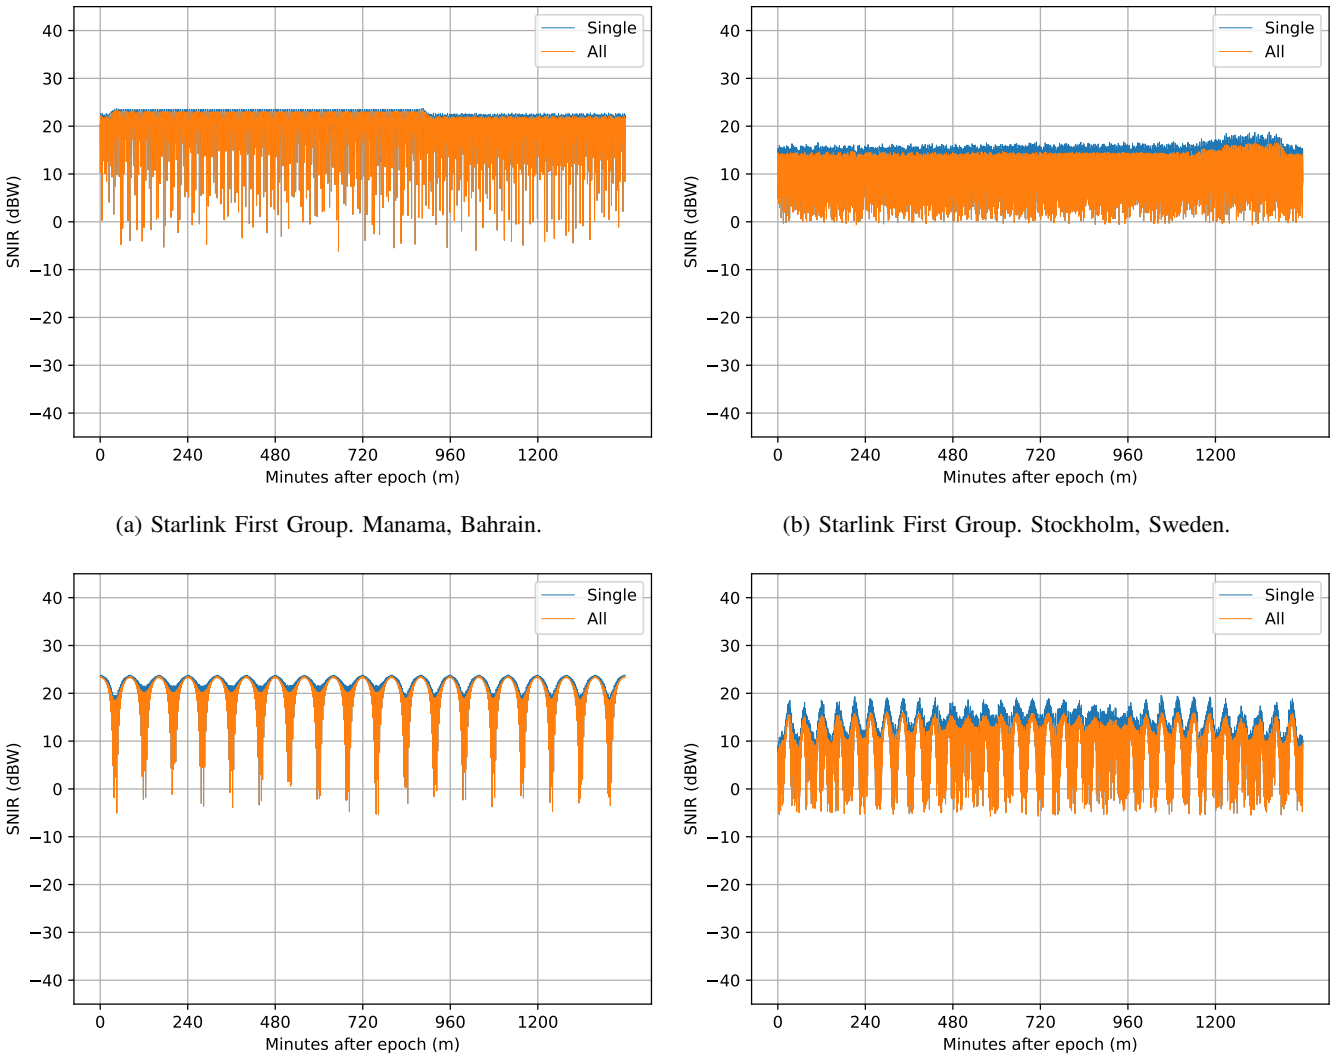

(a) Starlink First Group. Manama, Bahrain.

(b) Starlink First Group. Stockholm, Sweden.

(c) OneWeb. Manama, Bahrain.

(d) OneWeb Phase 2. Manama, Bahrain.

Fig. 4: SINR measured at ground receivers over a 24-hour period, for single satellites and the attacking constellation as a whole. Each point represents a 10 s window. Jamming threshold taken as < 10 dB SINR.

and thermal management constraints.

Finally, the results show that CubeSats simply do not have the power required to compete with most Ka-band services. However, communications are not the only space-based service in use, Global Navigation Satellite Systems (GNSS) being a key alternative. We describe in App. C an analysis of the potential for GNSS interference using CubeSats.

A. *Usage Analysis*

We have seen that constellation design plays a large part in the type of jamming available, with Starlink and the CubeSats having frequent but seemingly random spikes, contrasted with OneWeb and more PlanetLabs' periodic jamming. This could well be an advantage for an attacker, knowing when a given system might be down (if only briefly). However if the user/defender is attentive to this pattern, and does not require 100% uptime, then it may be very easy for the defender to adapt and instead use the system outside of these times. One disadvantage for the attacker is that they are completely constrained by orbital dynamics, and so cannot alter when they jam (though can choose to jam less than maximally possible). This will clearly work best if the defender can't choose to use the system outside of the times of jamming.

We can look more closely at how long each constellation is able to jam for by plotting a histogram of jamming time across all groundstations. Figure 5 shows this distribution for both Starlink constellations, as well as both OneWeb constellations. We can see in Fig 5a, Starlink never jams for longer than 90 seconds. This means that as the victim's usage is tolerant to interruption periods of 90 seconds. This is acceptable for messaging or file transfers, for example, but would render real-time communication via video or speech unusable.

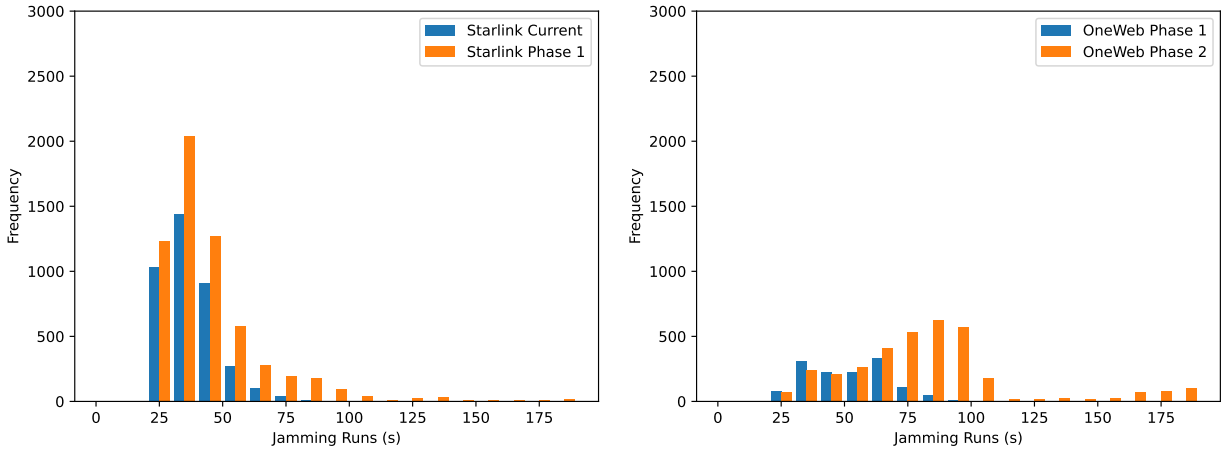

(a) Starlink vs Starlink Phase 1

(b) OneWeb vs OneWeb Phase 2

Fig. 5: Figures showing histograms of 'jamming runs'. These show the frequencies of time periods of constant interference (SNIR $<$ 10 dBW).

However, Starlink Phase 1 has (limited) interference capability up to almost 3 minutes, and OneWeb Phase 2 has an average period of 2 and half minutes, happening frequently. This presents a very serious problem for almost any use-case. Remotely managing systems with consistent two minute delays dramatically lengthens control feedback loops (i.e. between triggering a command and seeing the effect), which in industry may have expensive or dangerous consequences, especially if the satellite communications are used for warning or monitoring systems [5].

B. Attack Mitigations

The clearest opportunity for a user on the ground to reduce the effectiveness of the attacks considered here, is the use of a reception antenna with a more focused antenna pattern. This dramatically reduces not only the effect of 'secondary' interfering satellites, but also minimises the amount of time a primary constellation satellite spends at a 'high-gain' angle (i.e. sufficiently close to the peak gain of the antenna). This does not completely eliminate the possibility of interference, but limits how often it could cause serious disruption. We discuss the pattern specifics, both current and an ITU recommendation, in App. A.

More fundamental changes could also be made, either for new satellites or where existing assets can be modified. In particular, spread spectrum techniques could be applied, forcing the attacker to learn spreading or hopping patterns. These are commonly implemented on military systems [10], but would require access to appropriate spectrum and securely-managed keys.

VII. Conclusion

In this paper we have modeled interference, from a variety of current and planned constellations, directed at a number of groundstations. We found that jamming would already be possible, and its potential effectiveness is only going to increase as constellations grow larger. However, commercial constellations currently present the main threat as even the most effective constellation layout considered (OneWeb Phase 2) would require a minimum EIRP per satellite of roughly 30 dBW; two orders of magnitude more powerful than the biggest CubeSats (10 dBW). For Ku/Ka band, CubeSat interference does not seem to pose a problem, though it might have serious consequences for GNSS infrastructure.

Given the equipment requirements, it is unlikely that deploying a constellation solely dedicated to this attack would be feasible outside of exceptionally well-funded threat actors. However targeting an existing or planned constellation continues to be a viable option and is only exacerbated by growing interest in the space industry and the proliferation of off-the-shelf software and hardware components.

The security policies of satellite operators should reflect the consequences of this misuse. It is not enough to provide reports on avoiding interference through modulating power, direction, and operation conditions, as is currently required; these all rely on a benign operator. To most effectively minimise their impact on other space assets, comprehensive external code reviews, adversarial analysis, and continual monitoring would be required to ensure system integrity.

The potential mitigations mentioned in Section VI-B go some way towards alleviating the potency of such an attack, though these rely on more capable and therefore expensive receivers, which may limit their implementation.

REFERENCES

[1] Celestrak. [Online]. Available: https://www.celestrak.com/NORAD/elements/

[2] Newspace index: Swarm technologies. [Online]. Available: https://www.newspace.im/constellations/swarm-technologies

[3] Qb50: An international network of cubesats. [Online]. Available: https://www.qb50.eu/index.php/consortium.html

[4] "Planet labs tles," Jan 2022. [Online]. Available: https://ephemerides.planet-labs.com/

[5] E. S. Agency. Using space technology to monitor offshore oil and gas fields. [Online]. Available: https://www.esa.int/Applications/Technology_Transfer/Using_space_technology_to_monitor_offshore_oil_and_gas_fields

[6] Air Force Space Command. Range safety user requirements manual volume 3 - launch vehicles, payloads, and ground support systems requirements. [Online]. Available: https://static.e-publishing.af.mil/production/1/afspc/publication/afspcman91-710v3/afspcman91-710v3.pdf

[7] E. Alomari, S. Manickam, B. B. Gupta, S. Karuppayah, and R. Alfaris, "Botnet-based distributed denial of service (ddos) attacks on web servers: Classification and art," *International Journal of Computer Applications*, vol. 49, 08 2012.

[8] bigblu. Technical faqs. [Online]. Available: https://bigblu.ie/technical-faqs/

[9] C. Braun, A. M. Voicu, L. Simić, and P. Mähönen, "Should we worry about interference in emerging dense ngso satellite constellations?" in *2019 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2019, pp. 1–10.

[10] C. Brown and P. J. Vigneron, "Robust frequency hopping for high data rate tactical communications," Tech. Rep. [Online]. Available: https://www.sto.nato.int/publications/STO%20Meeting%20Proceedings/RTO-MP-IST-054/MP-IST-054-P12.pdf

[11] CCSDS, "Space data link security protocol (ccsds 355.0-b-2)," Consultative Committee on Space Data Systems, Tech. Rep., 2022.

[12] E. R. Committee, "Erc report 91 assessment of interference from unwanted emissions of ngso mss satellite transmitters operating in the space-to-earth direction in the band 1621.35 - 1626.5 mhz to gso mss satellite receivers operating in the earth-to-space direction in the band 1626.5 - 1660.5 mhz," European Conference of Postal and Telecommunications Administrations, Tech. Rep., 2000.

[13] K. Devaraj, M. Ligon, E. Blossom, J. Breu, B. Klofas, K. Colton, and R. W. Kingsbury, "Planet high speed radio: Crossing gbps from a 3U cubesat," Tech. Rep., 2019.

[14] FCC. Significant satellite rulemakings. [Online]. Available: https://www.fcc.gov/significant-satellite-rulemakings

[15] Furuno. Advanced ka-band antenna terminal designed for inmarsat global xpress. [Online]. Available: https://www.furuno.com/files/Brochure/306/upload/FV-110GX_EN.pdf

[16] ITU. Reference FSS earth-station radiation patterns for use in interference assessment involving non-gso satellites in frequency bands between 10.7 GHz and 30 GHz. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/s/R-REC-S.1428-1-200102-I!!PDF-E.pdf

[17] Jack Deasy (Director, UAS Programs). Inmarsat update. [Online]. Available: http://faculty.nps.edu/cdprince/mwc/docs/MWC_CONF/2011_-7_19-20_Conf/Deasy_Inmarsat_Intro_and_Program_Update.pdf

[18] Kymeta, "Link budget calculations for a satellite link with an electronically steerable antenna terminal," Kymeta, Tech. Rep., 2019. [Online]. Available: https://www.kymetacorp.com/wp-content/uploads/2020/09/Link-Budget-Calculations-2.pdf

[19] R. Li, P. Gu, and C. Hua, "Optimal beam power control for co-existing multibeam geo and leo satellite system," in *2019 11th International Conference on Wireless Communications and Signal Processing (WCSP)*, 2019, pp. 1–6.

[20] M. Manulis, C. Bridges, R. Harrison, V. Sekar, and A. Davis, "Cyber security in new space," 2020. [Online]. Available: https://doi.org/10.1007/s10207-020-00503-w

[21] NASA. (2017) Cubesat 101: Basic concepts and processes for first-time cubesat developers. [Online]. Available: https://www.nasa.gov/sites/default/files/atoms/files/nasa_csli_cubesat_101_508.pdf

[22] C. on National Security Systems, "Cybersecurity policy for space systems used to support national security missions (cnssp no. 12)," Committee on National Security Systems, Tech. Rep., 2018.

[23] OneWeb, "Oneweb non-geostationary satellite system (leo) phase 2: Modification to authorized system attachment b," Tech. Rep., 2019. [Online]. Available: https://fcc.report/IBFS/SAT-MPL-20200526-00062/2379706.pdf

[24] ——, "Amendment to modification application for U.S. market access grant for the oneweb ku- and ka-band system," Tech. Rep., Jan 2021. [Online]. Available: https://fcc.report/IBFS/SAT-MPL-20210112-00007/3495551.pdf

[25] J. Pavur and I. Martinovic, "SOK: building a launchpad for impactful satellite cyber-security research," *CoRR*, 2020. [Online]. Available: https://arxiv.org/abs/2010.10872

[26] J. Pavur, M. Strohmeier, V. Lenders, and I. Martinovic, "In the same boat: On small satellites, big rockets, and cyber trust," 05 2021, pp. 151–169.

[27] B. C. S. Services. Inmarsat-5: Overview brief for arctic. [Online]. Available: https://www.ntia.doc.gov/files/ntia/boeing_satcom_for_alaska1.pdf

[28] A. Space. Totem nanosatellite sdr platform. [Online]. Available: https://www.cubesatshop.com/product/totem-nanosatellite-sdr-platform/

[29] SpaceX, *SpaceX Non-Geostationary Satellite System Attachment A*. FCC, Nov 2018. [Online]. Available: https://fcc.report/IBFS/SAT-MOD-20181108-00083/1569860.pdf

[30] H. Wang, C. Wang, J. Yuan, Y. Zhao, R. Ding, and W. Wang, "Coexistence downlink interference analysis between leo system and geo system in ka band," in *2018 IEEE/CIC International Conference on Communications in China (ICCC)*, 2018, pp. 465–469.

[31] T. Wang, W. Li, and Y. Li, "Co-frequency interference analysis between large-scale ngso constellations and gso systems," in *2020 International Conference on Wireless Communications and Signal Processing (WCSP)*, 2020, pp. 679–684.

APPENDIX

A. Receiver Gain

Here we provide a description of the angle-dependent receiver gain modelling used in our simulation, based on the European Radiocommunications Committee's (ERC) report on analysis of Inmarsat receivers [12] (with $G_{\max} = 44$ [15]) which for 19.2 GHz (Inmarsat's downlink frequency) gives:

$$G(\phi) = \begin{cases} G_{\max} - \left(\frac{D}{\lambda} \cdot \frac{\varphi}{20}\right)^2 & \text{for } 0 < \varphi < \varphi_{11} \\ G_1 & \text{for } \varphi_m < \varphi < 100\frac{\lambda}{D} \\ 52 - 10\log\frac{D}{\lambda} - 25\log\varphi & \text{for } 100\frac{\lambda}{D} \leq \varphi < 48 \\ -10 & \text{for } \varphi > 48 \end{cases} \tag{3}$$

A plot of the receiver gain, following this model, can be seen in Fig. 6 (blue), alongside the ITU recommendation [16] that would provide much lower interference.

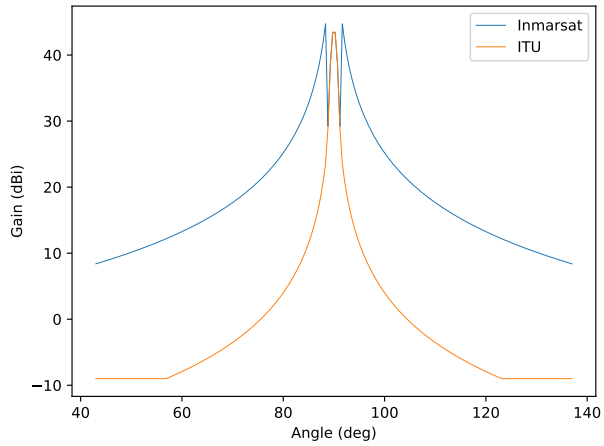

Fig. 6: Current Inmarsat pattern (blue), ITU recommendation (orange).

B. Satellite Details

Here we describe the characteristics of the satellites from each of the considered constellations.

a) Dove Satellite: Planet Labs have been extremely open about both the optics and communication details of their satellites. Given the vast quantity of data produced by the cameras (80 GB per pass), EIRP is more important for these imaging satellites than might be initially expected. Their paper gives clear numbers for EIRP: 8.2 dBW [13].

b) CubeSat: For the example cubesat, we used the EXA Kratos 1U [4] Cubesat Platform, with a 6 dBW patch antenna, and 1 W transmission power, providing an EIRP of 6 dBW.

c) Starlink: Based on the FCC Starlink Attachment, the highest usable EIRP density from Starlink satellites is 15.70 dBW/4kHz [29]. Since we are assuming bandwidth of 250 MHz, that gives an EIRP ≈ 39.68 dBW.

[4] https://www.cubesatshop.com/product/kratos-1u-cubesat-platform-1-step-solution/

d) OneWeb: The OneWeb FCC filing [23] gives a maximum downlink EIRP density of -2.7 dBW/4kHz. Using the same bandwidth as above, that gives a maximum EIRP ≈ 45.26 dBW.

C. GNSS Attacks using CubeSats

The European Space Agency (ESA) and others use S-Band (2220-2290 MHz down) for Telemetry, Tracking, and Command (TT&C) signals. This is relatively close to the L1-Band (1575.42 MHz) used by GNSS. A COTS SDR (for example the TOTEM Nanosatellite SDR Platform [28]) covers 70-6000 MHz, and therefore a satellite could reasonably use this SDR for operational or potentially malicious purposes. While interference wouldn't be effective, since GNSS is built to deal with extracting a weak signal from background noise, it does still leave open the option of a replay or spoofing attack to be explored in future work, since when simulated the CubeSat constellation still gives a more powerful signal than GNSS as seen in Fig. 7.

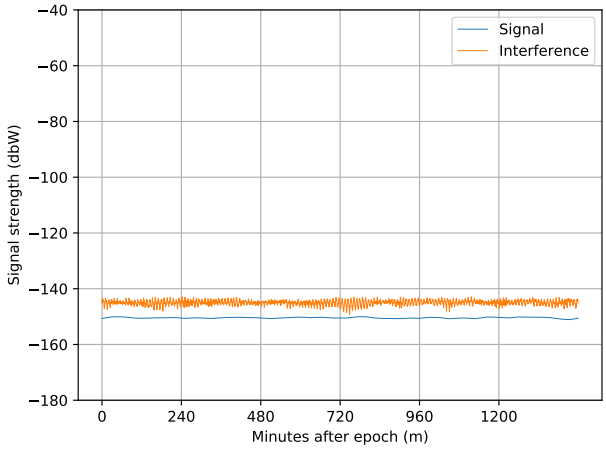

Fig. 7: A comparison between average simulated GPS signal strength (blue) compared to average CubeSat signal strength.

D. Plots for all Architectures

Fig. 8 shows comparative curves of jamming time as produced by each constellation considered in the analysis.

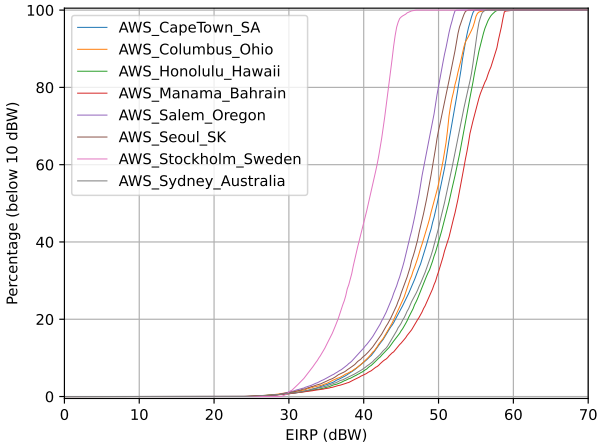

(a) Starlink

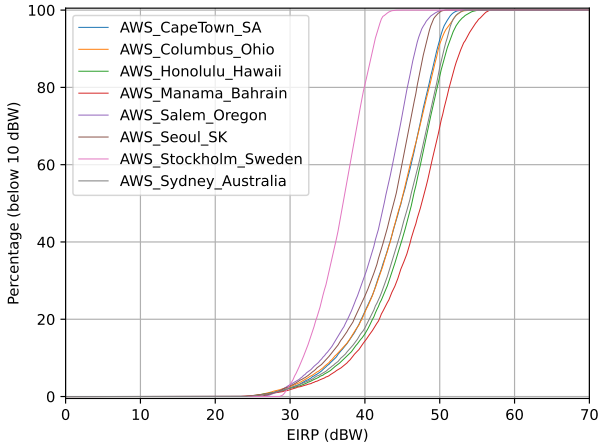

(b) Starlink Phase 1

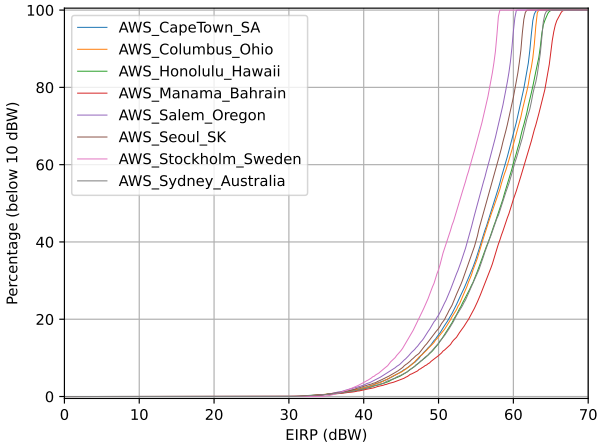

(c) OneWeb

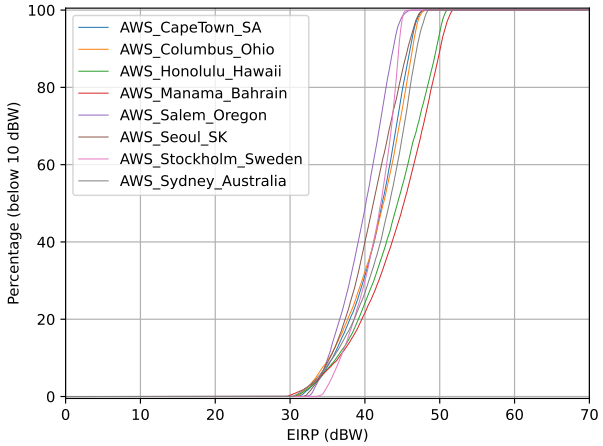

(d) OneWeb Phase 2

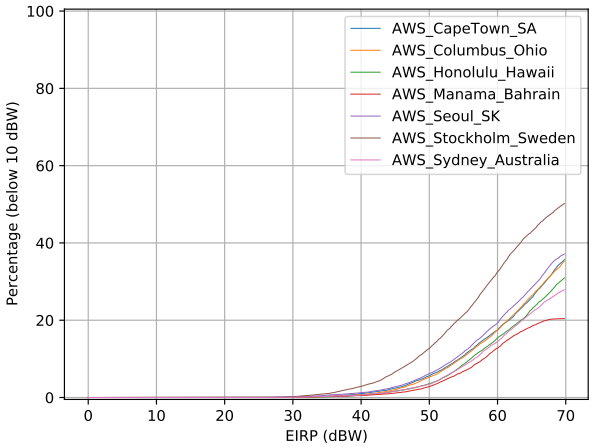

(e) Planet Labs' Dove

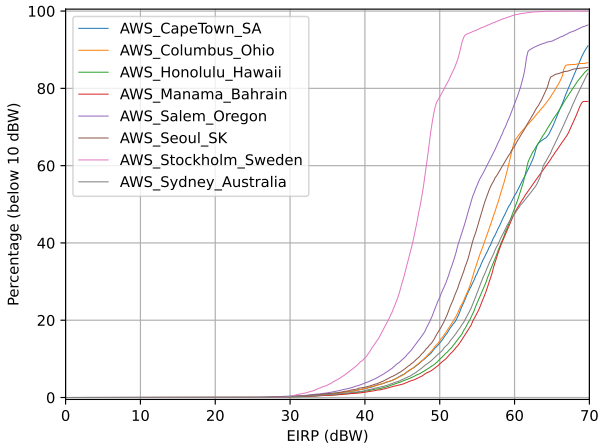

(f) CubeSat

Fig. 8: Groundstation curves for all constellations, showing percentage time jamming as a function of satellite power, with a threshold of 10 dBW.