

Securing New Space: On Satellite Cyber-Security

James Pavur

Wolfson College
University of Oxford

*A thesis submitted for the degree of
Doctor of Philosophy*

Hilary 2021

Abstract

Satellites offer critical services impacting the lives of billions around the world. However, the cyber-security properties of space systems are poorly understood. As the next generation of space missions begins to launch, there exists an acute need for robust and open research on space systems security.

This thesis offers a first step towards meeting that need, presenting a general method for uncovering novel cyber-physical security problems at the intersection of outer space and cyber-space. The foundation for this contribution is an analysis of historical space security incidents and emergent trends in space technologies and attacker capabilities.

Our method is developed through in-depth analysis of the security and privacy properties of modern satellite broadband services. Using real-world experiments, we identify previously unknown vulnerabilities impacting the security and privacy of millions of satellite customers, including many of the world's largest businesses. We further isolate underlying physical causes of these vulnerabilities and develop original techniques for their mitigation.

Building on this research, we systematize our method into a four-step process represented by the acronym RCMA (Recognize, Connect, Motivate, Adapt). This approach is then applied to further topics in space security, including Space Situational Awareness integrity and launch integration safety. In doing so, we find similar cyber-physical gaps between status-quo security practices and emergent threats. This recognition allows us to present novel solutions which bolster space mission security.

Over the course of this thesis, we identify and mitigate numerous technical vulnerabilities that impact hundreds of space platforms. More importantly, we propose and validate a flexible method for effective security research at the intersection of outer space and cyberspace. The intention throughout is to offer a launchpad for sustained and necessary work in an unusual domain.

Securing New Space: On Satellite Cyber-Security



James Pavur
Wolfson College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary 2021

Acknowledgements

So many incredible people were involved in making this thesis possible.

First among these is my wife, Casey, who followed me across the sea and has been there with love every step of the way. She has read each of my papers, commiserated and/or congratulated me through unceasing rounds of peer-review, and patiently endured the countless hours of soul-searching, brainstorming, whinging, and dreaming that make up a DPhil.

Second only to Casey is my supervisor, Ivan Martinovic. It takes a special kind of person to, upon learning that someone wants to hack satellites, invite them into his research group with open arms. I consider myself incredibly fortunate to have crossed paths with Ivan and cannot imagine a better supervisor or group.

Over the course of this research, I collaborated with many brilliant researchers. My frequent co-authors Daniel Moser, Martin Strohmeier, and Vicent Lenders from armasuisse have been an incredible intellectual resource throughout.

I wish to thank my assessors for their valuable feedback and guidance: Kasper Rasmussen, who helped this research evolve from its earliest conceptions in Transfer through to Confirmation and the Viva; Michael Goldsmith, who helped ground my early thoughts in the Transfer; Sadie Creese, who helped me organize everything in my Confirmation; and Greg Autry, who generously offered his time and expertise as an external examiner for the Viva.

Outside of work, I'd like to thank my friends and family for their support. My parents and siblings (Elisabeth, Gigi, and Sam) have sat through many a Google Hangout to send love, advice, and support halfway across the world. Likewise, D&D, board games, and dinners out with friends have helped keep me sane over the last few years.

Finally, I would not be here if not for the generosity of the Rhodes Trust and the people who supported my application to the scholarship. I'd especially like to thank my scholarship advisor, John Glavin, and recommenders: Samuel Visner, Eric Burger, Micah Sherr, Julia Lamm, Henry Quillian, Cole Ashcraft, Lisa Singh, Chen Gu and Ronak Shah. I'd also like to thank all the wonderful people at Rhodes House for their constant support, guidance, and commitment to the scholar experience.

Abstract

Satellites offer critical services impacting the lives of billions around the world. However, the cyber-security properties of space systems are poorly understood. As the next generation of space missions begins to launch, there exists an acute need for robust and open research on space systems security.

This thesis offers a first step towards meeting that need, presenting a general method for uncovering novel cyber-physical security problems at the intersection of outer space and cyber-space. The foundation for this contribution is an analysis of historical space security incidents and emergent trends in space technologies and attacker capabilities.

Our method is developed through in-depth analysis of the security and privacy properties of modern satellite broadband services. Using real-world experiments, we identify previously unknown vulnerabilities impacting the security and privacy of millions of satellite customers, including many of the world's largest businesses. We further isolate underlying physical causes of these vulnerabilities and develop original techniques for their mitigation.

Building on this research, we systematize our method into a four-step process represented by the acronym RCMA (Recognize, Connect, Motivate, Adapt). This approach is then applied to further topics in space security, including Space Situational Awareness integrity and launch integration safety. In doing so, we find similar cyber-physical gaps between status-quo security practices and emergent threats. This recognition allows us to present novel solutions which bolster space mission security.

Over the course of this thesis, we identify and mitigate numerous technical vulnerabilities that impact hundreds of space platforms. More importantly, we propose and validate a flexible method for effective security research at the intersection of outer space and cyberspace. The intention throughout is to offer a launchpad for sustained and necessary work in an unusual domain.

Contents

List of Selected Abbreviations	xi
I A Case for Space Security	1
1 Introduction	3
1.1 An Emergent Need	3
1.2 Structure and Methods	5
1.3 Key Contributions and Impact	9
1.4 Publications and Outputs	12
1.5 Research in Context	15
2 Building on 60 Years of Space Security Knowledge	17
2.1 Threat Modeling in Context	19
2.2 Learning from History	27
2.3 Defending The Signal	34
2.4 Defending Space Platforms	42
2.5 Defending Satellite Ground Systems	45
2.6 Holistic Security Models	48
2.7 Lessons and Opportunities	51
II Threats and Defenses in Satellite Broadband	55
3 Why Start with Broadband?	57
4 Old Vulnerabilities, New Applications: DVB-S MPE Threats	61
4.1 A Legacy of Exploitation	62
4.2 DVB-S Broadband Architectures	64
4.3 Full Horizon Survey: Experimental Design	67
4.4 Vulnerabilities and Findings	70
4.5 Improving Encryption: Status Quo Shortcomings	74
4.6 Summary	76

5	A Tale of Sea and Sky: Exploiting Maritime VSAT Services	79
5.1	Related Work on Maritime and Space Security	81
5.2	VSAT Broadband Applications and Architectures	84
5.3	Sustained VSAT Observation: Experimental Design	87
5.4	Threat Model and Attacker Capabilities	96
5.5	Broad Findings and Vulnerabilities	97
5.6	Findings: Physical Safety and Operations	102
5.7	Findings: Passenger and Crew Privacy	107
5.8	Active Attacks on VSAT Services	110
5.9	Underlying Causes and Next Steps	115
5.10	Summary	116
6	Making VPNs Work in GEO: The QPEP Architecture	119
6.1	The Need for New Approaches to GEO Encryption	121
6.2	System Design Requirements and Related Work	123
6.3	The QPEP System	131
6.4	QPEP Implementation	141
6.5	Secure PEP Testbed	147
6.6	Evaluating QPEP	148
6.7	Next Steps for Secure PEPs	162
6.8	Summary	163
7	GEO Broadband Security Lessons in Context	165
III	Cyber-physical Threats Beyond the Signal	169
8	Through the Lens of Physicality: Putting RCMA into Practice	171
9	Deceptions and Truths in Space Situational Awareness	177
9.1	Eyes on the Sky: SSA in the Status Quo	178
9.2	Threat Modeling	184
9.3	Scenario: Targeting SSA Projections	189
9.4	Scenario: Classification Deceptions and Defenses	196
9.5	Summary	216
10	Big Rockets, Small Satellites, and Cyber-Trust	219
10.1	Background: The Practice (and Politics) of Sharing Rockets	221
10.2	Threat Models for CubeSat Integration	223
10.3	Adversarial Analysis of Launch Safety Controls	225
10.4	Attack Simulation and Evaluation	229
10.5	Summary	236

IV	Conclusion	239
11	Conclusion: What's Next in Space Cybersecurity?	241
11.1	Research Summary	242
11.2	Summary of Key Contributions	245
11.3	Future Work in Space Cyber-Security	248
11.4	Final Remarks	251
V	Appendices and References	253
	Appendices	
A	Chronology of Significant Satellite Hacking Incidents	257
B	GSExtract Implementation Details	277
C	QPEP Testbed Configurations	283
	References	285

x

List of Selected Abbreviations

ACK	Acknowledgment (in TCP messaging)
AFSPCMAN	Air Force Space Command Manual 91-710
AFTS	Autonomous Flight Termination System
AIS	Automatic Identification System
API	Application Programming Interface
ASAT	Anti-Satellite Weapon
CCSDS	Consultative Committee for Space Data Systems
CDS	CubeSat Design Specification
COTS	Commercial off-the-shelf
CSIS	Center for Strategic and International Studies
DITL	Day in the Life (testing)
DNS	Domain Name System
DVB-S	Digital Video Broadcasting - Satellite
ECDIS	Electronic Chart Display and Information System
ESA	European Space Agency
FEC	Forward Error Correction
FTP	File Transfer Protocol
GEO	Geostationary Earth Orbit
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GSE	Generic Stream Encapsulation
GSO	Geosynchronous Orbit
GTO	Geosynchronous Transfer Orbit
HTTP	Hypertext Transfer Protocol
ICMP	Internet Control Message Protocol

IMINT	Imagery Intelligence
ISP	Internet Service Provider
LEO	Low Earth Orbit
LV	Launch Vehicle
MAC	Media Access Control (address)
MPE	Multi-Protocol Encapsulation
MPEG	Moving Picture Experts Group (also, a data format)
NASA	National Aeronautics and Space Administration
NAT	Network Address Translation
NATO	North Atlantic Treaty Organization
PDU	Payload Data Unit
PEP	Performance Enhancing Proxy
PLR	Packet Loss Rate
PLT	Page Load Time
POP3	Post Office Protocol Version 3
QPEP	QUIC Performance Enhancing Proxy
QUIC	A transport layer protocol (not an acronym)
RCMA	Recognize, Connect, Motivate, Adapt (research process)
RF	Radio Frequency
RFI	Radio Frequency Interference
RSO	Resident Space Object
RTT	Round Trip Time
SDR	Software Defined Radio
SINR	Signal to Interference plus Noise Ratio
SSA	Space Situational Awareness
SSL	Secure Sockets Layer (protocol)
SSN	Space Surveillance Network
SYN	Synchronization (in TCP messaging)
TLE	Two-Line Element Set
TLS	Transport Layer Security (protocol)
TS	Transport Stream (as in MPEG-TS)
VPN	Virtual Private Network
VSAT	Very Small Aperture Terminal

Part I

A Case for Space Security

Reader, can you suggest whence the story begins?

*The narration may border on the limits of incoherency
and triviality, but it possesses considerable zest.*

— Cao Xueqin, *Hongloumeng* (trans. Joly)

1

Introduction

Contents

1.1	An Emergent Need	3
1.2	Structure and Methods	5
1.3	Key Contributions and Impact	9
1.3.1	Radio Link Security Contributions	9
1.3.2	Contributions Beyond Broadband	11
1.4	Publications and Outputs	12
1.5	Research in Context	15

1.1 An Emergent Need

Space is changing. A coalescence of innovations in computing, launch systems, and remote sensing have brought tremendous opportunities. Over the next decade, we can expect an order-of-magnitude increase in the number of operational satellites in orbit seeking to reap these benefits [1]. As the scale and substance of space missions evolves, ensuring that such systems are secure against cyber-attack will prove vital to sustained development.

To date, the security properties of space systems are not well understood. Although digital attacks against satellites have been a recognized threat for more than 50 years, the high cost and complexity of such attacks has reduced the

perceived urgency of these risks. What expertise does exist is obfuscated by heavy government classification and proprietary commercial practices. As attacker capabilities and knowledge continue to evolve, the space community will need to go beyond mere obfuscation and incorporate security into the design and operational practice of their missions.

The primary objective of this thesis is cultivating foundations for open and accessible space cyber-security research. In contrast with the broader space domain, where open academic research has historically played a key role in facilitating public and private sector innovation, little published work exists concerning space systems security. Through both direct and methodological contributions, this thesis presents the case for sustained systems security research in the domain.

Within this broad mandate, we begin with an in-depth security analysis of long-range satellite broadband services. This dominant application of commercial space platforms is a key driver for next-generation space missions and a critical infrastructure for billions of people around the world. In studying the topic, we identify severe security vulnerabilities which have arisen as unintended consequences of the physical constraints of space platforms. We further invent novel techniques to overcome these factors and adapt terrestrial communications-security best practices to the unique requirements of space.

Next, we generalize the lessons learned from this analysis to the broader subject of space systems security. In the process of doing so, we discuss diverse topics ranging from rocket launches to radar astrometry. Throughout, we see how the demanding physical requirements of space systems can give rise to domain-specific technical adaptations with easily overlooked security implications. Moreover, we show how security practices can be effectively adapted to overcome these vulnerabilities. By the end of this thesis, the reader will not only have a deeper understanding of several technical challenges in space security, but they will also be equipped with a general approach for identifying novel security problems at the intersection of outer and cyber space.

1.2 Structure and Methods

The thesis begins with a systematization of the limited, but rigorous, body of existing academic knowledge on the topic (Chapter 2). Due to the relative paucity of publications, we supplement this literature base with our own original archival research, building a chronology of more than 100 significant satellite security incidents spanning the past 60 years of spaceflight. In doing so, we build strong empirical foundations for understanding domain norms and practices. This is presented as a taxonomy of three distinct technical sub-topics: radio-link security, ground system security, and space platform security.

In Part II: *Threats and Defenses in Satellite Broadband* we explore the first of those sub-topics: radio-link security. Our analysis of long-range satellite broadband security in this section makes up the bulk of the thesis, consisting of deep technical analysis into real-world systems. In it, we evaluate the capability of low-resourced cyber-adversaries, such as individuals and non-state actors, to engage in wireless attacks targeting customers of modern satellite broadband networks. Our experiments demonstrate that attackers today can access and manipulate deeply sensitive information transmitted by dozens of distinct satellite internet service providers (ISPs). Further, we show that these vulnerabilities have direct implications for the safety and security of a wide range of satellite-dependent systems — such as electrical grids and maritime vessels.

In the process of identifying these vulnerabilities, we draw out a number of underlying physical properties of satellite broadband which make it particularly challenging to defend against such attacks without performance degradation. On the basis of this analysis, we develop a novel hybrid between traditional virtual private network (VPN) tools and satellite performance enhancing proxies (PEPs). We validate its performance and security characteristics through simulations in a replicable testbed environment. The tool is published as freely available open-source software designed to empower individual satellite customers to protect their traffic without ISP involvement.




Figure 1.1: The RCMA method for applied space systems security research.

In reflecting on Part II, we will note aspects of our satellite broadband security analysis with broader utility to space security research. In particular, we see how physically-driven divergences between space and terrestrial technologies give rise to novel security lacunae. By rooting our analysis of space security in these cyber-physical dynamics, we can uncover domain-specific issues that may otherwise be overlooked.

We present a four-step research method, represented by the acronym *RCMA* and summarized in Figure 1.1. The purpose of RCMA is to facilitate the discovery and resolution of unique cyber-security challenges in space missions.

In the first stage, *Recognize*, the researcher isolates a physical aspect of space operations which has a direct effect on either the design or capabilities of domain technology. For example, the remoteness of orbit means that physical access to operating satellites is cost-prohibitive.

In the second stage, *Connect*, the researcher identifies how these physical effects may implicate standard security practices. In our previous example, the lack

of physical access means that any digital forensics conducted as part of security incident response requires some degree of trust in the claims made by software running on a compromised device.

In the third stage, *Motivate*, the researcher grounds these theoretical security implications in terms of the real-world operation of space systems. For instance, they may demonstrate an exploit in a common satellite real-time operating system (RTOS) that permits tampering with on-board audit logs, making their compromise of the satellite functionally undetectable.

In the fourth and final stage, *Adapt*, the researcher considers how existing research/practice around systems security can be effectively adapted or extended to mitigate these threats. In our scenario, they may consider adapting write-once read-many (WORM) file systems to the limited storage capacity of satellites or blockchain-based audit-logging to the reduced computation capabilities of satellites.

The RCMA method arose from post hoc analysis of the research in Part II and is presented as a helpful reframing of traditional hypothesis-testing approaches to security research. A demonstrative mapping of the various chapters in Part II to the RCMA method appears in Table 1.1. One benefit of RCMA is that it explicitly encourages researchers to consider cross-domain linkages. RCMA draws on aerospace perspectives in the first and third steps and information security perspectives in the second and fourth. This helps prevent the premature dismissal of topics which appear trivial from any one perspective but are complex in their interaction.

Rather than merely asserting the utility of RCMA, Part III: *Cyber-physical Threats Beyond the Signal* applies it to the two other space-systems sub-domains from Chapter 2's taxonomy. This is done through two smaller experimental analyses, one focused on critical data used by satellite ground systems and the other on a threat involving compromised satellite platforms during rocket launches.

The ground systems study (Chapter 9) starts by *recognizing* a number of unique physical challenges involved in tracking resident space objects (RSOs), such as satellites and pieces of debris. We *connect* these barriers to the security implications of trust-based dependence on centralized space-surveillance information sharing

Table 1.1: Demonstrative Mapping of RCMA Method to Part II: *Threats and Defenses in Satellite Broadband*.

Research Step	Example from Part II
Recognize physical aspects of space technologies which differ from comparable terrestrial systems.	Vast transmission distances mean that geostationary broadband has extremely high latency and broad signal footprints. (Ch. 4)
Connect these dynamics to their implications for traditional security approaches.	Broad signal footprints expose data to long-range wireless eavesdropping attacks and high latency has led to ISP adoption of PEP TCP accelerators which are incompatible with customer-operated VPN software. (Ch. 4-5).
Motivate the need for security improvements by demonstrating these impacts in a domain-specific and realistic context.	Low VPN adoption in real-world satellite broadband networks exposes sensitive data to long-range eavesdropping attacks, harming safety and security for critical infrastructure, maritime and aviation customers. (Ch. 4-5).
Adapt proven security approaches to better account for these domain-specific requirements.	We build and evaluate a novel consumer-oriented VPN/PEP hybrid which combines security with necessary TCP optimizations for satellite broadband communications. (Ch. 6).

systems in the status quo. Next, we *motivate* the need for improvements by developing cross-disciplinary threat models at the intersection of international relations theory, astrophysics, and systems security that abuse this trust. Finally, we *adapt* anomaly detection techniques from problem domains like network intrusion detection to design a system that helps bolster the integrity of space situational awareness data.

In the space-systems study (Chapter 10), we *recognize* how the physical demands of space launches have led to inexpensive secondary payloads, known as “CubeSats,” riding into orbit on the same rockets as critical primary payloads. We *connect* this to security by characterizing the ways in which physical redundancy and safety regulations for CubeSat integration assume non-maliciousness. We *motivate* the need for improvements through an adversarial analysis of these standards and a series of dynamic physical simulations demonstrating the physical viability of space-to-space radio frequency interference (RFI) attacks from an outwardly compliant but

functionally malicious CubeSat during launch. Finally, we propose initial policy steps to *adapt* existing CubeSat safety regulations to encapsulate adversarial security risks.

In Part IV, the thesis concludes by synthesizing lessons from across these experiments and highlighting their implications for further research in space cybersecurity. In addition to highlighting direct technical contributions from this research which may be of practical use to space operators, we also present the case that this work can be used as a launchpad for continued applied security research in an unusual domain.

1.3 Key Contributions and Impact

At a high level, the thesis proposes and evaluates an approach (the RCMA method) for conducting applied security research in the space domain. Our method is validated through real-world experiments, each of which also makes direct technical contributions to status quo understandings of space security. The intention is to provide a useful starting point for others interested in identifying and remediating novel security vulnerabilities in space systems.

In terms of direct technical contributions, we outline a number of previously unknown or under-studied threats to modern space missions. Rather than focus on individual pieces of hardware or software, we prioritize systemic issues impacting hundreds of missions. These threats are validated experimentally, contributing practical technical evidence to a topic that has been, to date, largely conceptual. Rather than simply raise problems, we further develop and evaluate defenses against these attacks. While academic novelty is one motivation, throughout we prioritize experiments and approaches which also have intuitive practical relevance to space operators.

1.3.1 Radio Link Security Contributions

The satellite communications work in this thesis has raised cross-industry awareness of critical security and privacy vulnerabilities. In direct response to these experiments, the United States Federal Bureau of Investigation (FBI) and the United

States Coast Guard have issued emergency notifications warning maritime vessels of significant risks to the security of their satellite networks. Through responsible disclosure, we have also informed several of the world's largest companies of the potential harms of insecure satellite broadband practices.

At first glance, the experiments presented in Chapter 4 may seem banal, even in light of its significant real-world impact. After all, the security and privacy risks of unencrypted radio communications links are both intuitive and well understood. However, as we delve deeper into more modern communications schemes (as in Chapter 5), we will develop a more nuanced perspective on the underlying complexity of implementing both these attacks and defenses against them.

This thesis demonstrates a substantial change in threat model, showing how sophisticated satellite eavesdropping capabilities can be achieved using inexpensive home television equipment. In doing so, we re-frame an historically abstract threat from nation-state agencies into a real and present danger impacting satellite customers. To quote one responsible disclosure contact on learning of these findings: “Your research just emphasizes the point to always revisit and recalculate your risk analyses. When we started the assumption was that it would take \$x00k to set up a listener on the ground. You proved it’s now 1/1000 of that today.”

These attacks are non-trivial and require more than simply purchasing the correct hardware. Modern satellite broadband leverages complex modulation and encoding schemes which are beyond the assumed physical reception capabilities of our equipment. In Chapter 5, we present a technique for overcoming these barriers through forensic reconstruction of corrupted partial IP packets in radio signal recordings. This offers as an independent contribution in demonstrating an easily overlooked characteristic of communications threat modeling: that attackers do not necessarily require comparable physical capabilities to legitimate users.

In addition to characterizing vulnerabilities, we assess their impact on real-world users. By analyzing more than a dozen different satellites in geostationary earth orbit (GEO), Chapters 4 and 5 contextualize these vulnerabilities as endemic to long-range satellite broadband services — rather than quirks of one specific ISP.

In Chapter 4, we discuss the impact of the threat model on home internet users and critical infrastructure operators. In Chapter 5, we show how the eavesdropping threat intersects with operational technology for maritime customers, with severe implications for the security and safety of ship and crew.

Finally, we contribute solutions. Chapter 6 outlines the underlying physical, commercial, and technical drivers of inadequate encryption in these networks. Our direct experiences with responsible disclosure help us to identify critical “pain points” for end-to-end GEO broadband encryption. This motivates the invention of a hybrid satellite Performance Enhancing Proxy (PEP) and customer-oriented VPN. Our open-source software implementation of this approach, called *QPEP*, is one obvious contribution, offering the first non-proprietary and verifiable tool of its kind. In the effort of building and evaluating QPEP, we also contribute a general purpose experimental framework and simulation software testbed to help other researchers both in replicating our findings and testing their own ideas.

1.3.2 Contributions Beyond Broadband

The remainder of the thesis builds upon the lessons learned in Part II to explore topics outside of space communications. This serves dual purposes: first, to validate the RCMA approach and second, to contribute novel technical insights on space systems security.

In the context of ground systems security, we present a previously unconsidered threat model to space missions in the forms of stealthy attacks on Space Situational Awareness (SSA) databases. These attacks leverage orbital mechanics and physical constraints on sensor capabilities to deceive satellite operators — the vast majority of whom lack the capacity to verify third party SSA data.

While the threat model presented in Chapter 9 is itself novel, our explicitly interdisciplinary approach further demonstrates how broader perspectives on inter-state dynamics in space can elevate a relatively banal technical attack (database exploitation) into a much more sophisticated vector (SSA deception). We leverage

technical research methods to bring a compelling source of evidence to bear on political debates concerning space surveillance priorities.

Additionally, we develop new mitigations to SSA deception attacks — adapting traditional machine-learning based anomaly detection strategies to orbital dynamics. In doing so, we provide a replicable method for weaker nation-state space powers and even individual satellite operators to verify third-party SSA claims and reliably detect deception attempts without the use of a single telescope.

In Chapter 10 we take a look at satellite platforms themselves, developing, to our knowledge, the first technical model of space-to-space cyber-mediated attacks and attacks targeting launch operations. Our cross-disciplinary adversarial analysis of safety engineering standards for Cube Satellites allows us to identify exploitable coverage gaps at the intersection of safety and security. The recommendations made at the conclusion of Chapter 10 have particular pertinence as a contribution to ongoing policy debates regarding the relevance of cyber-security certification for low-capability secondary payloads. For academic researchers, our explicit treatment of the challenges involved in conducting open security research in a highly regulated subject area (rocketry) may be of general interest. By leveraging model-based engineering techniques, we demonstrate steps for evaluating threats with only limited public information.

1.4 Publications and Outputs

The majority of contributions made in this thesis have been presented in peer-reviewed publications or are currently pending peer review. Several individual case studies and chapters within are drawn directly from these publications, albeit with some textual modifications appropriate to the style, format, and broader remit of this thesis.

In all of these publications, I was the first author. I was responsible for the direct research and writing components. Other authors provided greatly appreciated logistical, editorial, and supervisory support.

A summary of these publications and their appearance in the thesis is as follows:

- James Pavur and Ivan Martinovic. Building a Launchpad for Satellite Cyber-Security Research: Lessons from 60 Years of Spaceflight. (*Under Review*). Comprises Chapter 2.
- James Pavur, Daniel Moser, Vincent Lenders, and Ivan Martinovic. Secrets in the Sky: On Privacy and Infrastructure Security in DVB-S Satellite Broadband. (*Appears in: ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec), 2019. pp. 277-284*). Comprises Chapter 4.
- James Pavur, Daniel Moser, Martin Strohmeier, Vincent Lenders and Ivan Martinovic. A Tale of Sea and Sky On the Security of Maritime VSAT Communications. (*Appears in: IEEE Symposium on Security and Privacy (S&P), 2020*). Comprises Chapter 5.
- James Pavur, Martin Strohmeier, Vincent Lenders and Ivan Martinovic. QPEP: An Actionable Approach to Secure and Performant Broadband From Geostationary Orbit. (*Appears in: Network and Distributed System Security Symposium (NDSS), 2021*). Comprises Chapter 6.
- James Pavur and Ivan Martinovic. The Cyber-ASAT: On the Impact of Cyber Weapons in Outer Space. (*Appears in: 11th International Conference on Cyber Conflict (CyCon), 2019.*). Comprises a portion of Chapter 9.
- James Pavur and Ivan Martinovic. Lost in Space: Detecting Deception in Space Situational Awareness. (*Appears in: ACM Asia Conference on Computer and Communications Security (Asia CCS), 2021*). Comprises a portion of Chapter 9.
- James Pavur, Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. In the Same Boat: On Small Satellites, Big Rockets, and Cyber-Trust. (*Appears in: 13th International Conference on Cyber Conflict (CyCon), 2021.*). Comprises a portion of Chapter 10.

Beyond the publications which directly appear in this thesis and associated conference presentations, I have engaged in related research-outreach activities beyond the academic systems security community. Some notable instances of these activities include:

- Briefings on the work contained in Chapters 4, 5, and 6 at the Black Hat USA 2020 and DEFCON 28 security conferences.
- Briefing on the work contained in Chapter 9 at the DEFCON Aerospace Village, 2020.
- Coverage in various popular press articles/interviews on satellite cyber-security, including *Forbes*, *Ars Technica*, *Heise*, and others.
- Open source publication and maintenance of open source security software developed in Chapters 5 & 6.
- Several invited talks, including:
 - Briefing at DEFCON’s China Party 2021 virtual event (2021)
 - Presentation on space security for the Cybersecurity Oversight Office at the UK Civil Aviation Authority (2021)
 - Webinar on space security and the internet of things with the Space Generation Advisory Council (2020)
 - Presentation on satellite broadband at the Cambridge Cyber Risk Conference hosted by the Judge Business School (2020)
 - Presentation to students at Monash University — Malaysia (2020)
 - Keynote at CySat 2021 hosted by CYSEC SA (2021)
 - Panelist at CENSIS Cybersecurity Workshop on Remote and Hazardous Environments (2020)
 - Panelist at Connected Aviation Intelligence Inflight Hacking Prevention Roundtable (2021)
- Assisted in the development of content for *Mission Alenium*, a virtual introduction to satellite cyber-security led by the California Cybersecurity Institute and targeted towards high-school students.
- Featured in a short documentary by *Tomorrow Unlocked* on the intersection between space and cyber-security.

1.5 Research in Context

By the end of this thesis, the reader will have been exposed to an unusual and evolving niche in systems security. Through diverse case studies, this thesis will show how existing cyber-security approaches can be effectively adapted to the requirements of a novel domain. The focus throughout is on the challenges presented by the unique physical and logistical characteristics of space systems. While each case study makes specific technical contributions which directly bolster the security of space missions, the broader intention is to demonstrate a process by which cyber-security research can interface with space.

Much work remains to be done in order to ensure that space technology remains trustworthy and to fully realize the magical potential of space exploration and development. This thesis aspires to make a small, but significant, first step towards security beyond the Earth's mesosphere.

I insist [...] that there be altogether only three ways of fighting. Sand bombs, scientific wrestling and fencing with spears.

—Ferenc Molnár, *The Paul Street Boys* (trans. Rittenberg)

2

Building on 60 Years of Space Security Knowledge

Contents

2.1	Threat Modeling in Context	19
2.1.1	The Rise of Satellites	19
2.1.2	Emerging Threat Landscape	20
2.1.3	Threat Classes	22
2.1.4	Unique Technical Security Challenges	24
2.2	Learning from History	27
2.2.1	1957-1979: Early Days	29
2.2.2	1980-1989: Piracy and Spoofing	29
2.2.3	1990-1999: Broadcast and Flight Control Systems	30
2.2.4	2000-2009: Organized Attackers	31
2.2.5	2010-Present: Evolving Threats	32
2.2.6	General Trends and Developments	33
2.3	Defending The Signal	34
2.3.1	Eavesdropping Attacks	35
2.3.2	Signal Injection Attacks	38
2.3.3	Signal Spoofing Attacks	39
2.3.4	Future Directions in Space Signals Security	40
2.4	Defending Space Platforms	42
2.5	Defending Satellite Ground Systems	45
2.6	Holistic Security Models	48
2.6.1	Operational Frameworks	48
2.6.2	Policy and Legislative Frameworks	50
2.7	Lessons and Opportunities	51

From the launch of Sputnik in October 1957, space technology has played a critical role in the emergence of the information age. Today, satellites are far more than simple scientific demonstrations, instead underpinning essential services that define our lives. As the satellite industry undergoes a market renaissance driven by miniaturization and declining launch costs, defending these systems against cyber-attacks will only increase in importance.

Today, satellite cyber-security is a disparate and ill-defined topic, with critical contributions scattered across disciplines ranging from history and security studies to aerospace engineering and astrophysics. This chapter seeks to distill these interdisciplinary contributions and systematize status quo knowledge on security for space systems.

The process begins with threat modeling — unifying dozens of prior efforts to characterize threats to space systems into a single matrix linking attackers, vulnerabilities, and motivations. This model is supported with an exhaustive historical time of satellite incidents, found in Appendix A, where our own archival research is added to contributions from Fritz and Manulis et al. [2, 3]. The end result is an empirical and evidenced foundation for those making the case for space systems security research.

We build on this foundation to propose a natural taxonomy of four sub-domain topic areas: RF-link security, space platform security, ground systems security, and mission operations security. For each, we apply our threat modeling process to recent technical and academic developments, helping to uncover unsolved questions relevant to the safety and security of space missions. This includes the explicit presentation of promising research directions in each sub-domain. We use this analysis not only to motivate the technical research in this thesis, but also as a launchpad for future work in the domain.

2.1 Threat Modeling in Context

A robust understanding of the means and motivations of attackers is a key starting point for understanding the security requirements for space systems. As such, we can begin by contextualizing previous work into a high-level model of threat actors and their preferred techniques for targeting space systems. To do so, we will consider the state of the modern space industry, the historical behaviors of threat actors in the domain, and emerging technological shifts in space mission design.

2.1.1 The Rise of Satellites

Today, more than 2,000 operational satellites orbit Earth, supporting a market worth more than \$150 billion annually [4, 5]. They underpin a wide range of vital services, including: more than 10 TB/s of global internet capacity, media broadcasts to over 100 million customers, terabytes of daily earth observation data, and precise global positioning services [5]. Their importance will only increase. By 2035, satellite broadband is anticipated to exceed 100 TB/s globally and the direct industry value will exceed half a trillion dollars annually [5].

Around 40% of operational satellites are used for business communications and 30% support a mix of civilian and military government operations, with the remainder largely dedicated to mixed-use remote-sensing, meteorological, and navigational missions [6]. However, this balance will likely shift in response to demand for ubiquitous broadband service and remote sensing capacity. The emerging sector rising to meet this demand is widely referred to as “New Space” [7]. Among the most prominent New Space missions are mega-constellations proposed by organizations like Blue Origin, SpaceX, and OneWeb. If successful, these projects will increase the number of Low Earth Orbit (LEO) satellites by at least an order of magnitude.

A key driver of these changes is diminishing launch costs. Modern launch vehicles have reduced the cost-per-kilogram to LEO to under \$2,000 [8]. This is radically more affordable than NASA’s famous shuttle missions (at around \$54,500), and almost 90% cheaper than the average cost of all missions from 1970-2000 (around

\$18,500) [9]. For the first time, the deployment of satellite payloads is within the means of a vast array of new industry entrants.

Concurrent improvements in computing capabilities — particularly with respect to miniaturization — have compounded these effects. As computer hardware grows smaller and less power-demanding, increasingly complex and light satellites become feasible. This has resulted in the emergence of “small satellites” — a wide range of sub-500 kg devices, with many weighing less than 1 kg.

The emergence of commercial off-the-shelf (COTS) satellite components has further driven growth. The availability of ready-made satellite flight hardware decreases procurement costs, allowing New Space entrants to accept larger technical and commercial risks. Indeed, it is now possible to purchase a fully assembled 1 kg “Cube Satellite” for as little as \$16,000 [10].

2.1.2 Emerging Threat Landscape

As the usage of space assets grows, the threat environment they face has shifted. Historically, satellites have benefited from a sort of “security through obscurity” whereby system complexity and equipment costs dissuaded all but the most sophisticated adversaries. The combined effects of COTS components and constellations with thousands of identical satellites mean that this diversity and complexity of implementation is unlikely to endure.

Putting cyber-security aside for a moment, general motivations for harming satellites are well studied and intuitive. In a military context, space systems underpin Command, Control, Communications, Computer, Intelligence, Surveillance, and Reconnaissance (C4ISR) capabilities [11, 12]. Adversaries seeking to “level the playing field” against great powers have strong incentives to undermine these capabilities by harming space systems [13]. Further, civil society depends on satellites for essential navigational, communications, and meteorological services. Attackers motivated by societal disruption may view satellites as attractive “single points of failure” in critical infrastructures [14].

Table 2.1: Overview of Threat Actors Proposed in Literature.

Attacker Type	Example Motivations	Capabilities & Resources	Selected References
Military	- Space Control - ASAT Weapon	Very High	[5, 13, 15, 16]
Intelligence	- Counter-Intelligence - Technology Theft - Eavesdropping	Very High	[16]
Industry Insider	- Sabotage - Technology Theft	High	[16, 17]
Part Supplier	Sabotage	Moderate	[18, 19]
Organized Crime	- Eavesdropping - Ransom	Moderate	[5, 16]
Terrorist / Militant Org.	- ASAT Weapon - Message Broadcast - Notoriety	Low to Moderate	[5, 15]
Commercial Competitors	- Sabotage - Technology Theft	Low	[16]
Individual Hackers	- Notoriety - Personal Challenge	Very Low	[5, 15, 20]
Political Activists	Message Broadcast	Very Low	[15, 16]

Regarding potential attackers, a 2016 report by Chatham House, a prominent UK policy think-tank, taxonomizes threat actors into four broad categories: states seeking military advantage, organized criminal efforts for financial gain, terrorist groups seeking recognition, and individual hackers proving their skills [5]. This can be supplemented with the list of threat actors published by the Consultative Committee for Space Data Systems (CCSDS) [16]. CCSDS represents a consortium of national space agencies from eleven member states and thirty-two observer nations and is one of the most influential technical bodies for the development of space protocol and systems standards. Beyond overlaps with Chatham House's model, CCSDS adds: foreign intelligence services, political activists, commercial competitors, agency insiders, and business partners [16]. Independent authors within the military strategy and civil space science domains have further suggested supply-chain threats from equipment manufacturers [18, 19]. Table 2.1 offers a

composite summary of threat actors from these and other reports as a starting point for the development of threat models [15, 17].

It is worth noting that our research has been restricted to English-language resources, which tend to show a western bias in threat. For example, the Center for Strategic and International Studies (CSIS), a Washington DC political and security think-tank, isolates four main state belligerents in orbit: Russia, China, Iran and North Korea [15]. In terms of cyber-security, CSIS contends that Russian cyber-capabilities against satellites are particularly sophisticated and have been demonstrated in historical attacks on critical infrastructure and space systems. With respect to China, CSIS highlights the fact that the People's Liberation Army Strategic Support Force (SSF) has organizational responsibility over both China's counterspace weapons and offensive cyber operations — creating natural cross-over opportunities. Other sources note that Chinese military reports have explicitly advocated for the use of digital counterspace against US space assets [17]. Less information is provided to motivate the Iranian and North Korean threat, but CSIS notes sustained Iranian interest in cyber-attacks against the related ballistic missile defense (BMD) domain and North Korean cyber-attacks against terrestrial critical infrastructure. Very few English-language sources offer deep threat assessments of US and EU offensive capabilities in space, but it is perhaps not unreasonable to assume similar means and motivations.

2.1.3 Threat Classes

In addition to understanding *who* might be interested in harming satellites, it is important to consider *how* they might go about doing so. A high-level starting point can be found in the security studies and international relations fields, where scenario modeling is a common component of strategic analysis. Chatham House groups cyber-attacks to satellites into two broad categories: attacks which target satellites themselves (e.g., via control system exploitation) and attacks which target satellite ground stations (e.g., via traditional network intrusion) [5]. The European Space Agency (ESA) brings civilian governmental perspectives, outlining additional

threats to scientific missions including signal intercept and jamming, denial of service attacks, and supply chain malware [19]. Further technical specifics can be gleaned from CCSDS, which adds replay attacks, access-control failures, social engineering, data corruption, and meta-data analysis on encrypted traffic [16]. In their research on the intersection between space and military law, Rendleman and Ryals raise the novel additional threat of satellite hijackers who steal orbiting satellites to bolster their own space capabilities [21]. Finally, multiple researchers across the communications and systems security domain have considered the threat of signal piracy and spoofing attacks [2, 18, 22].




Figure 2.1: A Subsystem Approach To Satellite Threat Classification.

We can bring structure to this diverse array of perspectives by expanding Chatham House’s systemic taxonomy. Specifically, we propose three broad categories of attack surfaces: those relating to satellite signals, those relating to the space-platform, and those targeting satellite ground systems (Figure 2.1). Some alternative taxonomies further divide ground systems into “Customer” and “Mission” segments, but technical threats are fairly similar for both applications [3].

While our classification system still results in some overlap, such as the case where RF-signals are used to send malicious flight commands, it has two key

benefits. First, it aligns closely with common organization paradigms around space missions. Satellite missions are multi-stakeholder processes, where distinct organizations are often responsible for the on-orbit operations, communications, and ground segments. By mapping vulnerabilities to these domains, we can better clarify which organization has responsibility for defending against which threats. A second benefit of this approach is that the technical skills required for systems security research in each domain are intuitively distinct: on-orbit defenses draw from embedded and control systems topics; signals defense requires networking and radio expertise; and ground systems leverage traditional operational technology (OT) and information technology (IT) perspectives.

We further can combine this subsystem taxonomy with the actors outlined in Section 2.1.2. When combined with the historical analysis in Section 2.2, and the threat mapping from Table 2.1 we can produce a matrix associating vulnerabilities with technical means, attacker capabilities, and empirical context, as shown in Table 2.2.

2.1.4 Unique Technical Security Challenges

A superficial reading of these threats may suggest that satellites pose few novel challenges for systems security researchers. After all, terrestrial variants of all the listed threats easily come to mind. Indeed, many researchers — especially from the commercial space sector — contend that traditional IT security approaches offer sufficient coverage, advocating for the use of NIST controls and generic security information and event management (SIEM) tools [23–26]. This viewpoint is commercially appealing as it allows for the direct use of widely available security tools (and cross-domain hire of experts in those tools) as the main line of defense for satellite missions [27].

This viewpoint is not without detractors. Byrne et al., speaking primarily from the perspective of aerospace academia, argue that “the assertion that existing controls will protect against risk is sometimes accepted without reasonable supporting data or, even worse, is accepted where the lack of data is used as proof” [28].

Table 2.2: Satellite Threat Matrix.

Threat Type (From Figure 2.1)	Posited In	Example Attack Scenario	Relevant Subsystems (From Figure 2.1)	Empirical Examples	Sophistication (Refer to Table 2.1)
Denial of Service	[19]	Force satellite to enter "Safe Mode"	Payload	None to date	Very High
Hardware Backdoor	[19]	Inject malicious commands on hardware bus	Payload Ground	None to date	Very High
Privilege Escalation	[29]	Send flight control commands from payload software application	Payload	None to date	Very High
Bespoke Malware	[15]	Exploit vulnerability in satellite firmware or ground telemetry software	Payload Ground	[30, 31]	Very High
Payload Hijacking	[21]	Maneuver satellite to undermine sensor readings	Payload	Possibly: [32, 33]	Very High
Sensor Injection	[34]	Blind imagery sensors with long-range laser signals	Payload	[35]	High
Jamming	[19]	Block satellite phone reception in remote conflict zone	Signal	[36]	Low-Moderate
Eavesdropping	[16]	Intercept sensitive internet traffic from satellite signals	Signal	[37]	Low
Metadata Analysis	[16]	Identify classified satellite based on radio spectrum behavior	Signal	[38]	Low-Moderate
Replay Attack	[16]	Re-issue intercepted commands to cause harmful maneuver	Signal	None to date	Moderate-High
Signal Injection / Hijacking	[2, 18, 22]	Overwrite legitimate signal with falsified broadcast	Signal	[39]	Low-Moderate
Generic Malware	[16, 19]	Compromise space-related system with generic ransomware	Ground Payload	[40]	Low
Social Engineering	[16]	Phishing campaign used to access satellite design documents	Ground	[32]	Very Low
Physical Access	[5, 41]	Theft of laptop w/ flight software	Ground	[42, 43]	Low-Moderate
Data Corruption	[16]	Damaging stored imagery data to prevent intelligence use	Ground	None to date	Low

Falco, a computer science academic, takes this further, arguing that attempts to map traditional IT security to the space domain has created harmful technical knowledge gaps and discouraged specialization [14].

Falco further isolates six reasons that satellite cyber-security requires unique consideration [14]. First, satellites represent a single point of failure for other critical infrastructures, increasing the number and capabilities of attackers who may be interested in harming them beyond those obviously relevant to mission function. Second, there is comparatively little regulation guiding satellite cyber-security, creating uncertainty regarding the controls appropriate to a given system. Third, complicated supply chains both risk malicious backdoors and make it difficult to assign organizational responsibility for security. Fourth, the widespread use of COTS hardware integrated with bespoke systems creates a unique situation where vulnerabilities likely apply to many platforms, but patches may require bespoke platform-specific adaptations. Fifth, the specialized nature of aerospace means that few cyber-experts understand satellites sufficiently to adapt best practices to the domain. Finally, satellites are compute-constrained devices with limited resources, and security/performance trade-offs are more acute than in terrestrial systems.

The second point, regarding the shortcoming of existing regulatory standards, is further supported by Fidler, writing for the Council on Foreign Relations — an international relations policy think-tank [17]. In particular, he contends that mappings of IT standards to space systems amount to little more than “paper-shuffling” [17]. Bardin suggests that industry is unsure what would even constitute a cyber-attack against space systems due to lack of comprehensive threat modeling [41]. This may be attributable in part to overuse of the term “hacking” in media and policy circles to describe any disruption to satellite operations [2, 41]. For example, technical authors often treat radio jamming as an unrelated topic while policy analysts explicitly consider it a cyber-attack vector [5, 15, 19].

Falco’s third and fourth points, regarding supply chains, have been subject to some discussion as well. Space missions have uniquely complex bureaucratic structures. Many distinct organizations may share some device resources (e.g.,

communications systems), while operating others independently (e.g., on-board sensors). Excepting the largest players, satellite operators do not control the entire mission lifecycle. Launch vehicles, orbital injection, operation, and retirement are frequently handled by distinct entities. Some service providers (e.g., satellite television services) may have no ownership stake in the space platform at all, but instead simply lease radio access. The result is that operators cannot necessarily trust each other and may not share security priorities. Any given member of the mission ecosystem can potentially compromise others [19]. This threat is particularly acute for “New Space” systems, which rely heavily on third-party COTS equipment [18, 25].

Finally, Falco’s “expertise vacuum” is widely recognized as a significant barrier. Niche components of satellite systems lack direct terrestrial equivalents (e.g., star-trackers), impairing the development of a general body of knowledge for securing these devices [25]. In academic contexts, the cross-disciplinary mixture of engineering, astrophysics, computer science, and security studies complicates the search for appropriate venues and communities for publication and peer-review. For example, expertise in cryptography may not be directly useful without additional hardware and astrophysics knowledge, as extra-terrestrial radiation can induce random bit-flips in cryptographic key storage and requires special attention [44].

Ultimately, space systems appear to be much more than mere “computers in the sky.” Well-regarded terrestrial security practices often fail to transfer to space systems for unintuitive reasons. The result is that relatively little work, especially within systems security, has been conducted on space technologies.

2.2 Learning from History

Given the dearth of academic satellite cyber-security research, the threat may appear distant and hypothetical. Indeed, few satellite hacking incidents over the past half-century have received significant public attention, and one might be tempted to argue that satellite cyber-security is more an invented problem than present danger.

However, a deeper look at the history of operations targeting satellites reveals an unconventional but voluminous body of knowledge. Indeed, cyber-attacks against




Figure 2.2: Countries Involved in Satellite Hacking by Year of First Entry. This chart includes countries with governments and/or citizens implicated as attackers in satellite security incidents.

satellite systems have been occurring, almost unnoticed, for decades — perpetrated by attackers from across the globe (Figure 2.2). In this section, we present an overview of empirical data with a focus on long-term trends and unsolved security problems. This analysis builds on the prior work of Fritz and Manulis et al. [2, 3].

In conjunction with this research, we have developed an annotated chronology which details more than 100 significant satellite hacking incidents from 1962 to present day (Appendix A). To our knowledge, this chronology represents the most exhaustive record of satellite hacking incidents to date. Derived from original archival research synthesizing unclassified primary and secondary source materials, it offers evidence-based technical insights into the evolution and practice of satellite exploitation.

Before delving into the survey, it is worth clarifying its scope. In particular, the topicality of RF interference has been subject to much debate — with some regarding it as an issue of electronic warfare as opposed to cyber-operations. We have elected to include some of the most notable instances of such attacks in our analysis for two reasons. First, a willingness to engage in jamming suggests that an attacker values the ability to “virtually” deny satellite access to their victims

— offering potential insights into threat models for digital counterspace. Second, the hardware and expertise involved in jamming operations often has significant cross-over with more obviously topical signal-hijacking and injection attacks.

2.2.1 1957-1979: Early Days

In the earliest days of human spaceflight, the principal information security concerns revolved around the ability of adversaries to compromise satellite flight control signals. One of the first public discussions of satellite information security was a 1962 US congressional hearing to determine if private companies should be allowed to operate in space [45]. It was suggested that commercial missions would be more vulnerable to jamming and replay attacks from Soviet adversaries, while higher-altitude military satellites were presumed secure due to the complexity of the requisite equipment.

The subsequent two decades saw no major satellite hacking incidents. However, an important political debate was brewing over satellite broadcast abuse. The US had begun transmitting anti-communist propaganda on satellite beams directed into Soviet territory. In response, the USSR put forward a UN proposal in 1972 asserting a sovereign right to jam illegal radio signals in their territory [46]. To this day, state sovereignty over radio emanations from foreign satellites remains contentious. Modern norms on interstate jamming and eavesdropping attacks can be readily traced back to this 1972 dispute.

2.2.2 1980-1989: Piracy and Spoofing

The first major satellite hacking incident is generally thought to have occurred in 1986. An industry insider and satellite-dish salesman pseudonymously dubbed “Captain Midnight” hijacked an HBO television broadcast destined for satellite TV customers in Florida and replaced it with a message chastising network executives for new signal-scrambling copy protection technology [47]. Interestingly, this attack almost exactly mirrored a fictional short-story from a satellite enthusiast magazine the previous year — although no formal association has been proven [48]. The next

year, a similar attack took place wherein an employee of the Christian Broadcasting Network replaced a satellite stream operated by The Playboy Channel with biblical verses chastising viewers for not attending church on Sunday [49].

1986 also marked the first major satellite eavesdropping case, wherein the government of Indonesia was accused by an American satellite imaging company of illegally intercepting earth observation data without paying for a subscription to the satellite's service [50].

Terrestrially, the 1980s marked the first major attack against satellite ground systems. In 1987, a group of West German teens compromised top secret NASA networks by means of a Trojan horse program which concealed a keylogger [51]. These networks were reported to include information on classified military space missions and to have the capability to cause direct harm to satellites. Upon intercepting a mail-box message indicating that the compromise had been discovered, the teenagers voluntarily turned themselves in.

2.2.3 1990-1999: Broadcast and Flight Control Systems

Both satellites usage and exploitation accelerated throughout the 1990s. As satellite television became commonplace, states began using jamming attacks to control the flow of information across their borders. Iran began jamming foreign satellite television stations in 1994, a practice which continues today [52, 53]. In 1998, Indonesia became the first country to deliberately use a satellite to jam signals from a neighboring satellite as part of a dispute with Hong Kong over orbital slot access [2, 54]. By the end of the 1990s, commercially available satellite jammers emerged on the market, including a \$4,000USD Russian-made device capable of disabling GPS signals over a 200 km radius [55].

The 1990s saw the widespread emergence of cryptographic systems for satellite television piracy — kicking off an ongoing battle between satellite pirates and media companies which began with simple smart-code sharing networks and escalated into sophisticated cryptanalysis [56, 57]. From 1993 onwards, reports detail an

essentially annual cycle of hackers breaking TV protections, media companies designing improvements, and governments making related arrests.

Finally, a number of attacks against satellite ground stations occurred over the 1990s. These included high-profile incidents where hackers claimed to have accessed systems which could allow the issuance of flight control commands to orbiting satellites. Most notable among these incidents is a 1998 scenario wherein hackers, widely believed to be Russian-government affiliated, gained access to flight control systems in NASA's Goddard Space Flight Center [2, 32]. During this incident, the German-US ROSAT x-ray telescope inexplicably altered its orientation to point optical sensors directly at the sun - leading to irreparable hardware damage [2, 32]. Although details surrounding the incidents are highly classified, this is often cited as the first cyber-attack which caused physical damage in orbit.

2.2.4 2000-2009: Organized Attackers

The 2000s saw more incidents than the previous forty years combined. One major trend was the emergence of organized non-state attackers. Notable incidents of this nature included signal hijacking attacks by Falun Gong (a Chinese religious and protest movement) from 2002-2005, similar attacks by the Tamil Tigers (a Sri Lankan militant organization) from 2007-2009, and eavesdropping attacks compromising US military drone video feeds by Iraqi insurgents in 2009 [58–62].

Government-led jamming operations continued unabated. Most notable among these were an instance of Iranian jamming of signals directed to Turkey in 2000 and Cuban jamming of signals destined for the Middle East in 2003 [63, 64].

Significant attacks against ground stations during this period include complete flight control takeover of two NASA satellites in 2007 and 2008 [33, 41, 65]. These attacks were originally reported as signal jamming but later linked to a Chinese government compromise of NASA ground stations [65].

The 2000s also saw the first public case of a malware infection in orbit. In 2008, a Russian cosmonaut introduced Windows-XP malware to systems aboard

the International Space Station (ISS). This incident is widely believed to have been accidental [33, 40, 66].

Although not directly related to cyber-security, a major space security incident occurred in January 2007 when China demonstrated an anti-satellite weapon (ASAT) [67]. Not only did this generate a significant amount of space debris, it also demonstrated emerging state interest in offensive counterspace technology. This ASAT demonstration was preceded by a less well-known “virtual” attack in 2006, when a Chinese ground-based laser system was used to blind sensors aboard a classified US military satellite [35].

2.2.5 2010-Present: Evolving Threats

The accelerating usage of cyber-operations in space has continued over the most recent decade. In particular, a wave of jamming incidents in the Middle East and North Africa were kicked off by the Arab Spring protest movements in 2010 and have continued thereafter. This caused the list of countries with demonstrated satellite jamming capabilities to more than double with the addition of Egypt, Jordan, Bahrain, Ethiopia, Saudi Arabia, Eritrea, Syria, Azerbaijan, and Israel — along with renewed jamming from Libya and Iran [68–79]. Outside of the region, North Korea also began a sustained jamming campaign against South Korean military GPS in 2010 [80].

More sophisticated signal-related attacks also emerged. This included signal intrusion attacks by Hamas against Israeli news stations and academic research demonstrating weaknesses in satellite data services [39, 81–85]. In 2014, Russia was accused of launching a “stalker sat” which followed other satellites in orbit to intercept uplink signals, representing the first publicly acknowledged instance of satellite-to-satellite eavesdropping [86].

Attacks against ground stations and satellite control systems grew more sophisticated as well, with many being linked to state actors. In particular, China has been accused of compromising US space control systems in 2011, 2014, and 2017 [87–90]. This is perhaps unsurprising given that, in 2014, an internal US audit

of the Joint Polar Satellite System (JPSS) ground stations found more than 9,000 “high-risk” security issues, many of which remained unpatched from prior audits [91]. Commercial ground systems were also demonstrated to have severe vulnerabilities, including many hard-coded passwords and backdoors [92, 93].

This period has also seen the first organized criminal abuse of satellite systems. In 2015, the Russian advanced persistent threat (APT) actor dubbed “Turla group” was found to be abusing satellite internet signals to anonymously exfiltrate data from compromised computer systems [94]. In Chapters 4 and 5 of this thesis, we demonstrate techniques which may have been used by this APT actor to implement such attacks.

New attention has recently been paid to the satellite cyber-security field. In 2020, the US Air Force hosted an online “Hack-A-Sat” competition which explicitly sought to introduce cyber-security professionals to the world of satellite cyber-security and to uncover vulnerabilities in real space systems [95]. Similarly, in 2020, DEFCON hosted its first “aerospace village,” a sub-conference which included a briefings track focused exclusively on space systems security [96].

2.2.6 General Trends and Developments

In sum, there has been a clear general trend towards increased use of cyber-capabilities that target satellite systems (Figure 2.3). Over the past 60 years, and especially over the past 20, the number of actors willing and able to attack satellites in cyberspace has increased dramatically.

Today, almost 30 states have demonstrated some degree of cyber-offensive counterspace capabilities, including many which lack corresponding space faring capabilities. Moreover, there has been a distinct rise in the frequency, complexity, and magnitude of attacks instigated by non-state actors (Figure 2.4). Contrary to common perception, there is little historical evidence indicating that non-state actors are less willing or able to engage in digital-counterspace than state counterparts. However, this may be due to a reporting bias whereby non-state incidents are widely covered but nation-state attacks are classified.




Figure 2.3: Satellite Attack Types by Decade. Data from 2020 is included for completeness on the basis of publicly available reports as of August 15, 2020. In practice, there is often substantial lag between intrusions, detection, and reporting.

Together, these trends clarify the need for research combating cyber-security threats to satellites. Attacks against satellites are happening and have been for decades. As attackers grow more sophisticated and prevalent, increased awareness of present attacker behavior is a key first step towards contributing meaningful technical solutions.

The remainder of this chapter delves deeper into this chronology to identify unsolved technical questions in satellite cyber-security. These historical incidents are contextualized vis-a-vis the threat matrix outlined in Section 2.1.4 and organized on the basis of technical subsystems (RF, Space, Ground, and Mission).

2.3 Defending The Signal

More than two-thirds of historical satellite incidents in our review related to attacks on RF communications.

A significant portion of these are best classified as “jamming” attacks, which tend to require physical mitigations such as frequency hopping. As our focus is on digital counterspace, as opposed to electronic warfare, we will not delve deeply into




Figure 2.4: Satellite Attacks Associated With State and Non-State Actors. NB: Attribution is often uncertain and subject to dispute. Further detail regarding the attribution of any particular incident can be found in Appendix A.

jamming, but it is worth noting that the anti-jamming field is well-developed [97, 98]. Jamming nevertheless does provide context regarding attacker motivations and equipment capabilities. Additionally, in Chapter 10 of this thesis, we will see how jamming capabilities might be leveraged as a second-stage for attackers seeking cyber-physical effects from compromised systems.

Beyond jamming, we outline three general categories of communications attacks. The first, eavesdropping, relates to the interception and interpretation of signals by an unauthorized party. The second, signal injection, relates to the encapsulation of malicious data inside an otherwise legitimate transmission. The final category, signal spoofing, relates to malicious hijacking or overwriting of legitimate radio signals.

2.3.1 Eavesdropping Attacks

The eavesdropping challenge for satellites is primarily one of scale. Signals from a single geostationary (GEO) satellite can encompass an entire continent due to transmission distances. This means that attackers across a wide range of jurisdictions may be capable of receiving transmissions. In our historical review, we find that the scope and frequency of eavesdropping incidents has increased significantly. It

has been suggested that this is largely due to widespread access to the requisite equipment — such as Software Defined Radios (SDRs) — at reduced costs [99].

Despite the intuitive benefits of encryption in these environments, practical limitations on satellite crypto-systems are substantial. Satellite signals travel immense distances and are subjected to significant packet loss and latency as a result [100]. The naive addition of terrestrial end-to-end schemes, such as virtual private networks (VPNs), can have severe negative performance impacts. By some estimates, this can amount to as much as 80% reduction in perceived performance [2]. In Chapter 6 of this thesis, we characterize these performance trade-offs more closely and design adaptations suited to broadband systems.

Most solutions to this dilemma focus on ground-based tools, treating the satellites as “bent-pipes” for relaying encrypted signals. While this approach works for broadband applications, situations such as satellite telecommand necessitate the ability to encrypt and decrypt data on-board. Here, harsh orbital radiation and compute-constrained hardware act as significant barriers to adopting simple terrestrial approaches [44]. Indeed, research has suggested that attackers might even abuse poorly-implemented cryptographic defenses as a denial-of-service vector by overwhelming satellites with large quantities of deliberately invalid data [100].

Encryption in Broadcast Networks

In broadcast networks, there has been an enduring “game of cat-and-mouse” between television providers and signal pirates. This has given rise to some unique cryptographic systems.

One of the most widely used is the Common Scrambling Algorithm (CSA), which encrypts Digital Video Broadcasting (DVB) streams with a hybrid combination of stream and block ciphers [101]. CSA has been shown to have severe weakness which make it possible to crack streams in real-time on consumer hardware [101–103].

Alternative schemes are often proprietary and rely on smart-cards or specialized receivers with pre-distributed keys. An example is the DigiCipher format, which accounted for around 70% of encrypted satellite broadcasts in North America

in 2012 [41]. Another popular system is the PowerVu, which is used by the American Forces Network [41]. In 2014, it was demonstrated that PowerVu root management key entropy could be trivially reduced to a 16 bits, enabling real-time attacks on the system [104].

Beyond the direct cryptographic issues with proprietary ciphers, key management over one-way broadcast networks is difficult. Pirates will often emulate or copy smart cards to share one legitimate subscription among hundreds of users. This works because satellites broadcast signals with a single key which all customers with the footprint must be capable of deriving. As satellite broadcasts are at continent scale, the harms of key leakage are substantial. France et al. proposed a process by which individual keys could be revoked without re-issuing cards to all legitimate customers [105]. This fits within a broader body of academic work looking at key-revocation in satellite broadcast, but the problem remains unsolved in practice [100, 106–109].

Encryption in IP Networks

For internet and broadband, encryption is more complex. Due to speed-of-light latency, particularly in long-range GEO networks, TCP can suffer several negative performance effects [100, 110]. Satellite ISPs mitigate these issues and preserve limited bandwidth through the use of active traffic manipulation [100, 111]. This requires ISPs to have direct access to customer TCP headers. As a result, the use of VPNs and customer-implemented end-to-end encryption results in significant performance reductions.

Several solutions have been proposed to protect traffic over-the-air while maintaining acceptable performance. For example, Roy-Chowdhury et al. suggests the use of a multi-step SSL variant reveals certain header information to ISPs while leaving payload data encrypted [100]. Duquerroy et al. proposed a modification of IPSec called SatIPsec, which provides a layer-three encrypted tunnel with support for multicasting encryption [111, 112]. However, this solution also granted ISPs access to some customer traffic and required pre-shared secrets [100]. In Chapter 6,

we propose our own open-sourced solution which leverages the UDP-based QUIC protocol for over-the-air encryption [113].

In practice, many satellites ISPs use none of these solutions, instead sending sensitive customer traffic in clear-text. In Chapters 4 and 5, we will see how this impacts the security and privacy of home internet customers, critical infrastructure systems, and maritime vessels [114].

2.3.2 Signal Injection Attacks




Figure 2.5: Demonstrative Signal Injection Attack. At time 1 in the figure, the user requests a secure download file from a trusted server. Both the ISP (2a) and the eavesdropper (2b) receive this transmission concurrently. While the ISP routes the request over the internet (3a), the attacker transmits a malicious response to the customer antenna (3b).

Significantly less research has been conducted on signal injection attacks, a demonstrative example of which appears in Figure 2.5. Historically, satellite companies operated under the assumption that the cost of requisite equipment to alter or misuse legitimate satellite signals was beyond the means of most attackers [45]. However, our chronology turned up recent attacks which required little to no specialized equipment. For example, the Turla group attacks uncovered by Kaspersky in 2015 demonstrated that simply transmitting normal web-requests to IPs in a satellite network could be used to inject untraceable malware command and control communications in satellite broadcasts [94, 115]. Similarly, a security researcher

demonstrated that consumer SDRs could transmit specially crafted packets on the Globstar network, despite their use of complex Distributed Spread Spectrum (DSS) signals [83]. Further, the theoretical threat has been suggested, but not demonstrated, that an attacker could inject packets directly into a user's receiving antenna — potentially allowing them to bypass network firewall restrictions [111]. Lane et al. argues that carefully crafted packets may even be used to trigger vulnerabilities in the networking hardware stacks on the satellite itself [18].

To the best of our knowledge, no general defense against signal injection has been proposed. However, many of the encryption protocols discussed in Section 2.3.1 could have the benefit of bolstering the general integrity of satellite signals. Additionally, emerging telemetry standards, such as Space Data Link Security (SDLS), would intuitively complicate these attacks [116].

2.3.3 Signal Spoofing Attacks

The final category of signals-based attacks in our chronology is signal spoofing. The form and severity of these attacks has varied widely between incidents. However, the most common are attacks targeting media broadcasts — generally satellite television signals. Here, attackers typically replace the attacker's uplink signal with a more powerful malicious radio transmission [5].

As broadcast satellites often operate as dumb “bent-pipes,” they will dutifully relay any incoming transmission on the correct frequency. The most intuitive protection against such attacks is on-board verification of incoming signals. To the extent that such mechanisms exist in the status quo, they rely on proprietary trade secrets which have not been well characterized. To our knowledge, no public on-board verification standard exists. Such a system is non-trivial due to compatibility requirements with legacy ground stations, high cost of replacing orbital hardware, and general difficulties with encryption in space (see Section 2.3.1).

One variant which has received more academic attention relates to the spoofing of Global Navigation Satellite System (GNSS) signals, such as those from the Global Positioning System (GPS). Because GNSS signals are weak by the time they reach

Earth, attackers can overpower these transmissions locally using inexpensive and widely available equipment. GNSS spoofing has been studied since at least the late 1990s, but the recent emergence of consumer-grade SDRs has made it possible for even hobbyists to spoof GNSS signals [22, 55]. Indeed, in 2016, SDR-enabled wireless GPS spoofing attacks were used by players of the mobile game Pokemon GO as a cheating mechanism [117].

The simplest GNSS spoofing attacks target terrestrial receivers with simulated GNSS signals [118]. More complicated attacks seek to evade detection by, for example, correcting time synchronization discrepancies or modifying known valid signals rather than simulating from scratch [118]. The most sophisticated attacks may go further, simulating the spatial distribution of legitimate GNSS satellites to replicate expected physical characteristics [119].

Dozens of defenses against GNSS spoofing have been proposed. These range from sanity checking GNSS readings with additional sensor data (e.g., using an accelerometer to identify GNSS motion that does not correspond to physical motion) to detecting spectrum anomalies against an historic baseline [120, 121]. A full treatment of GNSS counter-spoofing could easily exceed the length of this thesis. As a starting point, Jafania-Jahromi et al. provide an accessible but deep survey of more than a dozen different classes of GPS defense techniques, including ones which allow individuals to determine locations accurately in the presence of an attacker [122].

Beyond GNSS spoofing, little research attention has been paid to the spoofing of satellite broadcasts. These range from satellite internet services to specialized critical infrastructure communications links. Given the relative maturity of the GNSS security community, it is possible lessons learned there may prove applicable to related challenges. Future research which considers the utility of GNSS counter-spoofing for non-GNSS transmissions may be a promising avenue.

2.3.4 Future Directions in Space Signals Security

At a high level, the dominant security challenge for satellite RF links is their inherently public nature. While similar issues have been mitigated terrestrially (e.g.

Table 2.3: Example Research Directions for RF Domain.

Security Challenge	Domain-Specific Obstacles	Selected Areas of Relevant Expertise
Unencrypted satellite broadband signals can be intercepted across vast distances.	Speed of light latency creates compatibility barriers between network optimizations and traditional VPN encryption tools.	Network Security, Delay Tolerant Networking, Physical-Layer Encryption
False messages can be injected into legitimate transmission relays via SDRs.	Satellites are resource limited and lack capability to cheaply verify incoming transmissions for relay.	Light-Weight Encryption, PKI and Signature Systems
Satellite signals can be overwritten with fake or malicious transmissions.	Satellites are far away and signal strength on the ground is very low leading to attacker advantages.	RF Scrambling Techniques, PKI and Signature Systems

in cellular networks), the unique hardware and environmental constraints of orbit mean few terrestrial solutions are “drop-in” compatible with satellites. This has impeded the widespread adoption of link-layer encryption. Even when defenses are widely employed, such as in broadcast television, they often depend on proprietary “black-box” encryption schemes which have repeatedly proven vulnerable.

Without robust and open security protocols which consider the unique demands of space, attackers will continue engaging in sophisticated eavesdropping, injection, and spoofing attacks. This goes beyond the academic task of inventing crypto-systems. Many well-studied proposals, such as SatIPSec, have been largely ignored [112]. Effective future work must incorporate not only technical systems security perspectives, but also pragmatic treatment of operational needs.

Based on our analysis in this section, Table 2.3 outlines three demonstrative examples of research challenges which exhibit unique characteristics in the context of space and areas of systems security expertise from which solutions to these problems might be derived.




Figure 2.6: A simplified satellite architecture with example compromise scenarios for onboard sub-systems.

2.4 Defending Space Platforms

When compared with the RF domain, only a small amount of literature exists on the defense of satellite payloads themselves. This dearth of research likely results from a few factors. First, satellite payloads have historically been highly bespoke systems [2, 14, 123]. Academics seeking broadly novel scientific findings may struggle to generalize from issues relating to any specific platform. This is further compounded by the proprietary and often restricted nature of satellite hardware, with export controls impeding trans-national collaboration. Finally, the industry acts as a “gatekeeper” to many of these components and often demonstrates skepticism or even hostility towards security research [27].

Also worth noting is that aerospace academia differs substantially from systems security academia. Even “simple” CubeSat projects are multi-year endeavors involving dozens or hundreds of collaborators [124]. A doctoral thesis in aerospace engineering may revolve around the design of a single sub-system for such missions [125]. Publications, especially prior to launch, tend to consist of narrow descriptions of engineering and implementation details, with less focus on broad theoretical generalizations and more focus on practical lessons and novel techniques. It is not unusual to encounter aerospace research describing the results of months or years of 3D-CAD modeling or simulation in less than five pages, with the model or satellite itself constituting the main contribution.

The relative verbosity, fast publication rhythm, and paper-first culture of systems security academia complicates cross-domain collaboration. On some topics, such as RF communications, this matters little; security academics can manage financially and technically without support from their aerospace counterparts. However, the satellites themselves are inordinately complex and expensive. Security researchers wishing to “go it alone” can struggle to make meaningful headway. Future work demonstrating collaboration that fulfills career and scientific objectives of participants in both areas is much-needed.

As satellite development undergoes significant changes, collaboration may become slightly easier. For example, the increased use of COTS components both lowers research costs and increases the chance for generalizable findings.

Despite present day barriers, there is a small body of prior work on payload security. Wheeler et al. proposes four general attack surfaces: input systems like sensors and RF antennae, output systems such as telemetry transmitters, internal communications such as Spacewire buses, and the underlying flight computer which integrates these components [34]. They further offer a “top-ten” list of payload vulnerabilities, ranging from malformed sensor data causing buffer-overflows to a maliciously triggered safe mode instigated by unhandled hardware states [34].

The primary status quo defense against these attacks is boundary delineation. Through RF encryption and specialized ground station hardware, satellite operators mitigate the risk of malicious individuals issuing control instructions to trigger unintended in-orbit behaviors.

Cohen et al. takes issue with this strategy, contending that it creates an “open-trust” environment in orbit [29]. Once boundary protections are overcome, lateral movement aboard the spacecraft and privilege escalation are trivial. This increases the threat posed by backdoors introduced into the spacecraft during its development. For example, Figure 2.6 provides a high-level schematic of on-board satellite subsystems and scenarios where compromise of each could escalate to mission failure.

This problem is not easily mitigated as cyber-attacks and environmentally-induced hardware malfunction are often indistinguishable to ground observers [34].

Table 2.4: Example Research Directions for Platform Security.

Security Challenge	Domain-Specific Obstacles	Selected Areas of Relevant Expertise
Malicious bus messages from third-party hardware can command flight-critical systems.	High availability requirements mean anomaly detection false positives are unacceptable. Limited cryptographic capabilities and packet sizes complicate terrestrially-viable techniques further.	CAN Bus and Automotive Security, Industrial Control System Security, Anomaly Detection
Logic error in read-only communications driver can be exploited to cause buffer overflow.	Zero physical access to reprogram or replace vulnerable component. Reprogrammability increases risk of platform hijacking or other attack vectors. Device physically reachable over RF by attackers in remote locations.	Formal Verification, Embedded System Security
Impossible to differentiate between radiation-induced hardware failure and cyber-attack.	Zero physical access makes logs irrecoverable in the case of system failure. Limited storage capacity and data transfer bandwidth makes robust logging cost-prohibitive.	Anomaly Detection, Intrusion Prevention Systems, Embedded Forensics

The remoteness of space means that forensic auditing capabilities must be built prior to launch and remain uncompromised following an attack [29]. Moreover, the limited bandwidth, data-storage, and compute capabilities of satellites means that it is uneconomical to store or transmit audit logs [29]. The transmission and storage of security data directly competes with core mission functionality. To mitigate these issues, both Cohen et al. and Wheeler et al. independently suggest the adoption of an on-board monitoring agent which detects behavioral anomalies and engages in autonomous intrusion prevention [29, 34]. This would facilitate clearer auditing and recovery in response to malicious behaviors, but, if not implemented correctly, could trigger harmful false-positives. Unfortunately, this approach lacks backwards compatibility, although some basic functionality (such as audit logging) may be applied to existing satellites [34].

In addition to security monitoring, it has been suggested that satellite hijacking attempts could be prevented through frequent, automatic re-imaging of satellite software [126]. By storing a verified secure copy of the satellite operating system on a trusted platform module (TPM), it may be possible to limit the amount of time in which an attacker might abuse the system. There are two notable downsides to this approach. The first is the addition of new hardware components, increasing satellite weight and power drain. The second is that it makes it difficult for satellite operators to patch other vulnerabilities as, in the case of a truly read-only verified firmware, they would be overwritten by the older vulnerable version.

Some attention has also been paid to the flight code itself. Wheeler et al. notes that more than 95% of the alerts raised by conventional code analysis tools triggered false positives on one demonstrative satellite, and they suspect many false negatives [34]. It has been suggested that formal verification may mitigate these issues, but no practical solution has been demonstrated to date [18]. Moreover, satellite software is rarely monolithic, incorporating third-party code for various components which increases the risk of software backdoors.

In sum, payload security is an important topic which has received very little technical research attention. Prior work has proposed a range of severe and, to date, unmitigated attack vectors. The barriers to research for payload security are particularly acute, making it a difficult area for systems security experts to contribute. However, many of the challenges identified in this domain could benefit greatly from such contributions. In Table 2.4 we present three example topics as demonstrative starting points.

2.5 Defending Satellite Ground Systems

Unlike space platforms, which suffer from esoteric hardware and limited access, ground systems benefit from the wealth of general cyber-security knowledge. Typically, satellite ground stations are not distinct from any other terrestrial computing network and, where they do differ, remain similar to terrestrial communications systems [24]. Although diversity of implementation exists, all ground stations

at a minimum consist of radio equipment and a computer which operates this equipment. Normally, the computer will run traditional operating systems with specialized software for satellite communications.

On rare occasions, this specialized software has been targeted. For example, in 2000 hackers stole copies of Exigent satellite control software for the purpose of reverse engineering [127]. More typically, attacks are byproducts of untargeted intrusions (e.g., in 1999 when a curious teenage hacker accidentally gained access to NASA flight control systems [128]). Because of this, very little academic literature focuses on ground station security. Nevertheless, some unique aspects are worth consideration.

First, satellite ground systems almost always represent the final security boundary against payload exploitation [29]. As discussed in Section 2.4, satellite software and hardware typically follow an “open trust” model whereby the ground station is trusted by all devices aboard the space platform. As such, ground systems represent a single point of failure for missions. In light of this problem, Llanos and Pearson suggest the development of redundant stations so that control can be regained in the case of compromise or loss [126]. This is one potential use for emerging “ground station as a service” offerings [129].

Second, satellite ground systems may be located in remote areas with limited physical security controls [41]. This arises because the main placement considerations relate to signal coverage and access to a particular orbit. Often, little to no staff will have a regular physical presence on-site [5]. Instead, day-to-day operations will be highly automated and controlled remotely from a centralized operations center [5]. This increases the threat of attacks leveraging physical access and contrasts with many other critical information systems.

Finally, satellite ground stations are generally the main “bridge” between the terrestrial internet and satellites. Due to heavy use of remote access, ground stations are difficult to fully “air-gap” [18, 130]. Prior security research has found numerous readily exploitable vulnerabilities in ground station software and demonstrated that ground terminals can be easily identified using IOT search engines like Shodan [92,

Table 2.5: Example Research Directions for Ground Security.

Security Challenge	Domain-Specific Obstacles	Selected Areas of Relevant Expertise
Destructive malware or denial of service attack against ground station can functionally isolate space mission.	High hardware costs mean that operators may only have a single point-of-contact with their satellites.	Cloud Security, Distributed / Shared Sensor Networks
Compromised ground station computer can issue trusted flight control commands to satellite.	Limited in-orbit verification capabilities mean verification is often on the basis of the ground station rather than its user.	PKI and Signature Systems, Industrial Control System Security
Backdoor in signal-processing hardware conceals critical data (e.g., photographs of a certain region).	Heavy use of proprietary protocols and hardware components. Single point-of-failure means attackers only need to manipulate <i>reception</i> of data, rather than transmission.	Supply Chain Security, Signal Processing

93]. Moreover, the relative normalcy of ground station hardware means that barriers to entry are low compared to other segments.

Generally, traditional enterprise security practices are prescribed to defend ground systems. For example, auditing malware on a ground station can be done with traditional forensic tools [29]. There are some systems which are unique to the satellite environment and may require special security treatment — such as long-range radio hardware [25]. However, our historical analysis has found no public instance of attacks targeting this equipment and limited academic study of these factors.

In sum, ground station security is typically considered an extension of traditional IT security. The critical difference is often in the severity of potential harms rather than the mechanisms of attacks and defenses. However, this maxim is far from

universal. Future offensive security work focused on unique satellite mission control hardware and software may uncover previously overlooked vulnerabilities. A few demonstrative examples of research directions for these dynamics appear in Table 2.5.

2.6 Holistic Security Models

While the subsystem taxonomization in this chapter is useful for identifying technical challenges and contributions, it neglects one key area of evolution in real-world satellite security practice. A sizable literature base has emerged discussing high-level organizational best practices and security frameworks. As this research tends to be more theoretical than applied, domain access barriers are less acute. This may explain the relative abundance of security frameworks compared to technical research.

Generally, these frameworks can be categorized into two broad classes: those which focus on the organizational practice of satellite operators and those which focus on the duties of policymakers. In this section, we will briefly discuss some core challenges facing such frameworks and some of the more consequential proposals.

2.6.1 Operational Frameworks

A key first step in developing any satellite security framework is to define its scope. Cunningham et al. argue that this is best done by dividing satellite missions into five broad phases and linking each phase to a distinct “cyber-security overlay” which promotes security by design. For example, in the “payload and subsystem development” phase, they suggest that satellite operators “incorporate security code and controls” [130]. Like many high-level frameworks, the core technical dimensions here are somewhat vague. However, the phase-oriented approach does bring a key benefit in delineating which organizations are responsible for given protections — an historical challenge discussed at length in Section 2.1.4.

An alternative framing is proposed by Zatti in which satellite security controls are tied to specific mission types with the addition of some generic controls common to all missions [33]. Vivero suggests a similar approach [131]. This framing seeks to balance

the diversity of satellite systems with the need for common best practices. One key advantage of mission-framing is in threat modeling. For example, the attackers interested in harming human spaceflight have radically different capabilities and motivations from those interested in compromising satellite television. Unfortunately, this framing leaves ambiguity in multi-stakeholder projects as to which organization is responsible for implementing which controls.

CCSDS suggests a hybrid approach [16]. This remediates the jurisdictional shortcomings of a pure mission-class approach while providing clearer threat-modeling. The proposal incorporates explicit consideration of mission-based attack probability mapped to lifestyle stages. While this is not an exhaustive framework, but rather a proof-of-concept, it is nevertheless among the most technically comprehensive examples to date.

The most commonly suggested approach, however, is to map preexisting IT security controls to satellite systems, though these suggestions rarely include specific mappings [23–25]. This is appealing because it draws on a set of generally accepted best practices. However, as noted by Knez et al., the uniqueness of space systems complicates this process and many controls are only superficially meaningful [23]. For example, IT security standards may impose requirements for user-account management and passwords, but satellite operating systems rarely incorporate the concept of users and accounts at all. This risks cases where large portions of standards are ignored or expensive software re-writes are required, which increases system complexity without guaranteeing meaningful security benefits.

Standard IT-security approaches neither consider the unique threat models targeting satellite systems nor the complexity of their multi-stakeholder ecosystem. When the resources for a ground system or space platform are shared by dozens of distinct business entities, applying IT security controls which assume single-owner computer systems can be difficult. Moreover, IT controls assume relatively static system lifecycles, but the security properties of satellites change significantly between lifecycle stages. For example, physical access controls which are relevant to satellites in a clean-room are meaningless for a satellite in orbit. In practice,

the NIST Cybersecurity Framework is widely employed, but it is unclear if it is fully fit for purpose [27].

2.6.2 Policy and Legislative Frameworks

Given the importance of satellites to modern information societies, it has been suggested that satellite operators may not adequately self-regulate. This is especially concerning for dual-use systems which are commercially owned but provide critical communications linkages to government operations. As such, it may be necessary to adopt regulations that re-balance incentive structures to better prioritize security.

One of the primary discussions is taxonomic. The question as to whether satellite systems are considered “critical infrastructure” remains unsettled and has a significant impact on the way in which companies and governments must protect them [19]. This may explain the relative paucity of satellite standards compared to similar infrastructure sectors [14, 24].

A general desire to classify satellites as critical infrastructure has been acknowledged by the US government since at least 2002; however, an explicit classification of this nature has yet to occur [132]. Such classification may force improvements, particularly with regards to redundancy and supply chain verification. However, industry actors have expressed resistance to rigid legal standards, contending that status quo requirements are adequate [27].

Beyond the critical infrastructure debate, an additional point of contention regards the legal rights of satellite operators to defend themselves. Rendleman and Ryals suggest satellite operators should be permitted to corrupt files and commit denial of service attacks (e.g., spectrum jamming) against attackers to regain control of their satellites [21]. They suggest the use of letters of Marque and Reprisal, a historical practice which allowed privateers to engage in combat against foreign vessels on the high seas [21]. This aligns with a broader trend applying maritime policy frameworks to space [133]. However, such “hack-back” rights are highly controversial [21].

One final notable genre of policy development centers on the international dynamics of satellite cyber-security. Housen-Couriel argues that status quo practice has created a legal lacuna in which it is unclear which international organizations and laws apply to satellite hacking incidents [134]. This suggests a need for new international law that either clarifies the applicability of existing frameworks or creates new frameworks specific to space systems [134]. Chatham House makes similar suggestions, pointing to the International Telecommunications Union (ITU) as the ideal regulatory body for such a regime [5]. They further suggest that this should incorporate interstate threat intelligence sharing due to trans-national effects of satellite failure — something which has been historically constrained by high classification levels [5]. Blount contends that cyber-intelligence sharing in space has promise due to existing collaborations (such as on debris tracking) [135]. While little progress has been made thus far, it remains possible that policymakers will seek technical input into the design of such systems.

Ultimately, we find a substantial body of policy research which has evolved more or less in isolation from relevant technical communities. The result is that many proposals appear aspirational rather than actionable. Much as in other areas, conscious effort by the systems security community to bridge this gap may pave the way for novel and impactful future work in both fields.

2.7 Lessons and Opportunities

Satellites are a vital component of modern life. From military communications to navigation and logistics to meteorological forecasting, billions of people rely on space infrastructure. This importance makes satellites attractive targets to a wide range of cyber-adversaries. Whether from an individual hacker seeking notoriety or a nation-state seeking an edge on the battlefield, satellites will face sophisticated, aggressive, and constantly evolving threats in cyberspace.

Our analysis finds a substantial legacy of digital counterspace operations spanning the past 60 years. Despite this legacy and the critical nature of these systems, the intersection between outer space and cyber space remains poorly understood.

What research does exist is scattered across diverse and isolated disciplines. In this chapter, we have synthesized these perspectives to draw out research problems which the systems security community can contribute towards solving. This is done through a taxonomy of satellite security into four sub-domains.

In the satellite RF domain, we find that the physical distance and visibility of space systems raise novel problems for broadcast and interactive communications. Attackers have exploited these dynamics to compromise sensitive data, transmit illicit broadcasts, and deceive satellite customers. The systems security community can contribute to mitigating these issues by considering the unique hardware and physical constraints impacting cryptography and verification in satellite communications and, indeed, is beginning to do so.

With respect to in-orbit platforms, we find that almost no technical research exists on the defense and monitoring of systems in orbit. Further, we surmise that this is due to substantial logistical barriers involved in accessing and conducting research with representative space hardware. While relatively few attacks have been conducted against platforms in orbit, such attacks can have particularly concerning consequences. Our analysis suggests that, with concerted effort to overcome institutional barriers, systems security researchers can make vital contributions for in-orbit anomaly detection, software verification, and forensic capabilities.

On the ground, we find that general IT security approaches are popular but that attackers still frequently manage to compromise important ground systems. The economic dynamics of satellite ground station operations, coupled with their geographic constraints, create single points of failure which are attractive targets for cyber-adversaries and difficult terrain for defenders. In thinking about shared infrastructure, monitoring, and verification, systems security research can help balance these dynamics.

Finally, from a high-level operational perspective, there exist many aspirational mission security framework proposals. There is an inherent attractiveness to the idea of “secure by design” satellite missions and a deep desire to develop best practices. However, turning these aspirations into practice is difficult. Existing security

frameworks have only limited relevance to space missions and technical diversity among satellites complicates one-size-fits-all approaches. As we move towards increased regulatory requirements for space infrastructure, incorporating technical research perspectives will be vital towards devising appropriate frameworks.

Thousands of new satellites will enter orbit over the next decade and these security questions cannot remain unanswered. Today, we sit at a critical inflection point between “old space” and “new space.” There is an opportunity for the systems security community to build upon the research of others and apply expertise that will secure the next era of human spaceflight.

The remainder of this thesis takes a preliminary step towards this objective. In Part II, we will delve deep into the topic of secure satellite communications — particularly with respect to geostationary satellite broadband services. Next, in Part III, we adapt our methods from Part II to additional topics suggested by our analysis of the space platform and ground systems domains. Finally, in Chapter 11, we will return to the questions raised here and reflect on how the approaches used in this thesis can be of broader utility for space security.

Part II

Threats and Defenses in Satellite Broadband

In what way did Aristotle suffer by having no electricity?

—Ranko Marinkovic, *Cyclops* (trans. Stojiljkovic)

3

Why Start with Broadband?

Our technical exploration of satellite systems security begins with a series of experiments focused on modern satellite broadband services. As global communications have been a key commercial driver of space development for decades, it is unsurprising that radio links have also been the favored target of cyber-adversaries attacking space systems (see Chapter 2).

Today, a market renaissance is taking place in the satellite communications sector — particularly around satellite broadband internet services. In many regions of the world, terrestrial infrastructure has proven incapable of delivering robust internet access. Whether the underlying causes of these issues arise from geographic, political, or economic factors, satellite services are uniquely suited to bringing the next billion internet users online.

The pursuit of this ambition has been a core driver of recent growth in the space industry. From novel launch vehicles to new radio systems, satellite internet has, in essence, bankrolled the “new space” revolution. The security of satellite internet thus offers an intuitive starting point for our exploration of the topic. Satellite broadband is a mature technology, almost as old as the internet itself, and is undergoing a period of evolution rather than pure invention. Understanding the state of status-quo systems can thus better inform industry efforts to design and develop their replacements.

We focus primarily on services from Geostationary Earth Orbit (GEO), which is located approximately 30,000 km above the Earth's surface. GEO is the dominant orbit for today's satellite broadband market due to its wide-coverage areas and high visibility. In theory, a single GEO satellite can service as much as 30% of the Earth's surface. This makes GEO attractive for services targeting vast and remote regions, such as the world's oceans.

New systems under construction are beginning to eye Low Earth Orbit (LEO) as an alternative due to its lower altitudes (approx. 1,000 km). The proximity of LEO offers benefits in the form of reduced latency and per-satellite launch costs. However, these advantages also require vastly more satellites to provide continuous coverage of any given location. Proposed next-generation constellation designs will thus require hundreds or thousands of satellites.

While LEO services may, someday soon, represent a significant slice of the market, GEO security is still an important topic. First, from a pragmatic perspective, mature GEO services exist at scale and applied real-world research on these networks is possible in a way that research on hypothetical future constellations is not. Second, GEO services are likely to remain critical to many industrial and transport sector operators, who rely on the integration of costly communications equipment in long-lived technological platforms. Even after LEO services are widespread, GEO offers significant coverage and cost advantages for certain markets and use-cases, such as latency-insensitive IoT networks. Perhaps most importantly, the design of GEO services can tell us a great deal about the state-of-the-art in satellite networking practice. Studying these networks and their security properties is thus an opportunity to ensure that the mistakes of the past do not become legacy characteristics of future systems.

Our investigation of satellite broadband security begins with a broad experimental survey in Chapter 4. In it, we scan the emissions of more than a dozen satellites in GEO orbit over Europe for satellite broadband signals. Together, these satellites provide internet services to more than 100 million square kilometers of the Earth's surface. The focus of Chapter 4 is a legacy protocol stack which is similar to some

of the first mechanisms used to transmit internet traffic over satellite feeds — an interactive derivative of the MPEG video streaming format. Despite the protocol's age, it is still used by thousands of industrial and home customers and is one of the dominant formats for satellite internet service offerings. We find that many of these service offerings include no provision for over-the-air encryption, leaving security up to individual customers who are either unaware of that burden or unable to meet it. The result is that critical infrastructure systems and individual web-browsing data are exposed to severe long-range eavesdropping threats which can be exploited using widely available home television equipment.

In Chapter 5, we consider a more modern enterprise-oriented protocol designed for sophisticated hardware costing hundreds of thousands of dollars. We delve deeper and more systematically into emissions from two major providers of maritime satellite broadband services over the North Atlantic, Baltic, and Mediterranean regions. Through the development of our own signal processing software, we demonstrate a technique to repurpose inexpensive home television equipment into a sophisticated satellite eavesdropping platform — representing a substantial shift in threat model. By forensically reconstructing corrupted signal recordings, we show that even low-resourced attackers can engage in attacks against modern enterprise satellite broadband offerings. The potential harms of this threat model are evaluated in the specific context of maritime technology, and we find that thousands of vessels fail to protect safety and privacy-critical data over these network links.

Chapter 6 focuses on the underlying causes of these security shortcomings — both in terms of ISP failure to encrypt traffic in the air and customer failure to use encrypted protocols to protect their own traffic. We find that specific physical characteristics involved in high-latency satellite communications are improperly handled by standard communications security tools, such as VPNs. Motivated by this finding, we invent, implement, and evaluate a novel security tool which empowers individual satellite broadband customers to encrypt their traffic without suffering performance reductions. This tool is evaluated through realistic and reproducible simulations in a purpose-built satellite security testbed. Both the

testbed environment and security tool are further developed into freely available open-source software to facilitate future research in this topic area.

Overall, this section tackles a significant security issue impacting millions of people who depend on satellite broadband. The scope and harms of these threats are demonstrated experimentally in a process of applied vulnerability research and disclosure. Rather than simply outlining security failings, we delve into the underlying physical drivers of these issues. This helps us design an effective mitigation tailored to the needs of both ISPs and customers, without imposing unreasonable performance and security trade-offs.

Moreover, our experience researching the topic of satellite broadband security serves as a starting point for continued research on space systems security in general. Over the course of this section, we will see how the physicality of space systems has direct and, at times, surprising implications for their security. The open and exploratory research process employed in this section is ultimately distilled into the repeatable *RCMA* template outlined in Chapter 1. In doing so, we supplement our direct contributions to the security of satellite networks with new research methods in service of the broader objectives of cyber-security for space missions.

Forsan et haec olim meminisse iuvabit.

*Perhaps, one day, it will be pleasing to remember
even these things.*

—Vergil, *Aeneid* 1.203

4

Old Vulnerabilities, New Applications: DVB-S MPE Threats

Contents

4.1	A Legacy of Exploitation	62
4.1.1	Motivation and Contributions	64
4.2	DVB-S Broadband Architectures	64
4.2.1	Data Visibility	65
4.3	Full Horizon Survey: Experimental Design	67
4.3.1	Deployment	68
4.3.2	Data Collection	68
4.3.3	Ethics, Data Privacy, and Legal Considerations	69
4.4	Vulnerabilities and Findings	70
4.4.1	Privacy	70
4.4.2	Infrastructure Systems	72
4.5	Improving Encryption: Status Quo Shortcomings	74
4.6	Summary	76

This chapter focuses on the ability of low-resourced malicious actors to undermine privacy and security in one widely used and established satellite broadband protocol: Digital Video Broadcasting – Satellite (DVB-S) traffic with Multi-Protocol Encapsulation (MPE). DVB-S MPE is a decades-old approach to the transmission of packetized data in satellite environments and is relevant to both critical infrastructure and consumer applications. As the next generation of satellite




Figure 4.1: Transmissions destined for locations roughly covered by the polygon on this map were analyzed at our collection site in Europe.

broadband communication protocols begin to emerge, understanding the security and privacy characteristics of mature implementations can pave the way towards improvements in incipient systems.

In this chapter, we present a broad experimental assessment of threats to modern DVB-S MPE broadband services. This is done through the analysis of radio emissions from 14 geostationary satellites with a combined signal footprint of more than 100 million square kilometers (Figure 4.1).

4.1 A Legacy of Exploitation

Given the historical longevity of DVB-S MPE broadband, there has been some prior work on the security of these networks. However, most relevant work took place in a radically different world — predating the emergence of smartphones and dramatic growth in infrastructure connectivity. As a result, the principal threat model in prior work focuses on home-internet users and tends to prioritize individual ISPs rather than underlying practices.

For example, in 2005, a team of researchers from Ruhr University Bochum published what appears to have been the first experimental security study of DVB-S satellite broadband [136]. They collected traffic from a single Astra satellite (since deprecated) and observed that the satellite transmitted DVB-S encapsulated web-browsing data in the clear. The researchers suggested that end-users who rely on

satellite broadband employ encrypted alternatives to many vulnerable application layer protocols (e.g., replacing POP3 with POP3S for email).

A few years later, at Black Hat 2009, private security researcher Adam Laurie presented on a traffic interception experiment which used modified equipment provided by a satellite ISP [137]. This was followed at Black Hat in 2010 by Leonardo Nve Egea, who demonstrated satellite internet sniffing using DVB PCI cards in the Ka-Band [82].

Some peripherally related academic work has emerged in the form of various standards revisions for satellite internet [111, 112, 138]. Moreover, some related non-academic work on other aspects of satellite systems (such as the security of software on satellite terminals) has taken place [92, 93]. However, the ensuing years have seen very little study of satellite broadband transmissions and, at the time of this research, more than a decade has elapsed since the last major academic consideration of these networks.

Beyond academia and security conferences, hobbyist and criminal communities are the primary source of modern domain-specific research. For example, online communities dedicated to the receipt of free satellite television have developed various high-quality tools for scanning, intercepting, and interpreting DVB-S signals [139, 140]. Online forums dedicated to the illegal cracking and cloning of private keys associated with satellite television networks are active hotbeds of informal offensive security research [141]. Finally, criminal groups have demonstrated the usage of satellite internet connections for exfiltrating data to undetectable command-and-control stations [94].

In short, only a small body of high-quality but dated academic research on DVB-S MPE security exists. Researchers in the 2000s suggested severe security shortcomings, but no assessment has been performed to determine if the modern situation has improved or changed [82, 136, 137].

4.1.1 Motivation and Contributions

This chapter seeks to make several further contributions over that provided in prior work. Rather than focus on a single satellites and or ISP, we implement a broad “whole sky” survey, representing the first large-scale multi-satellite security study of DVB-S broadband. While the vulnerabilities we scan for were posited in the early 2000s, there is little information as to the modern day characteristics of DVB-S based satellite broadband. The internet of 2005 and the internet of today are radically different and it is hard to intuit how DVB-S services may have evolved over this time. For example, satellite ISPs may have reacted to prior work and adopted new techniques to protect these networks - perhaps in response to evolution in data protection best practices and regulations. Similarly, in the early 2000s, embedded and internet-of-things systems were relatively rare, and prior issues were primarily relevant to home internet users. Today, connected infrastructure and operational technology makes a significant component of the satellite broadband market and there is little prior work as to the security characteristics of these connections in practice.

In short, this chapter begins by replicating prior work and updating it for a modern context, but it seeks to go beyond that in terms of scale and depth. As detailed in Section 4.3.1, we develop our own software to automatically identify internet-relevant satellite feeds at scale, allowing us to search through more than 300 satellite transponders and many thousands of program streams automatically. This enables us to speak not just to the security behavior of connections belonging to a single ISP, but about endemic industry practices and security issues. Finally, this chapter lays the foundation for our consideration of the previously unstudied traffic types in Chapter 5 and provides broad motivations for the defenses proposed in Chapter 6.

4.2 DVB-S Broadband Architectures

The long distances in geostationary earth orbit (GEO) satellite networks (72,000 km from customer to ISP) have resulted in protocols designed specifically for satellite

radio broadcasts. Among these are the widely used Digital Video Broadcasting-Satellite (DVB-S) and DVB-S version 2 (DVB-S2) protocols [142]. DVB-S is widely used in the provision of satellite television services. This has resulted in the emergence of numerous free tools and protocol analyzers, allowing individuals to receive satellite television without purchasing subscriptions. Moreover, DVB-S is the de facto global standard for satellite broadcast and IP services, particularly from GEO [143, 144]. When compared with more proprietary protocols for which comparatively little public tooling exists (e.g., Inmarsat’s BGAN system), this makes DVB-S a particularly attractive target for attackers [145].

The DVB-S standard typically transmits data per the Moving Pictures Experts Group (MPEG) standards in the form of MPEG transport streams (MPEG-TS) [146]. While MPEG-TS is primarily used for media broadcasting, the standard has also been extended to support many other types of data. In particular, one such extension, called “Digital Storage Media Command and Control” (DSM-CC) was developed to provide interactive features on Video Cassette Recorders (VCR) [147]. As demands for interactive satellite broadband services grew, DSM-CC was repurposed to relay arbitrary packetized data to and from satellite internet customers via an additional encapsulation layer called Multiprotocol Encapsulation (MPE) [148]. A subsequent revision of the MPE method called Unidirectional Lightweight Encapsulation (ULE) has also been created which allows for the transmission of packetized data without the use of DSM-CC tables [148]. In both protocols, IP packets destined for many distinct customers are transmitted on the same stream and then extracted by customer equipment on the basis of address information in the IP header or ISP-assigned MAC addresses in the MPE/ULE headers.

4.2.1 Data Visibility

The specific topology of satellite networks can have significant impacts on an attacker’s ability to understand broadband transmissions. Our threat model takes into account two of the most common network topologies.




Figure 4.2: A single-band two-way satellite internet setup. Attackers at position 1 and 2 will see downstream and upstream traffic respectively.

The first configuration leverages two-way linkages from the customer's satellite dish and is ideal in remote locations (such as rural areas or naval vessels) where no terrestrial connectivity is available (Figure 4.2). The customer transmits a web request directly to a satellite, which then relays these requests on another beam towards an ISP-controlled ground station. At the ground station, this request is relayed across the wider internet and the response is subsequently transmitted back the customer. An attacker listening to the downlink-to-consumer connection (Position 1 in Figure 4.2) could intercept responses from the internet, while an attacker listening to the downlink-to-ISP connection (Position 2 in Figure 4.2) would be able to view requests made to the internet. Depending on the specific geographic location of the attacker in either scenario, additional traffic may also be visible through interception of signals emanating from antenna side lobes.

The second configuration offered by some specialized ISPs such as BROADSAT's Opensky combines satellite and terrestrial linkages (Figure 4.3) [149]. Uplink requests are transmitted via a terrestrial connection, which typically has better latency (Position 1 in Figure 4.3). Requests are sent to a proxy operated by the satellite ISP (Position 2 in Figure 4.3) and responses are relayed back to the user via satellite (Positions 3-5 in Figure 4.3). Such configurations are ideal in cases where uplink latency is of higher priority than uplink bandwidth and for customers with extant but inadequate terrestrial service. Here, an attacker would be incapable of observing uplink traffic over the air as it is transmitted via wire rather than radio.




Figure 4.3: A combined terrestrial uplink, space downlink satellite internet setup.

4.3 Full Horizon Survey: Experimental Design

For this chapter, we sought to discern the capabilities of a single malicious individual rather than a larger organization or nation state. As such, we restricted our equipment selection to hardware that was readily available for purchase online and only employed free software tools. The total cost of necessary equipment was under €300, as demonstrated by the budget in Table 4.1.

Table 4.1: Hypothetical Attacker Budget.

Equipment	Cost
Selfsat H30D Satellite Dish	€85
TBS 6983 Satellite PCI-E Card	€197
3-Meter Coaxial Cable	€3
Total	€285

It is worth noting that equipment quality can have a meaningful impact on capabilities. For legitimate customers, specialized hardware targeted to their ISP is used in the form of a satellite receiver/modem. Relying on generic equipment can increase processing errors – especially for complex modulation modes such as 16 and 32-APSK.

4.3.1 Deployment

Two inexpensive satellite receiver assemblies (of the sort intended for RV-camping) were deployed to simulate this threat model. One consisted of a 75 cm, flat-panel satellite receiver dish and a TBS-6983 DVB-S receiver. The other consisted of a 60 cm flat-panel dish, a motorized targeting assembly, and a TBS-6903 DVB-S receiver. The 75 cm dish remained in a fixed position while the 60 cm dish was repositioned to target many satellites over the duration of the study. These panels were configured to receive Ku-band transmissions between 10,700 MHz and 12,750 MHz with both vertical and horizontal polarizations. Both assemblies were located in Europe and, due to environmental constraints, could observe geostationary satellites positioned between 40 degrees East and 37 degrees West.

A set of 14 geostationary satellites were selected based on signal quality at the collection site. From these satellites over 350 transponders were identified using existing “Blind Scan” tools.

A collection of python utilities developed for the purpose of this study was used to analyze each of these transponders for signs of DVB-based internet transmissions on the basis of three criteria. First, a stream was deemed more likely to carry internet traffic if DSM-CC (MPE) services were listed in the stream’s program table. Second, streams which contained valid UDP or TCP packets, based on existing MPEG-TS dissectors in Wireshark, were flagged as candidates. Finally, streams were parsed against a list of regular expressions commonly seen in internet traffic. A total of 19 streams met at least two of these criteria and 4 additional matches were identified on the basis of the regular expression engine alone. From these 23 transponders, streams which appeared to carry IP-TV traffic or simple device firmware update services were discarded, along with streams with extremely low signal quality and throughput (anything less than 5 kb/s). This left 13 transponders for further study.

4.3.2 Data Collection

In June 2018, approximately five hours of traffic were recorded on each of the 13 selected transponders. These recordings were initiated automatically in sequence




Figure 4.4: Volume of internet traffic on each targeted transponder during the five-hour analysis window.

over the course of three days. Each recording was processed to remove NULL packets and irrelevant program data (such as program data related to satellite television). After this process, we were left with 50 gigabytes of satellite internet traffic. Depending on location, radio conditions, and transmission modes, the amount of data collected per transponder ranged from as low as 8 megabytes to as high as 10 gigabytes (see Figure 4.4).

As anticipated in Section 4.3, recordings included many data errors resulting from our use of general-purpose hobbyist equipment. Nevertheless, sufficient information could be extracted to give a general characterization of security concerns.

4.3.3 Ethics, Data Privacy, and Legal Considerations

Prior to our experiment, it was unclear what sort of information would be uncovered. As such, we assumed a worst-case scenario and treated all recorded radio signals as if they might contain sensitive information. Data was stored at the collection site in Europe and both physical and electronic access was restricted. Local laws relating to the interception and analysis of radio traffic were strictly adhered to.

We also made plans to responsibly disclose any security issues which warranted it to the appropriate authorities. After the study, all collected data was deleted.

4.4 Vulnerabilities and Findings

Across all thirteen frequencies included in the final phase of our study, broadband traffic was transmitted in plain-text. Of course, well-encrypted transmissions would not have been distinguishable from non-internet traffic and it is thus unclear to what extent these thirteen service providers are representative of the industry as a whole. Nevertheless, having the same issues appear over many distinct providers suggests that the earlier single-provider studies discussed in Section 4.1 were not merely anecdotal and the problems they identified have not yet been addressed.

The dangers of unencrypted wireless transmissions are well understood and, to some degree, academically uninteresting. However, unique properties of satellite broadband act as novel risk-multipliers.

The principal differentiator is scale. Our experiment included data from a coverage footprint of more than 100 million square kilometers (Figure 4.1). A handful of strategically located satellite dishes would allow an attacker to intercept broadband signals encompassing most of the globe. Furthermore, satellite interception offers a privileged ISP-esque vantage-point and enables eavesdropping on the entirety of a target's traffic.

4.4.1 Privacy

A surprising amount of sensitive information appeared in the collected data. Indeed, many of the same categories identified over a decade ago still appear in modern satellite traffic.

One significant improvement since the mid-2000s has been increased adoption of SSL/TLS encryption. While this protects against certain types of eavesdropping attacks, the very process of requesting and exchanging SSL certificates leaks potentially revealing information. Our data included over 52,000 SSL “wildcard” certificates from around 1,200 distinct domains (Figure 4.5). Information a user




Figure 4.5: The top 20 domains identified on a basis of SSL certificates. Number 8 has been hidden as it is a private subdomain range for a major energy provider.

might consider deeply private – such as TLS certificates or DNS responses from various adult websites – is, in fact, being broadcast across an entire continent. With collating data, such as knowledge of a user’s IP address, this risk becomes particularly severe.

While SSL usage is widespread, our naive string-matching analysis nevertheless uncovered thousands of unencrypted HTTP requests, file downloads, FTP sessions, torrent connections, VoIP conversations and emails. The chart in Table 4.2 indicates which of these general classes of sensitive information were identified on each of the 13 transponders. Due to the sensitive nature of our findings, specific service providers and satellite names have been withheld. These findings raise legal and business concerns regarding whether it is the responsibility of satellite service providers to protect customers using insecure protocols over DVB-S or if responsibility for encryption in transit falls to end-users.

Beyond this broad sense of information leakage, a number of individual narratives emerged during manual analysis. Although anecdotal, these incidents provide perspective on the need for communications security improvements. These incidents ranged from individuals who shared national identification numbers

Table 4.2: Observed Traffic Contents. Note the diversity of protocols and that, while some customers and applications encrypt their traffic with TLS, many feeds also contained unencrypted payloads.

Stream	TLS	HTTP	Email	Tokens	FTP	Files	Torrent	VoIP
1	No	Yes	No	No	No	Yes	No	No
2	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes
3	Yes	Yes	No	Yes	Yes	Yes	No	Yes
4	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes
5	Yes	Yes	No	Yes	No	Yes	No	Yes
6	Yes	Yes	No	Yes	No	No	No	Yes
7	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
8	Yes	Yes	Yes	Yes	Yes	No	No	Yes
9	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
10	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
11	Yes	Yes	No	Yes	No	Yes	Yes	Yes
12	Yes	Yes	No	Yes	No	Yes	Yes	Yes
13	Yes	Yes	No	Yes	No	Yes	No	Yes

with hotels via plain-text email messages to online shoppers submitting payment details in clear-text.

One particularly illustrative case relates to an individual who connected his iPhone to a WiFi network and synced his email client over IMAP. Potentially unbeknownst to him, this WiFi hotspot was connected to a satellite modem. As a result, he broadcast information in clear-text across all of Europe detailing the specific town in northern Spain where he lived, his full name, phone number, and both his office and personal addresses. If this were not concerning enough, the individual proved to be a defense lawyer. Included in the traffic were confidential communications between him and clients regarding ongoing cases (Figure 4.6). Our ability to intercept it raises serious concerns for attorney-client privilege and important questions regarding who should be ultimately responsible for the protection of this data.

4.4.2 Infrastructure Systems

Through manual inspection of intercepted traffic, we detected flows associated with electrical power generation facilities. The majority of these were wind and solar farms, but we also encountered facilities associated with the oil and gas industry.

```

...=3D"cs80D9435B"><span
class=3D"cs19=..C3E152">E-mail: <a
href=3D"mailto:[REDACTED]"><span
class=3D"cs2=..50A6940">[REDACTED]</span></a>
</span></p><p class=3D"csGB..%80D9435B"=..><span
class=3D"cs19C3E152">&nbsp;</span></p><p
class=3D"cs95E872D0"><span
c=..lass=3D"cs19C3E152">&nbsp;</span></p><p
cl.væx>µ»#7Á...-..E...@$....Ä,-."-..7#7+....%.....°
..G....>.*ass=3D"cs80D9435B"><span
class=3D"=..cs675EBA1">AVISO LEGAL</span></p><p
class=3D"cs80D9435B"><span class=3D"cs19=..C3E152">Este
mensaje va dirigido, de manera exclusiva, a su
destG...inatrio y c=..ontiene informaci=C3=B3n
confidencial y sujeta al secreto profesional; cuya
d=..ivulgaci=C3=B3n no est=C3=A1 permitida por ley.
</span></p><p class=3D"cs80D=..9435B"><spG...an
class=3D"cs19C3E152">En caso de haber recibido&nbsp;<span
este mensa=..je por error, le rogamos que, de forma
inmediata, nos lo comunique mediante e=..ste medio o a
trav=C3=A9s del tel=CG...3=A9fono (+34) 942 [REDACTED] y
proceda a s=..u eliminaci=C3=B3n. Asimismo, le
comunicamos que la distribuci=C3=B3n, copia=.. o

```

Figure 4.6: The footer of one email from a lawyer to his client which was sent in plaintext via satellite internet. Sensitive information has been censored.

These were not isolated to a single provider but appeared across several satellite internet services and terrestrial infrastructure operators.

While in some cases, such as an American solar power provider, TLS encryption was employed to protect infrastructure traffic, many operators used unencrypted HTTP and FTP connections. In the case of one specific software platform commonly used in the wind-power generation industry, over 5,000 plaintext requests were observed to various facilities and administration pages (Figure 4.7). Inside these requests we found credentials in the form of either HTTP Basic Authorization tokens or as session cookies that could be used to gain unauthorized access to the plants. Moreover, credentials belonging to a company which operates almost a fifth of the world's installed wind energy base appeared frequently in unencrypted FTP control flows. Vulnerable system administration pages and FTP servers were publicly routable from the open internet. This means that an attacker could sniff a session token from a satellite connection, open a web browser, and login to the plant's control panel.

Beyond electrical plants, other infrastructure traffic also appeared in our dataset.




Figure 4.7: The login page for one vulnerable wind farm we encountered in our traffic. Credentials to access this site were readily identified in satellite transmissions.

For example, on a handful of transponders we identified Modbus traffic — a popular serial protocol widely used in SCADA systems and PLCs. In another instance, a satellite transponder appears to have been reserved for the national postal service of a sizable Eastern European country, and several of their intranet credentials were transmitted in plain text. References to maritime products appeared in the streams, but an in-depth analysis of maritime communications protocols was well beyond the project’s scope. Chapter 5 takes a closer look at protocols in the maritime context.

4.5 Improving Encryption: Status Quo Shortcomings

Our findings suggest a need for improved communications privacy in satellite broadband. The current state of industry has resulted in the leakage of sensitive information from both individual and industrial satellite broadband customers.

The application of terrestrial encryption techniques to satellite environments is non-trivial. Satellite transmissions cover vast distances and are subject to speed-of-light latency effects (upwards of 500 ms for a round-trip transmission to GEO) and packet loss, which can impair the function of encryption schemes designed

for high-reliability terrestrial environments (e.g., by requiring re-transmission of corrupted key materials) [100, 150]. Moreover, satellites themselves are limited in terms of computing capabilities and any on-board cryptographic operation risks trading off with other mission functionality [100].

In television networks, “scrambling” algorithms which apply encryption to an entire MPEG-TS program are widely used to prevent piracy of premium television channels. However, these techniques are, at present, not well suited to internet traffic for two reasons. First, many of the dominant algorithms in this space (e.g., the Common Scrambling Algorithm or PowerVu) have been demonstrated to have severe cryptographic weaknesses [102–104]. While these may be acceptable for low sensitivity TV broadcasts, where the principal goal is often to simply increase the complexity of piracy, internet traffic and sensitive data may merit more robust protections. Moreover, because these scrambling algorithms apply at the level of entire streams, customers necessarily share access to a “master key” (often updated at regular intervals) which could be used to compromise the privacy of other customers whose traffic is multiplexed into the same transmission [104]. Future work which builds on these scrambling techniques with a focus on the needs of internet customers may represent a possible avenue for developing more robust DVB-S protections.

One alternative approach would be the use of tunneling mechanisms such as IPSec. In the short term, this is likely the best approach for individual customers and infrastructure operators. However, terrestrial tunneling technologies impose significant performance constraints over satellite connections [111]. Due to the significant latency in GEO broadcasting, satellite service operators have adopted connection acceleration techniques which help minimize these effects. Specifically, broadband providers widely make use of Performance Enhancing Proxies (PEPs) to simulate artificial TCP acknowledgments and TCP features which misinterpret satellite latency as network congestion [100, 110]. Moreover, although TLS functions over satellite networks, plain-text HTTP requests will often appear more performant to users due to the heavy application of HTTP accelerators [100]. Tunneling and

end-to-end encryption prevents satellite operators from inspecting the necessary packet headers to continue providing these services.

In the early 2000s, several solutions to these issues emerged, ranging from decrypting traffic at the satellite ISP ground station and re-encrypting for transmission over the internet to implementing revisions to the IPSec protocol itself to extract TCP headers needed for PEP usage; however, these have gained little traction [100, 106, 112]. While promising standards proposals exist for encrypted satellite data links in scientific missions, additional work adapting these to multi-user broadband ecosystems is needed [116].

Chapter 6 takes a deeper look at these challenges and presents a novel approach to PEP-compatible encryption which resolves many of these issues.

4.6 Summary

Our initial experimental analysis raises significant concerns for the security of DVB-S MPE broadband. Severe confidentiality shortcomings were identified across more than a dozen service operators, and several gigabytes of potentially sensitive web traffic were collected in a matter of hours. An attacker with cheaply available hobbyist equipment may compromise the security and privacy of individuals in an area encompassing tens of millions of square kilometers. Moreover, satellite eavesdropping provides a potential route to harming many connected critical infrastructure systems such as power-generation facilities.

Satellite encryption is not a trivial task, but it is a necessary one. This case study demonstrates the importance of finding cryptographic systems which balance satellite network performance and communications security. While DVB-S MPE is, in relative terms, a “legacy” standard for the transmission of IP traffic in satellite ecosystems, it is clear that many still rely on the protocol for vital and sensitive communications. Designing a protocol-agnostic encryption tool which can be employed in legacy satellite networks, as well as more modern implementations, may be necessary for protecting such users.

However, growth in satellite broadband will require more efficient and modern protocols than DVB-S MPE and the security properties of these alternatives may differ from legacy applications. In Chapter 5, we will further consider a more modern protocol and its exposure to eavesdropping attacks in a specific application domain (maritime).

Then the Frost his songs recited,
 And the rain its legends taught me;
 Other songs the winds have wafted,
 Or the ocean waves have drifted;
 [...] In a ball I bound them tightly;
 And arranged them in a bundle;

—Elias Lönnrot, *Kalevala* (trans. Kirby)

5

A Tale of Sea and Sky: Exploiting Maritime VSAT Services

Contents

5.1	Related Work on Maritime and Space Security	81
5.1.1	Motivation and Contributions	84
5.2	VSAT Broadband Applications and Architectures . . .	84
5.2.1	VSAT Network Design	85
5.3	Sustained VSAT Observation: Experimental Design .	87
5.3.1	Equipment, Targets and Recording	87
5.3.2	Data Extraction and Signal Interpretation	90
5.3.3	Collection and Forensic Performance	91
5.3.4	Extended Collection	94
5.3.5	Ethical and Legal Concerns	95
5.4	Threat Model and Attacker Capabilities	96
5.5	Broad Findings and Vulnerabilities	97
5.5.1	Applications and Protocols	98
5.5.2	Hosts and Vessels	99
5.6	Findings: Physical Safety and Operations	102
5.6.1	Navigation and Charting	104
5.6.2	Vessel Operations and Security	106
5.7	Findings: Passenger and Crew Privacy	107
5.8	Active Attacks on VSAT Services	110
5.8.1	TCP Session Hijacking	111
5.8.2	TCP Hijacking Requirements	111
5.8.3	Hijacking Implementation	112
5.8.4	Further Active Attacks	115
5.9	Underlying Causes and Next Steps	115
5.10	Summary	116

The maritime transportation industry has trended towards ever-larger vessels operated by ever-smaller crews, a change facilitated by the increasing digitization of modern ships. In December 2015 the *CMA CCM Benjamin Franklin*, with a crew of merely 27 members, brought more than \$985 million worth of cargo to the Port of Los Angeles in a single visit [151, 152]. Ships such as this have leveraged digitization to make the maritime industry a keystone sector in the global economy, transporting more than 80% of the world’s trade goods annually [153]. Moreover, the use of computing technology for marine operations is expected to grow for the foreseeable future — perhaps even progressing to fully autonomous vessels [154].

One of the critical drivers of this digitization revolution has been improvements in ship-to-shore communications. Through space-based radio transmissions, land-side operations centers remain connected to vessels traversing the remotest parts of the globe. However, despite the vitality of these connections, little research has been conducted on their security properties. This chapter makes an initial experimental contribution towards securing such increasingly critical linkages.

Specifically, the chapter focuses on one major ship-to-shore communications technology: maritime very small aperture terminal (VSAT) satellite broadband. We demonstrate that an attacker can intercept and even modify maritime VSAT connections using standard satellite television equipment costing less than 1% of state-of-the-art alternatives. Moreover, we present a purpose built forensic tool — GSExtract — designed to recover IP traffic from even highly corrupted maritime VSAT feeds collected on consumer-grade equipment.

GSExtract is used to conduct an experimental analysis of two major maritime VSAT providers offering services to Europe and the North Atlantic and encompassing a service area of more than 26 million square kilometers. These two providers rely on an underlying networking platform with more than 60% share of the global maritime VSAT market.

We find that status quo maritime VSAT communications raise serious security and privacy concerns. From more than 1.3 TB of real-world satellite radio recordings, we select a series of demonstrative case studies highlighting unique threats to maritime navigation, passenger and crew privacy, and vessel safety. Our contributions suggest that several of the world's largest shipping, freight, and fossil fuel companies rely on vulnerable VSAT networks which may be abused for the purposes of crime, piracy, and terrorism. The chapter concludes with a brief discussion of both immediate and long-term technical improvements which may address these issues.

5.1 Related Work on Maritime and Space Security

While, to the best of our knowledge, no experimental analysis of maritime VSAT radio signals has been conducted to date, a broader literature base on maritime cyber-security has begun to emerge. This sub-field is well characterized by DiRenzo et al., who synthesize a number of academic and governmental reports and outline theoretical attacks against several marine navigational technologies, including: Global Positional Systems (GPS), Automatic Identification System (AIS), and Electronic Chart Display and Information System (ECDIS) [155]. In a broad sense, the focus has primarily been on the impact of system compromise rather than the mechanism by which that compromise might occur.

Some practical consideration of attack vectors can be found in literature relating to GPS security. For example, in 2013 researchers at the University of Texas, Austin demonstrated the ability to spoof GPS position readings aboard the *White Rose of Drax*, a luxury yacht [157]. They further suggested that attackers might take advantage of GPS subsystems to alter ship coordinates and even hijack vessels. Reports of GNSS spoofing by Russian authorities in the Black Sea suggest that such attacks have been put into practice [158]. Beyond maritime, a much wider body of research surrounding the general topic of GPS spoofing and countermeasures exists [159].



Figure 5.1: A typical marine VSAT system [156].

Regarding AIS, a near-universally deployed maritime location reporting and collision prevention system, there is significant interest both within academic and hobbyist circles. Radio communities have emerged using software-defined radios to record AIS signals and develop open source maps tracking maritime traffic [160, 161]. Moreover, security-focused research has identified a number of vulnerabilities in AIS environments — including the ability to create non-existent vessels or false collision incidents [162].

In a less technical context, some work has been done to identify threat actors with motivation to harm maritime targets via cyber-mediated attacks. For example, Jones et al. contend that terrorist organizations might view a disabled or impaired oil tanker as a powerful weapon [163]. Furthermore, given the high value of typical cargo payloads (on the scale of hundreds of millions of dollars), information systems aboard ever more automated freight vessels may become targets of pirates [164, 165]. The recent kinetic attacks against Japanese and Norwegian oil tankers in the Gulf of Oman, almost universally attributed to state-sponsored adversaries,

demonstrate that modern nation states have the motivation to harm commercial maritime vessels [166, 167]. Moreover, given that no state has claimed responsibility for these acts, the plausible deniability and covert nature of cyber-operations may be particularly desirable to state actors.

Within the maritime industry, organizations appear generally confident in their ability to defend against cyber-attacks. A recent survey of maritime executives and cyber-security decision-makers found that almost 70% felt that the industry was “prepared in cybersecurity” [168]. Moreover, 100% of representatives from large maritime companies (those with more than 400 employees) felt that their company was already “prepared to prevent a data breach” [168].

Very little research exists specifically concerning the security properties of maritime VSAT. Most prominent are two conference presentations by a private security researcher from the firm IOActive at DEFCON and Blackhat which disclosed serious firmware vulnerabilities in the software of many widely used VSAT routers [92, 93]. However, the research did not extend to the radio signals transmitted to and from these devices and did not consider the capabilities of a terrestrial eavesdropper. Research into the general security of satellite broadband, such as that presented in Chapter 4, is not necessarily applicable as specialized marine systems use modern protocols geared to high-bandwidth enterprise use cases. However, given that specialized enterprise data links in other transportation sectors — such as aviation — have been found to have significant security issues, a closer look at maritime VSAT services is likely warranted [169].

The relative lack of research on maritime VSAT security may arise in part because the dominant service providers tend to leverage more complex transmission modes (e.g., 16 or 32APSK modulation) and more recent protocols (e.g., Generic Stream Encapsulation or GSE) compared to the traditional satellite broadband networks discussed in Chapter 4 [170]. While many open source and freely available tools exist for interpreting MPEG-TS recordings, to our knowledge no comparable software exists for GSE [139, 140]. Additionally, the equipment sold to maritime VSAT customers to receive and interpret these signals (such as the system in

Figure 5.1) can cost upwards of \$50,000 [171]. These high costs act as a significant barrier to entry for researchers.

5.1.1 Motivation and Contributions

In this chapter, we seek to address this lack of prior research and develop new techniques for observing and studying the security properties of modern maritime VSAT broadband services at scale. To our knowledge there is no publicly available software for the conversion of raw over-the-air recordings of GSE-encapsulated data into a packet capture format for IP and application-layer analysis. Rather, GSE decapsulation is a problem which is handled in the form of embedded software built into satellite customer modems and is designed under the assumption that the receiver is an authorized participant in the network. While it is possible that governments and intelligence agencies may have specialized tools to support signals intelligence collection targeted at GSE encapsulated feeds, this chapter focuses on developing such capabilities for low-cost consumer-grade equipment.

In doing so, we develop a novel tool for recovering data from low-quality GSE signal recordings and demonstrate a substantial shift in threat models. Further, we leverage this tool to ameliorate the lack of prior observational studies on the behavior and security properties of real-world VSAT networks. In doing so, we identify several significant, and previously unknown, security issues impacting thousands of maritime vessels.

5.2 VSAT Broadband Applications and Architectures

By enabling ships to remain connected to terrestrial computer networks, wherever they may be, VSAT has been a key driver of digitization. The specific utility of VSAT depends highly on the purpose of a given ship. For example, a cruise operator might use VSAT to provide broadband internet connectivity to their passengers whilst a fishing vessel might leverage cloud-based analysis of fishing yield data [172, 173].

There are, however, several common use cases for VSAT connectivity with broad applicability [171]. For example, marine transportation is highly regulated, and VSAT services allow ships across sectors to communicate with port authorities and land-based regulatory experts, far in advance of arrival. Moreover, modern fleet management products delivered over VSAT enable maritime companies to maintain situational awareness as to the state of their fleet, provide remote expert support, and optimize fuel efficiency and scheduling in response to weather changes [171, 174]. Finally, VSAT connectivity enables critical safety and navigational aids ranging from remote medical support to up-to-date navigational charts [171].

5.2.1 VSAT Network Design

To some extent, the term “VSAT” is a misnomer. While the acronym suggests “very small” terminals, products exceeding the size of automobiles are regularly sold as VSAT hardware [175]. Moreover, from a communications protocol perspective, the VSAT designation means very little. VSAT service operators use a wide range of protocols, many proprietary and undocumented, and generalizations applicable to the entire VSAT industry are difficult if not impossible.

Within the maritime context, however, VSAT services are more standardized due to the global nature of the shipping industry. Satellite service operators in one region of the world will enter into sub-licensing agreements with operators in other regions to provide global coverage, and this requires the use of inter-operable protocols. For example, both of the providers considered in this chapter rely on an underlying networking technology stack used in more than 1,200 VSAT networks globally and with more than 60% market share in the maritime domain [176].

In this chapter, we focus on satellite networks operating from geostationary earth orbit (GEO). When contrasted with low earth orbit (LEO), geostationary networks offer two main advantages for maritime VSAT. First, because the satellites appear stationary relative to a fixed point on the earth’s surface, receiving a signal is simpler than in LEO networks where satellites frequently pass over the horizon. Moreover, GEO satellites operate from an altitude of more than 30,000 km which enables vast

coverage areas measuring millions of square kilometers from a single satellite. These wide coverage footprints are particularly attractive to maritime customers operating in remote ocean waters. The principal disadvantage of GEO networks is that the long distances involved create speed-of-light delays that increase network latency.




Figure 5.2: The typical flow of data through a maritime VSAT network. The attacker in the diagram can eavesdrop on traffic from step 6 but has limited visibility into traffic at all other stages.

A maritime VSAT network is not significantly different from other satellite networking environments with respect to its basic architecture. As outlined in Figure 5.2, the customer sends web requests up to their provider's satellite which then relays those requests on a different frequency to a large ground station. This ground station then forwards customer requests across the open internet, receives the responses, and relays those responses back up to the satellite which then forwards those same responses back down to the customer. From geostationary orbit, speed of light signal propagation means that this process takes around 500ms in ideal conditions.

One unique aspect of eavesdropping in satellite networks that does not hold for most other wireless networks is that the geographic location of an attacker within the coverage area can have significant impacts on their ability to observe certain signals. For example, the attacker depicted in Figure 5.2 can easily observe responses from the satellite internet service provider (ISP) to the customer but would have




Figure 5.3: Signal Coverage Footprint. Traffic from ships across the entire shaded area (more than 26 million square kilometers) was observable from our collection site in Europe.

a much more difficult time intercepting the focused uplink requests transmitted by the customer. This means in our experimental analysis, the recorded traffic generally only contained “forward-link” packets received by satellite customers but not the “reverse-link” packets sent by customers to their ISPs. In theory, an eavesdropper physically located near a satellite ISP could intercept such packets, but the beams used to transmit this portion of the connection are much narrower and have smaller footprints than general broadcast signals. Additionally, the satellite to ground station link may operate over frequencies for which hardware is less widely available. A further discussion of this one-way eavesdropping threat model can be found in 4.2.1.

5.3 Sustained VSAT Observation: Experimental Design

5.3.1 Equipment, Targets and Recording

In order to assess the status quo state of maritime VSAT communications privacy, we developed an experiment to collect and analyze representative maritime VSAT emissions from two major service providers servicing shipping routes in the North

Atlantic, Baltic and Mediterranean regions. An approximate map of the signal footprints involved in our research can be seen in Figure 5.3.

As mentioned in Section 5.1, commercial maritime VSAT systems are expensive. Even if an attacker had sufficient funds to procure an installation, these systems are not generally sold direct to consumers but rather according to a business-to-business or “VSAT as a service” model (typically in the form of annual contracts costing thousands of dollars monthly). As such, an attacker might prefer to employ widely available and inexpensive satellite television equipment.

The use of a standard home-television satellite dish and inexpensive hobbyist satellite tuner gives rise to several issues. Consumer grade equipment is likely both smaller and less accurately targeted than maritime VSAT systems. This results in lower antenna gain and lower signal-to-noise ratios. The effect is that many frames will be lost in the signal processing stages. Moreover, the tuner hardware itself — normally an FPGA or ASIC based demodulator — may fail to maintain an acceptable rate of throughput when interpreting more complicated modulations. In maritime VSAT, 16 and 32-APSK modulations are widely employed for high-bandwidth connections. This contrasts with simpler QPSK and 8PSK modulations dominant in the terrestrial ecosystem and consumer-grade hardware.

Despite these issues, we hypothesize that a resource-poor attacker may nevertheless be able to intercept, demodulate, and interpret maritime VSAT streams. This is because an eavesdropper does not necessarily need 100% reliability to pose a threat, even if an eavesdropper misses half of all packets, the small portion which they do intercept may contain sensitive information. In order to test this theory, we restricted our experimental equipment to widely-available consumer-grade products with a total cost of less than \$400 (Table 5.1).

In our specific experimental setup, our equipment was capable of receiving DVB-S2 signals in the Ku-band frequency range (10.7-12.75 GHz). While maritime VSAT services are offered in many different spectrum ranges (particularly C-band due to rain-fade concerns at sea), we expect any findings in the Ku-band should hold across other frequencies. It is worth noting that our research is

Table 5.1: Experimental Equipment.

Item	Approximate Cost
TBS-6903 DVB-S2X PCI Card	\$300
Selfsat H30D Satellite Dish	\$88
3-meter Coaxial Cable	\$5
Total	\$393




Figure 5.4: Scanning for satellite streams across the Ku-band in two orientations. Distinct humps in the spectrum represent channels for potential analysis. NB: To maintain platform anonymity, the lower axis has been normalized.

restricted to DVB-S2 signals. While DVB-S2 is a dominant standard used by hundreds of satellite broadband operators, some proprietary alternatives exist. An entirely different technical approach (and possibly different hardware) would be required to analyze such products.

While the location of satellites which offer VSAT services are widely available public knowledge, the specific frequencies used are not. In order to identify




Figure 5.5: A simplified overview of protocol layers which comprise maritime VSAT streams.

frequencies, the attacker must scan the RF-spectrum of radio emissions from the satellite for channels and then ascertain which are used for VSAT services (see Figure 5.4). For this experiment, we identified a total of 15 VSAT streams on two geostationary platforms, mostly on the basis of signal modulation settings and strings detected in raw signal recordings.

5.3.2 Data Extraction and Signal Interpretation

Both of the targeted maritime VSAT operators in our study employed a modern protocol stack which combined the newer DVB-S2 standard (formalized in 2005 to replace the original 1995 DVB-S standard) with adaptive coding and modulation (ACM). Data was further encapsulated into generic continuous streams using the generic stream encapsulation (GSE) protocol first proposed by the European Telecommunications Standards Institute in 2007 [170, 177, 178]. An overview of this encapsulation method can be found in Figure 5.5.

Unlike older multi-protocol encapsulation (MPE) streams, to our knowledge no publicly available software for receiving and interpreting satellite data feeds in this format exists. As a result, we developed *GSEextract*, a set of python utilities that permit the extraction of arbitrary IP data from raw recordings of GSE continuous streams. For those feeds most commonly used in maritime VSAT service, *GSEextract* allows an attacker to reliably interpret a significant portion of broadcast data with comparatively low quality satellite television equipment.

It bears mentioning that *GSEextract* is not merely a naive implementation of the DVB-S2 and GSE standards. Rather, the utility leverages several assumptions about maritime VSAT implementations to enable the recovery of arbitrary IP packets in the presence of frequent signal processing failures. A detailed description of these assumptions and the technical implementation of *GSEextract* can be found in Appendix B. *GSEextract* would be ill-suited as a utility for operating a maritime VSAT internet service due to these assumptions, but it performs well as a forensic tool. Two of the core strategies employed are the use of a known valid MATYPE header as a “crib” for re-synchronization within corrupted streams and the intelligent




Figure 5.6: The degree to which GSE packets within a given stream were recoverable. Stream 4 was of significantly lower throughput than the others and was included to assess GSEextract’s performance in lower bandwidth contexts.

padding of internal payload data to construct valid packets when data fragments are missed by the radio receiver.

5.3.3 Collection and Forensic Performance

For an initial assessment of GSEextract’s performance, we elected to record 24 hours of data from the two transponders on each of the two targeted satellites which offered the strongest and most reliable signal (as indicated by signal-to-noise ratio) at our research site in Europe. In total, this amounted to 96 hours of maritime traffic recordings and approximately 300 GB of reconstructed packet captures. As anticipated in Section 5.3.1, recordings made with consumer-grade hardware were imperfect, with significant data loss. GSEextract interfaced with raw DVB-S Baseband Frame recordings made by the TBS-6903 card as no software was found capable of processing higher layers from the corrupted recordings. Nevertheless, GSEextract was able to extract between 40-60% of the GSE Protocol Data Units (PDUs) contained within the targeted streams and partially recover a further 10-25% of corrupted PDUs (Figure 5.6).




Figure 5.7: The overall proportion of successfully reconstructed IP payload bytes using GSExtract. These metrics were only calculable for successfully identified IP packet headers and do not apply to “unrecoverable” GSE packets lost in the signal processing stage (see Figure 5.6).

We lacked ground-truth regarding the quantity of internet traffic transmitted, which made it difficult to determine what proportion of a VSAT feed was successfully picked up by the employed hardware. However, a proxy metric can be derived based on the number of padding bytes injected by GSExtract into a recovered capture. In the case where a large number of IP packets were corrupted, it is expected that GSExtract will inject a correspondingly large number of bytes into the resultant .pcap file when reconstructing partial IP payloads. In the case where most IP packets are recovered successfully, GSExtract will not add many additional bytes. This metric suggests that, at the IP packet level, GSExtract recovered on average, approximately 92% of any given IP payload. However, in terms of overall data volume we estimate that GSExtract was able to reconstruct between 60% to 85% of bytes transmitted on a given frequency (Figure 5.7). Performance was roughly correlated with signal quality, with the lowest-quality data signal also showing significantly higher rates of data corruption using GSExtract. Additional variance in performance measurements may result from specific network properties and




Figure 5.8: The average percentage of a given IP packet which is unrecoverable by GSExtract. Each point represents the mean padding applied by GSE to IP packets within a 10-byte payload size interval and its diameter correlates with the number of packets captured of that size. As expected, GSExtract’s performance worsens as IP payloads grow, since BBFrames containing later fragments are more likely to be missed. Note that around 1.4kb, IP packets which are close to the size of a single BBFrame may be recovered more reliably from some providers as they are allocated the entirety of a single BBFrame with no fragmentation. The precise maximum size of a BBFrame is dynamically dictated by network conditions, possibly explaining the high variance observed around this threshold.

behaviors (e.g., use for video streaming vs. web browsing) across each signal.

This discrepancy between the average recovery rate in terms of packets compared to in terms of bytes results from the use of fragmentation in GSE. Specifically, the IP packets most likely to be recovered by GSExtract were smaller packets which could be transmitted entirely within in a single DVB-S baseband frame (BBFrame). This size varies, often minute-to-minute, depending on network traffic conditions. Generally, however, as an IP packet gets larger, the probability of fragmentation increases. The more fragmented an IP packet is, the more likely that one of those fragments is not picked up by the signal hardware. The strength of this relationship can be observed in Figure 5.8.

Even in the case of fragmented packets, however, GSExtract is often able to identify and recover significant portions of the lost payloads. While there is no state of the art for comparison, one would expect a naive decoder to have higher




Figure 5.9: A comparison of packet loss with and without GSEExtract’s PDU recovery. The solid line depicts a naïve decoder which employs GSEExtract’s basic re-synchronization strategy but no other forensic techniques. The dashed line depicts GSEExtract’s performance, indicating only those packets for which partial recovery was impossible.

errors at higher degrees of fragmentation. In contrast, GSEExtract breaks this positive correlation and allows for reliable rates of partial recovery regardless of fragmentation rates (Figure 5.9). Even in highly fragmented and unreliable streams, GSEExtract successfully identifies and partially reconstructs between 84% and 92% of received GSE PDUs. In essence, GSEExtract “makes use” of the vast majority of traffic which is successfully demodulated by the satellite hardware. It is only in cases when the IP header itself is not received by the satellite hardware that a payload is fully “unrecoverable” (see Figure 5.6).

5.3.4 Extended Collection

In addition to the four initial experimental feeds, we recorded a continuous week of traffic from each service provider. This was devised to support deeper measurements into traffic patterns and behaviors over time. In total, this provided approximately 1.3 TB of data and more than half a billion DVBS-2 messages for analysis.

Beyond storage costs, there is no practical limitation on an attacker’s ability to record data using this method. Even in the case of complete signal interruption or

loss (such as in the event of adverse weather), GSExtract is capable of automatically reconstructing and resuming analysis of broken GSE data streams. While beyond the scope of this security analysis, GSExtract may thus be well suited to multi-month longitudinal measurement studies of traffic trends within the maritime ecosystem. Additionally, while a single satellite dish can only tune to one channel at a time (acting as a practical constraint on the amount of data which can be collected), significantly more data might be captured through the use of multiple dishes simultaneously. VSAT-specific signals intelligence (SIGINT) collection platforms sold to nation-state security services likely also have this capability, albeit at costs far beyond the reach of our proposed threat model [179].

5.3.5 Ethical and Legal Concerns

On account of the real-world networks in which we conducted our experiment, all relevant legal regulations surrounding traffic collection and analysis in the jurisdiction of our research were strictly adhered to. Given that we had no prior indication of the sensitivity of information in maritime VSAT feeds, we treated all collected data as if it contained sensitive information. No information was stored longer than necessary, and we made no attempt to decrypt data — even in cases where encryption appeared weak or improperly implemented.

To the best of our ability, we have attempted to responsibly disclose these findings to service providers and individual maritime customers impacted by our research. Many contacted organizations have expressed surprise at the findings and an interest in taking steps to mitigate them. This research has led to conversations with C-Suite executives at some of the world's largest businesses, suggesting that this attack vector on ship-to-shore communications is novel and of particular concern to maritime industry participants. Our goal with this work is not to focus on specific companies, providers, or implementations but to raise awareness of insecure industry standards used for potentially sensitive data transmissions. As a result, we have elected to withhold the specific names and transmission frequencies of affected services providers to keep focus on the identified issues.



TLP:WHITE

Private Industry Notification

FEDERAL BUREAU OF INVESTIGATION, CYBER DIVISION

14 February 2020

PIN Number
20200214-001

Please contact the FBI with any questions related to this Private Industry Notification at either your local **Cyber Task Force** or **FBI CyWatch**.

Local Field Offices:
www.fbi.gov/contact-us/field

E-mail:
cywatch@fbi.gov

Phone:
1-855-292-3937

The following information is being provided by the FBI, with no guarantees or warranties, for potential use at the sole discretion of recipients to protect against cyber threats. This data is provided to help cyber security professionals and system administrators guard against the persistent malicious actions of cyber criminals.

This PIN has been released **TLP:WHITE**: The information in this product may be distributed without restriction, subject to copyright controls.

VSAT Signals Vulnerable to Low-Cost Device Exploitation

Summary

The FBI has identified a potential increased risk to data transmitted by Very Small Aperture Terminals (VSAT). Previously, the cost of the satellite equipment needed to intercept the data from these terminals served as a barrier for threat actors. However, recently conducted research discovered man-in-the-middle attacks against maritime VSAT signals can be conducted with less than \$400 of widely available television equipment,³ presenting opportunities to a wider range of threat actors to potentially gain visibility into sensitive information. VSATs are commonly used within the maritime and aerospace industries, predominantly seen in airplanes, cargo shipping, cruise ships, and offshore oil drilling platforms for a variety of services as well as point of sale systems in the retail sector. VSAT networks generally utilize Transmission Control Protocol (TCP), Internet Protocol (IP), and radio frequency (RF) channels to transmit data.

Figure 5.10: An FBI industry notification circulated regarding the research findings in this chapter.

Several months in advance of our publication of this research, a United States Federal Bureau of Investigation (FBI) threat intelligence report was circulated warning the wider maritime industry of specific technical details of our findings (Figure 5.10). While we had not contacted the FBI directly, it is possible one of our responsible disclosure contacts circulated a copy of our unpublished research to them.

5.4 Threat Model and Attacker Capabilities

This experiment focused on a threat actor with a relatively low degree of sophistication. Beyond the aforementioned assumption that the attacker was resource

constrained to consumer-grade equipment, we also assumed that the attacker was not capable of directly interfering with the operation of the satellite network itself. That is to say, the attacker is *passive* with regards to satellite signals and cannot directly inject, spoof, or interrupt radio emissions. Future experimentation considering the possibility of an active attacker may prove valuable but would be difficult to conduct safely and legally in real-world maritime VSAT networks.

While our threat model assumes a passive attacker in the satellite context, we grant the attacker the ability to engage in *active* attacks against internet-connected systems. For example, if the attacker observes confidential information in a satellite feed, we consider how that information might be abused to impact publicly routable maritime platforms.

Our threat model did not focus on any specific operational motive for the attacker beyond that of an honest but curious observer. However, as mentioned in Section 5.1, significant concerns have been raised regarding the threats posed by criminals, pirates, and terrorists to critical maritime systems. Throughout the chapter, we note findings that appear intuitively relevant to such specific threats.

5.5 Broad Findings and Vulnerabilities

All four maritime VSAT networks included in our experiment did not appear to apply encryption by default. Moreover, a superficial review of an additional 11 VSAT network streams did not uncover any fully encrypted maritime VSAT services. While we cannot determine the full extent to which the providers we selected are representative of the global VSAT industry, especially given our geographic focus on Europe and the North Atlantic, this suggests that a large portion of maritime VSAT signals transmitted using GSE are inadequately protected. Given that the underlying routing equipment used in these networks accounts for more than 60% of the global maritime VSAT market, and is used by eight of the ten largest VSAT providers, we expect that findings on these networks have wide-ranging applicability to the industry [180]. Moreover, one of the satellites included in our study was




Figure 5.11: The 20 most commonly observed protocols across all collected VSAT signals. Note that sessions are counted on a log scale.

launched within the past 3 years, suggesting that these findings are not merely representative of security issues in legacy systems.

5.5.1 Applications and Protocols

The principal protocols identified in our recordings are outlined in Figure 5.11. To some extent, traffic transmitted over maritime VSAT network is similar to that which would be observed by any other ISP. For example, maritime VSAT terminals are used by crew and passengers for the purposes of general web browsing, media streaming, and personal communications. Of course, it is unusual for an attacker to have the vantage-point of an ISP-level eavesdropper, especially over a coverage area of millions of square kilometers.

However, there are some important differences in the use and operation of maritime networks. Maritime VSAT services are sold as a component of internal business technical infrastructure as well as external connections to the wider internet. As a result, maritime VSAT traffic includes not only general access to internet services but also internal business communications. A traditional approach to designing and securing business networks by, for example, defending

the perimeter between the business LAN and the internet, may not easily translate to VSAT architectures.

The effect of this difference is demonstrated by contrasting the protocols used to access IP addresses within the satellite network with those located outside it (Figure 5.12). We observe a much higher usage of unencrypted protocols, such as HTTP and clear-text POP3 (as opposed to HTTPS or POP with TLS), when both participants are “local” to the VSAT network than when one of the participants sits external to the satellite environment. This may suggest that maritime operators consider VSAT networks to operate in a manner akin to a corporate LAN environment and are unaware that these networks are subject to over-the-air eavesdropping.

Less broadly, maritime networks differ from terrestrial networks in that communications serve several unique functional purposes in maritime environments. Thousands of specialized applications designed to enable the remote monitoring and operation of various ship components rely on maritime VSAT networks to communicate with terrestrial offices or other ships in a fleet. Given this technical diversity, it is difficult to exactly characterize which captured traffic belongs to which applications. However, an overview of some common maritime and terrestrial application functions observed in GSExtract’s captures appears in Table 5.2. More detailed case studies of specific services can be found in Sections 5.6 and 5.7.

5.5.2 Hosts and Vessels

Despite prior research suggesting that larger maritime organizations are more confident in their cyber-security controls than smaller ones, we observed sensitive data originating not only from small fleets, but also from some of the world’s most significant maritime operators [168]. These included three members of the Fortune Global 500 and at least six publicly traded entities with combined annual revenues exceeding \$700 billion [181]. In the cargo sector alone, we observed sensitive traffic from organizations which, combined, account for more than one-third of all global maritime shipping.

Table 5.2: Frequency Breakdown for Selected Applications.

Application/Protocol Metric	Observed Quantity
Electronic Navigational Chart (ENC) File Transfers	15,344 ENC Files
Automatic Identification System (AIS) Geolocation Update Messages	4,245,273 Messages
Session Initialization Protocol (SIP) Conversations (Voice over IP Protocol)	150,832 Sessions
Email Protocol Conversations (Both Encrypted and Unencrypted)	704,845 Sessions
Unique Email Addresses from Unencrypted POP, SMTP & IMAP sessions	17,501 Addresses
Connections to “Big 5” Owned IP Addresses (Google, Amazon, Facebook, Apple, Microsoft)	18,993,774 Sessions
Unique Hostnames from DNS Responses	278,337 Hosts




Figure 5.12: A comparison of the ratio of sessions using unencrypted protocols vs encrypted alternatives on the basis of whether a session is contained within the local IP range or reaches out to globally addressable IPs. A higher preference for unencrypted protocols is observed in “internal” VSAT traffic. Note that this ratio is expressed on a log scale.

In total, GSExtract identified more than 9,000 distinct hosts belonging to the VSAT network which participated in 50 or more sessions over the recording window. More than 4,000 participated in at least 500 sessions and more than 400 had publicly accessible IP addresses. Although ships may occasionally have multiple VSAT terminals aboard, these numbers suggest that thousands of distinct marine vessels were included in our traffic recordings. Due to overhead and latency concerns, and the general broadcast nature of satellite communications, VSAT networks generally rely on static IP address allocations (as opposed to, for example, DHCP). As a result, IP addresses roughly map to physical host routers or devices.

As every ship has distinct technologies aboard, fully automating the identification of ships based on their internet traffic is likely impossible. However, an attacker would naturally have an interest in linking intercepted traffic to a physical vessel at sea. In order to characterize the difficulty of this task, a random sample of 100 host IP addresses was selected from the traffic. The following basic metadata characteristics were then extracted:

- Top 10 Source and Destination Autonomous System Numbers (ASNs)
- Top 50 TLS Certificate Alternative Names
- Top 50 TLS Subject Common and Object Names
- Top 50 TLS Issuer Common and Object Names
- Top 50 DNS Query Host Names
- First 2000 Unique 7+ Character Strings Captured

Using this basic metadata, it was possible to glean significant information about individual vessels. For 62 of the 100 hosts, this data was sufficient to characterize what types of computing devices might be on board. In some cases (17), it was only possible to determine the general operating systems used by devices on board (e.g., Windows 10, Android). However, one could often determine individual software programs running on these hosts and even fingerprint specific software versions. Indeed, for three of the hosts, Common Vulnerabilities and Exposures (CVE) reports were identified as likely exploitable against specific software aboard the ship.

More practically, about a quarter of the analyzed hosts (26) could be tied to specific owners or fleets, permitting an attacker to target specific companies or industries. These organizations were spread over eight broad industries: Oil & Gas, Cargo, Chemical Shipping, Government, Fishing, Subsea Construction, Maritime Support, and Offshore Wind Power. Moreover, the companies hail from 11 different countries (Germany, United Kingdom, Netherlands, Korea, Norway, Spain, Bermuda, Pakistan, Switzerland, Poland, and Italy). The largest employs more than 70,000 individuals while the smallest operates only a single fishing vessel.

12 of these hosts could be further associated to specific vessels (or, in one case, a remote polar research station). These vessels are summarized in Table 5.3 and allude to the diversity of maritime organizations vulnerable to this threat.

Simple extrapolation suggests that, using only cursory manual analysis, a dedicated attacker could expect to identify more than 1,000 vessels in the sample traffic collected for this study. Moreover, this is likely a lower-bound. A deeper manual review of traffic from a given host may permit an attacker to identify the associated customer and ship with even greater reliability (albeit at the cost of increased investigation time).

This experiment was devised with dual purposes. First, to identify security issues that might endanger the physical safety of crew and ships using maritime VSAT connections. Second, to identify less serious but significant issues which might undermine the data privacy and network security of maritime VSAT customers. While there may be significant overlap between these two categories, we have attempted to divide our findings according to each for clarity.

5.6 Findings: Physical Safety and Operations

Section 5.1 notes a significant lacuna between prior work acknowledging the theoretical desire of cyber-attackers to target maritime vessels and technical research discussing the mechanism by which such attacks might manifest. Our experimental findings suggest that attacks against maritime VSAT communications may be one such mechanism and that securing maritime VSAT is not just important to

Table 5.3: Specific vessels identified from 100 randomly selected host addresses in case study.

Vessel ID*	Vessel Type	Gross Tonnage	Operator Industry	Operator Fleet Size	Example of Identified Client Software Information	Notable Traffic Observations
1	Subsea	22,000t	Oil & Gas	70 Vessels	Specialized Maritime Software	Unencrypted Netlogon Traffic
2	Container	150,000t	Shipping	250 Vessels	PLC Firmware Binaries	“Cargo Hazard A, Major” In Cargo
3	Icebreaker	9,000t	Research	Government	IT Support Software	Unencrypted SMB Fileshares
4	Firefighter	8,000t	Oil & Gas	70 Vessels	Specialized Maritime Software	Unencrypted SQL Database Replication
5	Seismic	8,000t	Seismic	10 Vessels	Antivirus Software & Version	Unencrypted Email Conversations
6	Chemical	5,000t	Shipping	1 Vessels	PLC Firmware Binaries	Unencrypted PLC Firmware Update
7	Outpost	(Island)	Research	N/a	OS Minor Version Numbers	Polar Island Research Station
8	Container	33,000t	Shipping	600 Vessels	Messaging Software	Unencrypted REST API Credentials
9	Fishing	1,300t	Fishing	1 Vessel	OS Major Version Numbers	Unencrypted Email Conversations
10	Chemical	17,000t	Shipping	10 Vessels	Specialized Maritime Software	Unencrypted Fileshare Credentials
11	Container	110,000t	Shipping	500 Vessels	Maritime Navigation Software	Unencrypted Email Conversations
12	Subsea	22,000t	Oil & Gas	70 Vessels	Firewall Software & Version	Vulnerable Windows Server 2003

*Note: Vessel names have been withheld and fleet sizes and tonnage are approximate due to privacy concerns.

protecting directly networked devices, but also to the wider physical safety of ship and crew. Specifically, we consider two targets: the navigational and charting systems used to safely route vessels at sea and sensitive operational information regarding cargo contents or security procedures aboard a vessel.

5.6.1 Navigation and Charting

In the context of ship navigation, maritime VSAT services are used to provide real time data regarding the location of other vessels, optimal routing plans, and accurate nautical charts. These critical operational links have a direct influence on the ability of modern vessels to operate safely and reliably. Attackers who could undermine the reliability of navigational data aboard their victim's vessels could cause serious harm to both their victims and the general public. For example, a terrorist organization which altered nautical charts to cause an oil tanker to run aground on a hidden reef would have a catastrophic environmental impact. Similarly, pirates with the ability to view, or even alter, planned routes for cargo vessels could determine an optimal time and location to attempt seizure. For example, traffic intercepted from a multi-million dollar yacht in the traffic captures included detailed itinerary plans for upcoming destinations.




Figure 5.13: A map of AIS positions reported in one VSAT stream with a heavy concentration of reported vessels near the Faroe Islands. A total of more than 4 million AIS messages were identified in the study.

As mentioned in Section 5.1, there is significant interest in AIS positional traffic. Our traffic captures included more than 4 million AIS messages describing the locations of various marine vessels. A map of some of these signals can be

found in Figure 5.13. These messages mostly appeared to be transmitted from terrestrial web-servers to AIS navigational appliances aboard various vessels. If an attacker managed to transmit additional AIS messages on these streams (see Section 5.8), they might maliciously conceal or artificially introduce vessels into the charting maps aboard a targeted ship.

It has been previously suggested that attackers might abuse Electronic Chart Display and Information Systems (ECDIS) to cause vessels to collide with undersea hazards [163]. However, to our knowledge, no practical mechanism for attacking such systems has been identified to date. ECDIS has come to replace paper nautical charts on modern vessels and is a vital component of safe marine navigation. One of the principal advantages of modern ECDIS systems compared to paper charts is the ability to have frequently updated and interactive data enabled by the use of VSAT connectivity. These updates include critical safety messages called Notices to Mariners (NMs) which relay details regarding developing nautical hazards.

```
> Transmission Control Protocol, Src Port: 21, Dst Port: 41573, S
v File Transfer Protocol (FTP)
  v 257 "/Inbox/chartdelivery" is current directory.\r\n
    Response code: PATHNAME created (257)
    Response arg: "/Inbox/chartdelivery" is current directory.
```

Figure 5.14: Traffic from an FTP-based ECDIS update. This system is likely trivially vulnerable to the attacks in Section 5.8.

While every ECDIS product is different, the traffic observed in our study suggests that several commonly used ECDIS platforms are trivially vulnerable as a result of information leakage over maritime VSAT networks. In several cases, ECDIS chart updates were transmitted over the unencrypted POP3 e-mail protocol. In many of these instances, files which were appropriately named and sent to the correct POP3 inbox are automatically downloaded and used by the targeted ECDIS. In other cases, updates must be manually copied by a crew member onto an external storage device from the e-mail inbox and inserted into the appropriate ECDIS device — often on a regularly scheduled basis. We also found several instances in which ECDIS charts were updated via insecure FTP connections with or HTTP APIs (Figure 5.14). Were

```

HTTP/1.1 200 OK
Transfer-Encoding: chunked
Content-Type: application/octet-stream
[REDACTED]
[REDACTED]
[REDACTED]
GAS PIPELINE
Anchoring and any use of gear towed on the bottom is prohibited
in the protection zone which extends 200 metres on each side of
the pipeline. Gas from a damaged pipeline could cause an explosion,
loss of a vessel's buoyancy or other serious hazard.

```

Figure 5.15: A captured NM which was transmitted via a clear-text HTTP API. This system is likely trivially vulnerable to the attacks detailed in Section 5.8.

an attacker to submit maliciously altered files via any of these update mechanisms they would be able to alter the nautical maps used to navigate the victim’s vessel.

A public standard for the cryptographic verification of ECDIS charts exists (IHO S-63) and would mitigate such attacks [182]. The S-63 standard was developed with the explicit goal of preventing malware from causing harm to vessels and is an addition to an older unsecured format (S-57) [183]. S-63 implements a public-key signing system to facilitate client-side verification of chart authenticity and integrity.

Nevertheless, catalog references to more than 15,000 charts in the unauthenticated S-57 format appeared in our traffic captures. Moreover, many popular charting services do not use either the S-57 or S-63 standards but instead use their own proprietary formats. A cursory inspection of two such vendor-specific formats suggested that no cryptographic verification system was employed. For example, Figure 5.15 depicts an NM alert which is transmitted via an unsecured web API.

Future systematic work investigating the robustness of these proprietary formats against data tampering may provide valuable context for maritime charting customers. Regardless, these findings provide a clear practical demonstration of the importance of employing S-63 or comparable verification standards, even for “air-gapped” or otherwise secured ECDIS with low risks of malware compromise.

5.6.2 Vessel Operations and Security

Beyond navigation and charting, many other aspects of day-to-day modern ship operations rely on VSAT connectivity and, in the context of unsecured VSAT

```
2nd
Engineer","phone":null,"createdDate":1555016097,"inactive":false,"pictureUrl":null,"presencelog":
{"id":"██████","crewmemberId":"██████","present":true,"date":1556830579}},
{"id":"██████","groupId":"██████","idealId":null,"badgeld":null,"order":39,"lunchOrder":null,"
firstName":"H██████","lastName":"D██████","job":"Chief
Stewardess","phone":null,"createdDate":1556961769,"inactive":false,"pictureUrl":null},
{"id":"██████","groupId":"██████","idealId":null,"badgeld":null,"order":40,"lunchOrder":null,"
firstName":"M██████","lastName":"K██████","job":"Stewardess","phone":null
```

Figure 5.16: A portion of the crew manifest from software aboard a \$50 million luxury yacht which was captured during the experiment.

transmissions, may present a security threat to the safety of ship and crew. Even simple data that does not appear intuitively sensitive, such as a manifest listing personnel aboard a vessel, can provide a dangerous advantage to pirates assessing their ability to overwhelm the crew of a targeted ship (Figure 5.16).

The regular transmission of cargo manifests and other information required by various port authorities could allow attackers to identify targets of interest. We regularly observed cargo manifests discussing the contents of vessels, normally in the form of e-mail attachments or encapsulated in the traffic of various proprietary fleet management software products. In one illustrative example, we observed a vessel transmit a report indicating it was transporting hydrogen sulfide (Figure 5.17). The Islamic State has previously attempted to manufacture or acquire hydrogen sulfide for the purpose of developing chemical weapons [184]. While the particularities of chemical weapons development are far beyond the remit of this research, the leakage of such information raises intuitive concerns.

5.7 Findings: Passenger and Crew Privacy

Like many large organizations, maritime companies frequently handle sensitive data concerning their customers and employees. Unlike other large organizations, a significant portion of this data is transmitted over-the-air and, in the case of VSAT connections, can be physically intercepted by attackers thousands of miles away. The general susceptibility of maritime VSAT connections to eavesdropping

```

<p class=3D"MsoNormal"><span
style=3D"font-family:&quot;Courier
New&quot;;t=
ext-transform:uppercase">THE CARGO MAY
POSSIBLY CONTAIN H2S. THE MASTER AND=
CREW SHOULD THEREFORE ENSURE THAT ALL
PERSONNEL HANDLING CARGO SHOULD BE M=
ADE FULLY AWARE OF THE HAZARDS AND
ADVICE RELATED TO H2S AS OUTLINED IN
THE LATEST EDITION OF ISGOTT. ALL
RELEVANT PRECAUTIO= NS, AS RECOMMENDED
BY THE LATEST EDITION OF ISGOTT, MUST
BE TAKEN WITHOUT E=
XCEPTION.<o:p></o:p></span></p> <p
class=3D"MsoNormal"><span
style=3D"font-family:&quot;Courier
New&quot;;t=
ext-transform:uppercase">&nbsp;<o:p></
p></span></p> <p

```

Figure 5.17: A portion of a risk assessment document captured during the experiment indicating the presence of hazardous materials aboard a vessel.

```

CID Number: [REDACTED] Rank: COFF Name: S [REDACTED] &nbsp;<br>
Passport: Z [REDACTED] Issued: 05 [REDACTED] Expiry: 04 [REDACTED] <br>
Seaman book: [REDACTED] Issued: 04 [REDACTED] Expiry: 03 [REDACTED] <br>
Nationality: [REDACTED] Date of birth: [REDACTED] Place of birth: [REDACTED] <br>
<br>
CID Number: [REDACTED] Rank: 2OFF Name: [REDACTED] JL&nbsp;<br>
Passport: R [REDACTED] Issued: 14 [REDACTED] Expiry: 13 [REDACTED] <br>
Seaman book: [REDACTED] Issued: 24 [REDACTED] Expiry: 23 [REDACTED] <br>
Nationality: [REDACTED] Date of birth: [REDACTED] Place of birth: [REDACTED] <br>

```


Figure 5.18: A redacted instance of passport and crew member data intercepted during the experiment.

thus raises serious privacy concerns and suggests that maritime VSAT traffic may be a target for cyber-criminals and identity thieves.

For example, ships crossing international borders must maintain information regarding the visa and passport details of their passengers and crew members. This data is frequently transmitted along ship-to-shore links in anticipation of arrival at a given port. Despite the sensitivity of this data, in a single 24-hour window, we were able to find more than a dozen instances of complete passport details transmitted in plain-text across VSAT connections (Figure 5.18).

Consumer-oriented maritime businesses, such as ferries and cruise ships, rely on the ability to sell goods and services to passengers as a component of their revenue stream. As such, they must handle and verify credit-card payment details while at

sea and VSAT technology is used to facilitate this service. Figure 5.19 depicts one of more than 12,000 messages observed from on-board credit card readers captured during the study. Reverse engineering the communications protocol employed by these machines was beyond the scope of this project, but the presence of this traffic suggests that sensitive financial data may not be adequately protected over VSAT links. Similar issues with secure transaction handling have been previously identified in the aviation sector over an unrelated terrestrial radio protocol [169]. This suggests that, despite the general availability of encryption technology for sensitive data, a lack of customer awareness regarding data link security for esoteric and domain-specific contexts may cultivate risky practices.



..116
[REDACTED]
[REDACTED] 02B34;37
.AMERICA N EXPRES
S.34 [REDACTED]
[REDACTED] 100.
[REDACTED]
[REDACTED]
00Y.NNN. 000000.0
[REDACTED]

Figure 5.19: A heavily redacted screenshot of traffic from a handheld credit card reader belonging to a major cruise line. More than 12,000 such messages appeared in the study.

Internal network traffic relating to the business operations of a maritime organization may also contain deeply sensitive information. While the majority of email protocol traffic was encrypted, more than 130,000 unencrypted email sessions were identified within the experimental recordings. This included deeply sensitive information such as a password reset link for the Microsoft account belonging to the captain of a multi-million dollar yacht and candid discussions between oil company leadership discussing a recent accident leading to the death of a crew member. That this information was broadcast in plain-text over an entire continent is deeply concerning.

Email was only one of many contexts in which sensitive business information was leaked over VSAT connections. For example, one organization used VSAT linkages to replicate employee intranet profiles across their vessels and, as a result,

leaked hundreds of employee emails, usernames, addresses, next of kin information, and password hashes. Likewise, more than 95,000 unencrypted FTP sessions were observed — many of which were used to propagate updated information about crew members and user accounts across an entire fleet. Although encrypted alternatives to these protocols are widely available, many maritime organizations do not employ them in practice.

One encryption protocol was widely employed, with TLS ranking as the third most common protocol in our dataset. However, even in this case, a cursory analysis identified frequent issues within implementations. Of the approximately 30 million TLS sessions observed, around 9% used cipher protocols generally considered to be weak or insecure [185]. Restricting the analysis to only “internal” traffic local to the maritime VSAT network, the prevalence of weak or insecure cipher suites increased substantially to 36%. Legal constraints prevented closer investigation into the practical exploitability of these ciphers but future work here may prove fruitful.

5.8 Active Attacks on VSAT Services

Beyond passive eavesdropping, an attacker may also wish to directly interfere with active VSAT communications links. However, for a low-resourced adversary, there are several barriers to doing so.

First, the non-broadcast components of the feed (e.g., the uplink connection from the ship to a satellite, or the downlink connection from the satellite to a ground station) are highly directional signals. To intercept or spoof these components would likely require the use of aerial vehicles sitting in the line-of-sight from a vessel to the satellite or ships which have been strategically deployed to listen on antenna side-lobes from the VSAT dish located on a target vessel. Moreover, successfully replicating the modulation states and signal characteristics of a satellite feed in real time would require access to expensive and sophisticated radio equipment. Given these constraints, the threat of an active attacker in VSAT environments has historically been of little concern.

5.8.1 TCP Session Hijacking

Using our experimental setup, we successfully demonstrated the capability of an attacker to arbitrarily modify traffic in a real-world maritime VSAT environment through TCP session hijacking. While the process of TCP hijacking is well understood, these attacks are rarely practical in terrestrial ISP networks due to challenging race-conditions.

The unique physical properties of satellite networks offer a substantial change to this threat model as an attacker is almost guaranteed to “win” the race to hijack the session (Figure 5.20). Speed-of-light delays over the satellite link are significant. For the 425 publicly routable hosts in our captures, the mean round trip time (RTT) was approximately 725 ms and the median RTT approximately 700 ms. This grants an attacker around 350 ms to send their malicious TCP responses. Even under ideal theoretical conditions, RTTs to geostationary orbit measure upwards of 500 ms.

5.8.2 TCP Hijacking Requirements

A maritime VSAT network is only vulnerable to TCP hijacking attacks under certain conditions.

Firstly, an attacker must determine public IP routes to both ends of a targeted TCP conversation. Generally, this requires vessels within the network to have public IP addresses. However, it may also be possible to identify IP mappings through a Network Address Translation (NAT), albeit with significantly more effort. For example, in the experimental captures, public IP routes to internal hosts were occasionally leaked inside SMB file paths and HTTP headers. Interestingly, many of these leaks originated from malware traffic scanning for vulnerable hosts, indicating that an organizational policy to use encrypted application layer protocols (e.g., HTTPS) may not be sufficient to fully hide IP mappings.

Unique to satellite ecosystems, there is also the risk that the observed TCP session over the air is not the same session as that observed by the receiving vessel and internet endpoint. This is due to the use of Performance Enhancing Proxies (PEPs). PEPs modify TCP connections and generate artificial ACK responses


in the TCP three-way handshake in order to prevent high latency from being misinterpreted as a sign of network congestion by the TCP protocol.

PEPs can vary significantly. First, they may modify traffic at either the client, the ISP gateway, or both. Additionally, they can either “split” traffic into distinct TCP sessions — generating unique sequence numbers and handshakes for both sides — or “snoop” into TCP sessions, operating invisibly and preserving TCP header information across the entire link. In the former case, TCP session numbers transmitted over the satellite link may not be the same as TCP session numbers expected by either or both of the session endpoints. This can either prevent a hijacking attack entirely (if the connection is “split” into three hops), or limit attacks to a single direction (if the connection is “split” into only two hops).


In our study, approximately 425 hosts, or around 5% of observed hosts, had publicly routable IP addresses. However, this is likely not a representative ratio as the provisioning of public IP addresses varies substantially between VSAT providers. Among 11 other VSAT service providers which were considered but not selected for long-term recordings, approximately a third provided clients with publicly routable IP addresses. For legal and ethical reasons, we did not attempt to fingerprint PEP software on individual hosts as this requires active port scanning and connections to customer endpoints.

5.8.3 Hijacking Implementation

To hijack TCP sessions, GSExtract monitors live VSAT traffic for TCP SYN connections from a specified internet host to a specified VSAT target. It extracts the appropriate sequence number from this intercepted data and uses it to transmit an artificial TCP SYN-ACK response to the internet host. This malicious response reaches the internet host hundreds of milliseconds before the legitimate response completes its 70,000 km journey through geostationary orbit. A similar process is used to intercept the final ACK response of the three-way handshake and all subsequent TCP packets.



(a) A TCP-SYN packet and associated sequence number sent from the terrestrial back office arrives at both the legitimate recipient and the eavesdropper simultaneously.



(b) The attacker generates a SYN-ACK response with the received sequence number and transmits it over a low-latency wired internet connection. Meanwhile, the legitimate recipient also generates a SYN-ACK response and sends it via the VSAT link. Due to speed-of-light effects, the attacker's response is virtually guaranteed to arrive first. At this point, the attacker has hijacked the TCP conversation.

Figure 5.20: Notional Overview of TCP-Hijacking in VSAT.

In order to responsibly assess this threat in a real-world VSAT network, we elected to hijack our own attempted connection to a closed TCP port aboard a remote vessel. Specifically, we generated malicious responses to our own HTTP requests sent to an IP address located within the VSAT environment. This allowed us to successfully generate traffic which appeared to be from a web server running aboard a vessel operating within the customer network. This sort of attack could be used to falsely report location details or other ship status information to a terrestrial operations center.

TCP session hijacking also enables other attack vectors, including command injection into telnet sessions and man-in-the middle attacks on certain SSH configurations. In the context of our aforementioned findings, TCP hijacking may represent a mechanism for maliciously altering ECDIS navigational charts, NM alerts, AIS area reports, or other operationally vital information. Additionally, a trivial denial of service attack can be achieved through the introduction of malicious TCP RST packets. An attacker could thus significantly reduce the reliability of all TCP connections to a maritime vessel. It may even be possible for an attacker to completely block TCP connectivity to a ship at sea.

We have only assessed our ability to intercept incoming connections from the internet to a host within the VSAT network. We did not interfere with any legitimate uplink connections from vessels as this risked interrupting critical communications and causing harm to end users. Nevertheless, we expect this attack would work equally well for intercepting uplink connections from satellite hosts to the broader internet. While in this direction the attacker's latency advantage would be reduced, the attacker would still have the time advantage of being able to reply immediately to the customer's request rather than routing the request over the open internet and awaiting a response. This suggests an eavesdropper may gain full-duplex access to VSAT TCP streams, despite having the capability to intercept only half of the connection over radio.

5.8.4 Further Active Attacks

Beyond TCP hijacking, other active attacks against VSAT systems appear intuitively possible. For example, at least 30,000 HTTP conversations with session tokens were identified and may be vulnerable in HTTP hijacking attacks. Similarly, DNS responses are regularly observed over the VSAT feed, while predicting DNS queries may be difficult (as these are sent over the uplink and thus not observed in signal captures), certain operating systems (such as older versions of Windows) generate predictable DNS transaction IDs and could accept a malicious response [186]. Further work assessing active attacks in maritime VSAT is likely warranted. However, this would require cooperation from VSAT customers and service providers to conduct ethically.

5.9 Underlying Causes and Next Steps

Increased awareness within the maritime industry is a vital first step to addressing these issues. Based on the content of observed traffic in this study, it appears that maritime VSAT customers are unaware that outsiders can listen in to traffic on their networks — especially when this traffic is logically routed within a LAN environment. In many cases, these issues would be substantially mitigated through the use of application-layer encrypted alternatives, such as requiring the use of TLS for POP3 email sessions or HTTPS for internal web traffic.

However, deeper issues such as the TCP hijacking and denial of service threat or application fingerprinting through the identification of TLS certificates are more difficult to resolve. While VPNs represent an intuitive solution, standard VPN products are incompatible with the aforementioned performance enhancing proxies (PEPs) which are vital to maintaining usable speeds in VSAT environments [100, 110]. The specific challenges of VPN usage in satellite networks are detailed in depth in Chapter 6.

While some proprietary solutions exist, these implementations are not well studied and their security properties are unverified beyond marketing claims [187].

Academic proposals have also been made, particularly around MPEG-TS based communications in the early 2000s, but these have not been updated for newer DVB-S2 and GSE standards [112]. Industry proposals for securing scientific space missions show promise but lack the key-management infrastructure and multiplexing capabilities for multi-user environments [116]. As such, a verifiable and open standard for modern encrypted satellite broadband is much-needed — both within the maritime VSAT context and more broadly.

In the shorter term, especially for sensitive information of the nature outlined in our case studies, maritime VSAT customers may need to accept the significant performance costs of employing IPsec and other end-to-end tunneling techniques over VSAT connections. Higher latency connections may not be desirable from a user-experience perspective, but they are preferable to an alternative which endangers ship and crew.

5.10 Summary

Historically, high costs of access to equipment and esoteric nature of maritime satellite protocols may have acted as significant barriers to entry for threat actors. However, this is no longer the case.

By leveraging inexpensive and widely available satellite television equipment, we have demonstrated that an attacker can eavesdrop on many marine VSAT connections at less than 1% of traditional equipment costs. Further, we have presented GSEextract, a forensic tool which enables the recovery and extraction of significant quantities of valid IP traffic from highly corrupted and incomplete GSE transponder streams. These tools were tested in a real-world environment and used to observe four major maritime VSAT streams providing coverage to Europe and the North Atlantic — together encompassing more than 26 million square kilometers of coverage area. These providers all employ an underlying technology stack used by more than 60% of the global maritime VSAT service industry.

Through this experimental analysis, we discovered that status-quo maritime VSAT networks lack basic link-layer encryption. These issues were contextualized

vis-a-vis their impacts on the safe navigation and operation of vessels and the security and privacy of passengers and crew. Further, we demonstrated the ability to even deny or modify certain ship-to-shore communications depending on VSAT network configuration. In short, the insecure nature of maritime VSAT enables a number of novel threats to marine vessels which may be exploited by a wide-range of relevant threat actors including pirates, criminals, and terrorists.

Our experimental findings suggest that the status quo poses significant risks to some of the world's largest and most vital maritime organizations. To the extent that maritime operators are unaware of the risk exposure caused by eavesdropping attacks on ship-to-shore communications links, we hope this research is a first step towards characterizing the threat. Moreover, we suggest the use of common encryption technologies in the short-term and the need for bespoke protocols in the longer term which handle the unique latency constraints of satellite networking environments.

Technologies linking sea and space have played a defining role enabling the global economy that has shaped modern life. Ensuring that these networks remain defended against ever more sophisticated and capable attackers will be key to preserving these benefits. In the next chapter, we will consider the specific technical barriers to achieving this objective and present a technique for overcoming them.

*Te canam, magni Iovis et deorum
nuntium curvaque lyrae parentem,
callidum quicquid placuit iocoso
condere furto.*

*I sing of you, messenger of mighty Jove and the gods,
father of the curved lyre, skilled at hiding whatever
you please with a witty trick.*

—Horace, *Odes* 1.10

6

Making VPNs Work in GEO: The QPEP Architecture

Contents

6.1	The Need for New Approaches to GEO Encryption . .	121
6.2	System Design Requirements and Related Work . . .	123
6.2.1	The Eavesdropping Threat Model	123
6.2.2	TCP Performance Over Satellite	125
6.2.3	Existing Security Approaches	128
6.3	The QPEP System	131
6.3.1	QPEP Design Contributions	132
6.3.2	Use of the QUIC Protocol	135
6.4	QPEP Implementation	141
6.4.1	System Architecture	141
6.4.2	Error Handling and Session Management	144
6.4.3	Limitations	145
6.4.4	Availability	146
6.5	Secure PEP Testbed	147
6.6	Evaluating QPEP	148
6.6.1	Experimental Setup	149
6.6.2	Baseline Performance	150
6.6.3	Performance Under Adverse Conditions	154
6.6.4	Performance in LEO	158
6.6.5	QUIC Optimizations	159
6.7	Next Steps for Secure PEPs	162
6.8	Summary	163

Historically, security and performance have often traded-off in satellite broadband

networks. As a result, many satellite internet service providers (ISPs) do not offer over-the-air traffic encryption, exposing sensitive customer data to eavesdropping attacks. This is because techniques used to optimize TCP connections in long-distance satellite links are often incompatible with commonly used encryption techniques, such as VPNs.

Since the early 2000s, academics and satellite operators have grappled with the challenge of offering both encrypted and performant TCP over satellite. In Section 6.2 we highlight notable proposals and discuss why they have seen limited real-world adoption. While some encryption systems exist, these follow a “black-box” model and are inaccessible and costly for smaller organizational and individual customers. Moreover, their proprietary nature makes security and performance claims difficult to verify and most permit ISPs to eavesdrop on traffic.

No open-source encryption tool exists for performant TCP communications over satellite links. Although academic proposals are numerous, these are often purely theoretical or lack replicable source-code. As a result, interested researchers must either repurpose outdated code to incorporate modern encryption or reinvent PEPs from scratch. The combined requirements of cryptography and low-level network programming create steep barriers to entry. Moreover, the lack of standardized testing environments makes comparing approaches difficult without privileged access to satellite infrastructure.

The end result is that satellite broadband users have no good options. They (or their ISPs) must purchase expensive and unvetted proprietary applications, accept the substantial performance hit caused by general-purpose VPNs, or transmit sensitive data in clear text over massive satellite footprints.

This chapter seeks to address both the lack of encryption options and high barriers to research in this domain. Its primary contribution is QPEP — an open-source and encrypted-by-default PEP. Unlike many proprietary encrypted PEPs, QPEP is designed for individual satellite customers and conceals traffic from both eavesdroppers and ISPs. Built around the open QUIC transportation protocol, QPEP benefits from robustly vetted cryptographic foundations and a

broad technical community. The system is implemented in Go, an accessible modern language, to facilitate contributions in future research.

As a secondary contribution, the chapter presents an open-source simulation testbed, built around the OpenSAND satellite networking engine [188]. This all-in-one dockerized environment is tailored towards rapid and replicable benchmarking for secure PEP applications. While it has immediate utility in our evaluation of QPEP, it is also designed to ease future system proposals and comparisons from others.

Within this environment, we demonstrate that QPEP achieves its design goals. It nearly halves average page load times compared to traditional VPNs and substantially improves on the performance of even unencrypted PEP applications. Additional simulations are conducted to assess QPEP's performance under various network conditions and in the presence of modifications to the QUIC standard. Ultimately, we find that QPEP represents a promising solution for performant over-the-air encryption in satellite networks while avoiding investment in new infrastructure or alteration of existing protocols.

6.1 The Need for New Approaches to GEO Encryption

The prevalence of security and performance trade-offs in modern satellite broadband networks may initially seem unintuitive. After all, the dangers of unencrypted wireless communications are well understood and have been robustly mitigated in systems ranging from home WiFi to cellular communications. Today, the decision of a terrestrial wireless ISP to offer unencrypted broadband services could, not unreasonably, be attributed to ignorance or incompetence.

This is not the case for long-range satellite communications. Severe security and privacy issues arising from the use of unencrypted broadband services from Geostationary Earth Orbit (GEO) have been known since at least 2005 [136]. However, our own experimental studies in Chapters 4 and 5 have found that, fifteen years later, tens of thousands of satellite customers still rely on unencrypted GEO links. Deeply sensitive data is readily observed by eavesdroppers with access to

simple home-television equipment — affecting customers ranging from individual home internet subscribers to massive corporations.

In the process of responsibly disclosing these vulnerabilities both to satellite ISPs and to their customers, our initial advice was simply to employ proven existing encryption techniques (e.g., IPsec). We quickly learned that this solution was somewhat naive and overlooked significant technical and cultural barriers unique to the satellite context.

In response to our disclosure efforts, satellite ISPs would often espouse the opinion that encryption was a duty which fell to individual customers. They were generally uninterested in deploying costly network-wide security protections, or already offered them, but only as a premium service add-on for particularly risk adverse clients (e.g., military customers). Occasionally, they would emphasize the importance of access to clear-text traffic headers in order to optimize network performance — increasing customer satisfaction.

When speaking with customers, they would acknowledge the value of encryption in the abstract, but were unwilling to accept substantial performance reductions caused by the use of end-to-end encryption tools such as VPNs. In some cases, they had already attempted to deploy VPNs but ended up removing them at the suggestion of their ISP to resolve these performance issues. Indeed, the support pages of many satellite ISP websites suggest the disabling of VPN software as a remediation for slow internet services [189, 190]. The most information-security conscious customers we contacted had attempted to employ piecemeal application-layer protections, such as replacing HTTP web-servers with HTTPS services. However, we found significant gaps in these defenses such as unencrypted DNS traffic or sensitive data from overlooked systems (e.g., a legacy FTP server or POP3 email service).

The principal motivation of this research is to present an approach which considers the unique technical and commercial requirements of these stakeholders. We design and implement a tool which empowers individual customers to encrypt the entirety of their satellite network connection by default, while maintaining

performance that is on-par with, or better than, the unencrypted services they use today. Critically, our design requires no network changes or satellite ISP involvement.

In addition to proposing this new system, the chapter delves into many of the performance characteristics of our implementation. Our motivation in doing so is to demonstrate that our approach offers meaningful performance benefits over traditional VPN tools. Beyond this core hypothesis, significant additional detail is provided to facilitate replicability and comparative benchmarking. We provide open-source implementations not just of our tool, but also for each of the experiments run in the chapter. This is because we recognize that, while our approach is a substantial and needed improvement, it is unlikely the only (or best) way to secure these networks. A key secondary motivation is thus to provide a framework and starting point for others interested in this topic area.

6.2 System Design Requirements and Related Work

Understanding the security/performance trade-off requires a closer look at TCP behavior over satellite. This section provides an overview of our threat model, key defensive challenges, and prior work to address them.

6.2.1 The Eavesdropping Threat Model

Our focus is on broadband provided from platforms in geostationary earth orbit (GEO). The basic operation of GEO broadband can be thought of as a “bent pipe” (see Figure 6.1). As GEO is located more than 30,000 km away from the Earth’s surface, a single satellite has line of sight to a vast area on the surface (theoretically as much as 40% of the Earth’s surface, but practically closer to 20% for broadband communications). This has the advantage of making GEO broadband a relatively inexpensive mechanism of providing global service. Only a half-dozen satellites are needed for almost complete Earth coverage (barring some polar areas).

An eavesdropping attacker is greatly aided by these coverage characteristics as emissions from GEO satellites are not targeted towards specific users. As a




Figure 6.1: Notional Overview of a GEO Network. A typical web request would travel from the customer to the satellite (Step 1) before being redirected down to the ISP's ground station equipment (Step 2). From there it would be routed as normal IP traffic to the internet (Step 3). This process then occurs in reverse traveling back from the internet to the ISP's ground station (Step A), up to GEO orbit (Step B), and down to the customer's dish (Step C). It's worth noting that the forward link signal (Step C) is typically sent on a wide beam, with footprints measuring on the order of millions of square kilometers.

result, the radio waves reaching an attacker's antenna could be carrying traffic intended for an entire continent of satellite customers. Since these are consumer-oriented networks, the equipment necessary for eavesdropping on such signals is inexpensive and widely available, as shown in Chapters 4 and 5. With the rise of software defined radios, even more complex protocols are within the reach of relatively unsophisticated attackers [83].

In light of this threat, it is not intuitively clear why status quo satellite broadband services fail to encrypt customer traffic. The main barrier is physical. Speed of light delays over the 30,000 km hop to GEO are substantial and round-trip latency can exceed 600 ms. Latency can be reduced with closer satellites in low Earth orbit (LEO) but this increases costs and complexity. While LEO offers as little as 50 ms in speed-of-light latency, satellites only maintain line of sight for a matter of minutes. Consistent global coverage thus requires hundreds of

satellites. Status quo LEO constellations can still experience round-trip delays of up to 1,500 ms depending on the route a message travels [191]. Thus, while in-development constellations have made ambitious claims, satellite latency will likely remain relevant for some time [192, 193].

6.2.2 TCP Performance Over Satellite

To understand how latency discourages encryption, one must consider its impact on TCP performance. In this chapter, we focus on standard TCP implementations on the assumption that forcing satellite customers to use alternatives (e.g., TCP-Hybla) is infeasible [194]. We outline two of the most prominent issues here but many others have been extensively characterized in prior work [195–198].

Barriers to TCP in Satellite Networks

The first challenge to TCP performance in satellite networks arises from the requirement that TCP data packets are responded to with an acknowledgment (ACK) message [198]. The effect is compounded by the three-way handshake which, in the best case, takes upwards of 1,500 ms to complete over GEO. When visiting a website with embedded images and related files, many three-way handshakes may be required - compounding delays. Although modern implementations may employ various optimizations to bundle or reduce the total number of ACKs, these are not tailored for satellite networks [195]. Further, in some legacy devices, ACKs may be elicited by every packet, greatly increasing perceived latency [199, 200].

The second challenge arises from TCP congestion control and TCP “slow-start” initialization [195]. TCP slow-start gradually increases the ratio of data segments to ACKs until a desired congestion window is reached. The time this process takes is thus a function of round-trip times (RTT) over the satellite link. Even once a connection has reached optimal window size, packet loss can be misidentified as a sign of congestion and cause the slow-start sequence to restart. While modern satellites are more reliable than in the past, packet loss is still common compared

to terrestrial networks. As a result, TCP sessions are both slow to maximize their bandwidth usage, and, once maximized, struggle to maintain that state.

These are but two factors among dozens, ranging from specific TCP option implementations to congestion control implications of link asymmetry [198]. Satellite network designs create a uniquely hostile environment for TCP.

PEPs

The most common approach to optimizing TCP traffic over satellite environments is the use of a class of appliances called “Performance Enhancing Proxies” or PEPs and loosely described in IETF RFC 3135 [201]. PEPs differ substantially and many implementations are proprietary and inaccessible to researchers. However, IETF RFC 3153 outlines a few basic principles that apply to most PEPs.




Figure 6.2: Split Distributed PEP Handshake Example

There are two typical PEP deployment options: integrated or distributed [202]. In integrated PEPs, a PEP appliance operates on a single endpoint - typically the ISP satellite gateway between the satellite network and the internet. In distributed

PEPs, a PEP appliance operates on multiple endpoints — typically the customer satellite modem and the ISP gateway.

In either deployment, the PEP intercepts TCP traffic and applies optimizations in order to compensate for satellite performance issues. Typically, PEPs do this in a manner which is invisible to conversation endpoints so that no modifications are required on consumer hardware. This is referred to as a “transparent” PEP [201]. However, the concept of transparency is somewhat misleading as, in many cases, PEP modifications are still detectable (e.g., altered TCP sequence numbers).

Beyond this, PEPs vary quite broadly. Modifications made to TCP packets are often proprietary and implementation-specific. One common approach is to “split” incoming TCP connections prior to transmission across the satellite link and issue local ACK messages immediately for received TCP packets [201]. This allows three-way handshakes and congestion control to be negotiated locally before the satellite hop but requires the PEP developer to handle errors across the split.

In distributed PEPs, this splitting approach is extended to create a tunnel between the individual PEP installations (see Figure 6.2). A TCP packet arriving at the client-side PEP (e.g., on the home satellite modem), is terminated locally as a TCP connection, and the payload is then forwarded through GEO using a modified TCP protocol (e.g., TCP-Hybla) or an alternative [194]. At the ISP gateway, a second PEP receives this modified packet, converts it back to normal traffic, and sends it along a locally-managed TCP connection to the internet.

Other PEP strategies can range from modifying TCP congestion control to bundling related packets into single transmissions. Commercial implementations often offer higher-level features such as inspecting HTTP payloads and combining requests for web-pages with their associated content [203]. A substantial body of existing work on PEPs covers these optimizations in detail not only for satellites, but also other latency sensitive environments (e.g., cellular networks) [195–197].

Security Consequences

PEPs have become a vital component of satellite broadband, and customers have come to expect the performance characteristics of PEP-accelerated networks. This has created unintended tension between broadband performance and security.

As noted in RFC 3135, PEPs break the end-to-end semantics of IP connections [201]. Specifically, they require that the PEP appliance transparently modify packets — essentially acting as a benevolent man-in-the-middle on all TCP connections. This creates inherent compatibility issues with most commonly used VPNs as the PEP is unable to “snoop” into the VPN traffic flow and identify ACK messages. Even VPNs which leverage TCP for the transport layer are not correctly accelerated, as ACK messages within the encapsulated connection are indistinguishable from other traffic. While most VPNs will function over PEPs, and in the case of TCP VPNs, observe modest improvements in initializing VPN sessions, functional browsing performance is roughly the same as if no PEP was deployed. The type of VPN employed may have marginal performance effects (e.g., UDP tunnels may receive better prioritization from the satellite ISP), but from a PEP-compatibility perspective, any VPN which does not leak the full TCP headers of a customer’s connection faces the same issues. As a result, end consumers are faced with a choice between the security of VPNs and the performance of PEPs.

6.2.3 Existing Security Approaches

In this section, we will briefly consider some of the more consequential approaches proposed in academia and industry to enable satellite broadband encryption at each layer of the TCP/IP protocol stack. This analysis better characterizes how, despite a long history of research, PEP-compatible security remains unsolved in practice.

Physical and Link-Layer Approaches

Many techniques for over-the-air encryption focus on the lower layers of the networking stack — before TCP/IP becomes relevant. For example, physical-layer techniques such as frequency hopping patterns derived from cryptographic




Figure 6.3: Simplified Overview of QPEP Architecture. A traditional “splitting” TCP PEP is combined with a tailored over-the-air QUIC tunnel which offers encryption and further performance benefits compared to TCP. This architecture can be employed by either individual customers in the form of software on their computers and enterprise routers or by ISPs in the form of ground station and modem software.

keys or direct sequence spread spectrum (DSSS) have been suggested as a mechanism for securing the entire satellite link [204]. Likewise, the injection of artificial noise as an alternative to key-based encryption has been proposed [205]. These schemes tend to focus on military systems as they often necessitate expensive modifications to hardware that would be commercially unpalatable.

At the link layer, proposals still incur hardware costs but costs are often more manageable and restricted to hardware-based decapsulation. For example, the Consultative Committee for Space Data Systems (CCSDS) has proposed Space Data Link Security (SDLS), a protocol with built-in encryption for telemetry commands to scientific space missions [116]. Likewise, the proprietary Common Scrambling Algorithm (CSA) has long been used to restrict broadcast access to paying satellite television subscribers using smart-cards, albeit with notable security weaknesses [102].

One challenge for link and physical layer encryption systems like these is the multi-user environment. As it is rarely economically feasible to allocate each customer a unique satellite channel, customers with the same ISP will generally have a key which allows them to receive traffic from other broadcast subscribers. In such systems, the customer modem will determine which packets are relevant on the basis of header information and drop other traffic from the multiplex streams.

An additional challenge with these systems is the process of key distribution and revocation over the broadcast medium.

Network and Transport-Layer Approaches

To provide over-the-air encryption with per-customer keys, a number of network-layer techniques have been proposed. In contrast with lower level approaches, interactions with TCP PEPs must now be considered directly. Many replicate traditional VPN software with bespoke modifications — for example by creating a modified IPsec with special encapsulating headers visible to PEP appliances [112, 206]. Proprietary “satellite VPNs” also exist which, while public information on their design is limited, are likely similar in design [207, 208]. Beyond concerns arising from proprietary encryption schemes, these non-standard layers can increase operator costs by limiting compatibility with existing networking equipment.

It makes intuitive sense to incorporate encryption within PEP appliances themselves — straddling the network and transport layers. This may be achieved by, for example, implementing an encrypted protocol over the satellite hop in a distributed PEP system. The transmitting PEP would first modify the TCP packets, then encrypt them. The receiving PEP would subsequently decrypt the received packets and forward them along the internet as normal. Many real-world PEP encryption products appear to employ this approach [209–211].

Most, if not all, encrypted PEPs are proprietary and not sold direct-to-consumer, making security claims difficult to verify. However, purported leaked manufacturer documents allude to built-in law-enforcement/intelligence back-doors in prominent examples [212]. Our own analysis of one satellite router with a pre-installed proprietary PEP found numerous cryptographic shortcomings, such as Diffie-Hellman implementations which are susceptible to man-in-the-middle attacks and key/IV reuse that permits replay attacks. Similar vulnerabilities have been alluded to in prior research [213]. Generally, the costs of adopting an encrypted PEP are undertaken by ISPs who may not perceive such purchases as value-for-money.

Application-Layer Approaches

An alternative approach would be the use of protocols which operate over the PEP-accelerated TCP connection. The widespread use of TLS encryption for websites, for example, has the effect of encrypting customer data over-the-air. However, this still leaks potentially sensitive data (such as the IPs a customer visits) over the massive radio-eavesdropping footprint of a satellite signal. Moreover, real-world observations of modern satellite traffic (see 6.2.1) have found that, while customers could, in theory, use TLS, many do not. While this decision is the customer's, satellite ISPs may nevertheless have a duty of care.

Another application-layer approach is tunneling traffic into an encrypted TCP stream and issuing local ACK messages before the data egresses from the client's computer. This differs from most SSL-VPNs which do not spoof the connection endpoint. Some commercial products appear to implement this approach [214]. However, the requirement of software installed on the client's computer limits compatibility with embedded devices and creates friction.

6.3 The QPEP System

To address some of these shortcomings we have developed QPEP, an open-source and non-proprietary tool which can be used by both individuals and ISPs to both encrypt and accelerate satellite TCP traffic¹. At its core, QPEP follows a distributed “snooping” PEP model similar to the methods described in Section 6.2.3. The QPEP client tunnels TCP traffic over the satellite link inside a stream that leverages the encrypted QUIC transport protocol. Tunneled traffic is decapsulated by a receiving QPEP server which then routes the decapsulated traffic over the internet as if it were the client. A high level overview of this architecture appears in Figure 6.3.

¹Source code and documentation for both our QPEP implementation and our OpenSAND-based testbed environment are available publicly (<https://github.com/ssloxford/qpep>). Example python scripts used to run all of the simulation scenarios presented in this research are provided.

6.3.1 QPEP Design Contributions

QPEP's principal objective is to enable customers of satellite internet services to defend against forward link eavesdropping attacks without suffering performance reductions. In contrast with prior work on this topic, QPEP expands its security model to consider both wireless eavesdroppers and the ISP as potential attackers. As a result, QPEP cannot take advantage of ISP PEPs and TCP optimizations but must instead integrate these performance optimizations within the context of its security design.

Architecturally, QPEP is most similar to a tunneling VPN appliance. Is it designed to offer an end-to-end encrypted link between a customer device on an untrusted satellite network (e.g. a computer aboard a maritime vessel) and an egress point into a trusted network (e.g. a corporate LAN environment). Unlike a traditional VPN, which typically tunnel traffic without modification, QPEP allows customers to apply PEP-style optimizations to tunneled TCP connections prior to transmission across the untrusted satellite link.

This hybrid-approach emerges from the recognition of interactions between two seemingly unrelated design factors of satellite networks which have not been considered in prior work. The first is the recognition that PEP applications are traditionally the prerogative of satellite ISPs for cultural and commercial, rather than technical reasons. There is no underlying reason that satellite customers could not choose to “bring their own PEP.” Rather, doing so has simply been redundant with the service satellite ISPs already provide. The second is the recognition that the TCP-tampering phase of a PEP's operation, which already breaks the end-to-end semantics of TCP connections, represents a minimally intrusive opportunity for transparently modifying the security properties of these connections. Combined, these two factors give rise to QPEP's unique integration of a transparent distributed PEP architecture with that of a transparent tunneling VPN appliance.

The fundamental security design of QPEP is deliberately straightforward. We use QUIC — a proven and popular standard — instead of designing our own scheme for authentication, key exchange, and session management over the satellite link.

Further detail on the specific implementation and implications of this security model can be found in Section 6.3.2. As mentioned in Section 6.1, our goal is to adequately resolve an urgent issue affecting real-world satellite networks. To the extent that existing protocols can be adapted to meet this need, arbitrarily complex variations serve little purpose beyond academic diversion.

That said, QPEP is far from the trivial combination of two pre-existing technologies. QUIC, as a transport layer alternative to TCP, is designed for traditional server-client conversations which carry data from individual applications. For example, QUIC might be used as an alternative to TCP when delivering the contents of website to a compatible web-browser. This requires both parties of the conversation to explicitly support QUIC in order to benefit from the performance and security enhancements offered by the protocol.

QPEP, on the other hand, is designed to transparently upgrade insecure and non-performant TCP connections to QUIC traffic. In doing so, QPEP allows users to reap the security and performance benefits of the QUIC protocol vis-a-vis the satellite link in an application-agnostic manner. In typical usage, we would expect *neither* the client or server of a given connection to have any special support for the QUIC protocol.

While limited prior work has considered QUIC performance in the context of individual applications served over a satellite connection [215], to our knowledge, QPEP is the first application of QUIC as a general purpose encapsulation layer for SATCOMs. This raises a number of technical challenges discussed further in Section 6.4. In particular, multiplexing unrelated connections within a QUIC session and ensuring that errors in individual streams do not impair overall performance requires the QPEP application itself to correctly monitor the state of individual encapsulated connections and propagate errors appropriately across the satellite link. Likewise, modifications must be made to both the client and server implementation of QUIC to enable longer-lived streams with sporadic traffic patterns — otherwise repeated initialization of the QUIC tunnel could quickly undermine performance gains.

Table 6.1: QPEP Comparison to Status-Quo PEP and Security Options.

	Accelerated TCP	Private From: Eavesdropper	Private From: ISP	Deployed By: Customer	Deployed By: ISP	Open Source Available?
Plain Connection	No	No	No	Yes (Default)	Yes (Default)	Yes (Default)
Traditional PEP (e.g. [216])	Yes	No	No	No	Yes	Yes (Rare)
Traditional VPN (e.g. [217])	No	Yes	Yes	Yes	No	Yes
Secure PEP (e.g. [187])	Yes	Partially	No	No	Yes	No
QPEP (this chapter)	Yes	Yes	Yes	Yes	Yes	Yes

In essence, QUIC is a web protocol, while QPEP is a novel application which leverages an optimized version of that protocol as one piece of its design. Notionally, QPEP’s split-tunnel architecture does not necessitate the use of QUIC and could be incorporate a variety of alternate tunneling protocols. The architecture contributed here would represent an intuitive starting point for doing so, and our QUIC-based implementation would offer a replicable benchmark to beat in such future work.

Ultimately, QPEP offers several direct and substantive improvements over prior work, as outlined in Table 6.1. QPEP leverages a novel tunneling strategy, an expanded security model which includes protection against ISP snooping, and, as show in Section 6.6.2 promising performance characteristics. Moreover, QPEP is, to our knowledge, the first open-source and verifiable satellite encryption software implementation for communications of this nature. To the limited extent that prior work on secure and performant PEPs exists, direct comparisons are impossible absent access to proprietary software belonging to ISPs or unimplemented academic proposals. Given this, our hope is that QPEP offers a useful foundation for replicable academic work in addition to serving as a direct technical contribution to the challenge of GEO broadband security.

In this chapter, we focus on the design and evaluation of a specific implementation of QPEP’s architecture. However, the core system architecture is a contribution which is agnostic to the encryption schemes and transport protocols employed. Indeed, the open-source code for both QPEP and its simulation testbed is designed to facilitate the swapping out of individual engineering components by researchers.

6.3.2 Use of the QUIC Protocol

The principal strategy employed by QPEP to support over-the-air security is the use of the QUIC protocol for tunneling traffic over the satellite link. While the QUIC standard is still evolving, it has already seen wide adoption in terrestrial networks due to its performance and security advantages over TCP. Several of these benefits make QUIC intuitively promising for secure PEPs.

QUIC is a proposed transport protocol which is presented as an alternative to TCP for modern web services [218]. One of the main advantages of QUIC is that it directly integrates a TLS 1.3 style handshake into its session initialization process. This contrasts with encryption schemes where encryption is an application-layer feature. In such instances, the establishment of a secure channel requires a separate handshake to occur after the TCP three-way handshake. By comparison, QUIC effectively combines the creation of a client-server connection and an encrypted channel into a single process. This is appealing as round-trips over satellites are costly.




Figure 6.4: Simplified Comparison of QUIC and VPN Initialization.

A detailed characterization of QUIC’s handshake protocol can be found in the IETF draft standard [218]. A simplified overview comparing the process

required to establish a secure link over a QUIC tunnel relative to a traditional VPN connection also appears in Figure 6.4. In typical usage, the QUIC handshake process allows for the creation of an encrypted channel from a QPEP client to QPEP server in a single round trip.

Our decision to incorporate QUIC contrasts significantly with other tunneling based PEPs. In the status quo, the dominant approach, used by commercial PEPs such as Tellitec’s “Enhanced TCP” (ETCP) product, is to implement a bespoke network/transport-layer protocol [209]. This limits compatibility with other network infrastructure (e.g., firewalls, switches, QoS appliances) and may thus require direct ISP participation and investment to deploy. QUIC, on the other hand, is notionally a transport-layer protocol but can also be implemented as an application on top of UDP. This allows for out-of-the-box compatibility with existing networks. An additional benefit to QUIC as opposed to, for example, TLS 1.3 secure channel tunnels, is that ISP PEPs largely ignore UDP traffic, limiting the risk of unexpected interaction with existing infrastructure [219].

While a TCP-over-QUIC-over-UDP architecture may initially seem rather contrived, it effectively allows us to develop a secure PEP which does not require ISPs to operate any decapsulation software on their gateway to ensure traffic compatibility with the wider internet. The QUIC layer provides transparent upgrades to the security and performance of the encapsulated TCP traffic and the UDP layer enables QPEP to operate transparently to ISPs. The end result is that individual customers can use QPEP to secure their traffic without requiring ISP involvement.

This ability for customers to protect their traffic between two arbitrary endpoints without trusting their ISP makes QPEP’s security properties most comparable to prior satellite VPN protocol proposals [112, 206]. However, QPEP’s design differs from these tools which still reveal limited portions of the TCP header to ISP PEPs for optimization (e.g., destination IP, port numbers, and TCP flags). An ISP snooping on a customer’s QPEP traffic would only see the IP address and UDP port of the customer’s upstream QPEP server. All information regarding the true TCP connections are hidden inside the QUIC tunnel. This means QPEP

can function in the presence of other ISP-installed PEPs without any special adjustments. From a service-provider perspective, it is no different from any other UDP-based application. Of course, QPEP could also be installed by ISPs on customer modems and network gateways just like with traditional PEPs, but trust in ISPs is no longer a design requirement.

QUIC & QPEP's Security Design

A full security analysis of QUIC's handshake and channel initialization process is both beyond the scope of this thesis and well-provided in prior work [218, 220, 221]. From a more pragmatic perspective, QUIC has benefited from real-world use in some of the world's most popular web applications for many years. To our knowledge, no attack against the protocol itself has been demonstrated which could give rise to security and privacy issues of severity comparable to those identified in Chapters 4 and 5. To the extent that relatively modest issues have been found and mitigated over the course of QUIC's development (e.g., as in [222]), we expect the QPEP architecture to be adaptable without issue. Indeed, this large existing research community is one of the principal advantages in adopting QUIC as our tunneling protocol rather than developing our own security layer for the satellite channel.

These caveats notwithstanding, there are some aspects relating to the interface between QUIC's security model and QPEP's architecture which do merit explicit consideration. Namely, QPEP is quite different from that of a typical QUIC-based web service. The "public-facing" ends of the QPEP system are TCP-based. The QUIC protocol itself is only used for the single two-way tunnel between a QPEP server and client middle-box. In practice, this means a QPEP server need only speak to a small handful of QUIC clients and a QPEP client need only expect to handle incoming QUIC traffic from a single server.

This has some benefits. For example, consider a typical QUIC reflection denial of service attack [223]. In this attack, an attacker spoofs a victim IP address and sends a small client hello message to many QUIC servers. Each of these servers responds to the victim's IP address with comparatively large server hello message,

potentially overwhelming the victim's link. As the QPEP server need only expect to handle QUIC client hello messages from a small number of clients, simple firewall policies could be implemented to restrict inbound client hello messages to a simple allow list. This significantly reduces the risk that QPEP servers could be abused to conduct QUIC reflection attacks against external victims.

While the QPEP server could still be tricked into generating spurious hellos to victim IPs which overlap its expected client IP space, QUIC's built-in protections against reflection attacks are likely sufficient to mitigate such a threat. Specifically, modern versions of QUIC restrict the size of server hello messages to a maximum of three times the size of a client hello message [224]. As the attacker would typically be limited to abusing only a single QPEP server to generate such traffic, the reflection attack could, at best, only triple the attacker's effective traffic. It is thus highly unlikely that the presence of QPEP offers a meaningful benefit to attackers compared to any number of simpler denial of service attack vectors (such as QUIC reflection from non-QPEP servers).

This property of QPEP's architecture makes it possible to employ more robust defenses against spoofed traffic than are practical in typical web services. Notably, the QUIC standard only requires that *clients* use TLS authentication to validate the identity of QUIC servers [224]. However, as QPEP servers need only validate a small number of broadly knowable QPEP client applications, it is feasible to mandate client-authentication against a particular certificate authority as part of the TLS handshake. In current proposals for QUIC mutual authentication, this would not mitigate the aforementioned risk of spoofed client hello messages as these take place before the authentication stage of the handshake. However, mutual authentication could detect and prevent spoofing at later stages. Moreover, a QPEP server could use mutual-authentication as a flexible alternative to IP filters for restricting use of the service to authorized users. In the SATCOMs context this could be particularly as mobile clients, such as a QPEP instance on a maritime vessel, may hop between multiple ISPs and networks as it traverses the globe. While the proof of concept

implementation in Section 6.4 does not include mutual authentication support, it represents one logical step for QPEP's evolution.

It is also worth noting that, in this thesis, our focus is on the threat model of a low-resourced passive attacker intercepting wireless transmissions from GEO. A number of proposed attacks against earlier versions of QUIC, namely related to performance degradation, required attackers to have the ability to manipulate in-flight traffic between QUIC endpoints [222]. In a wireless GEO broadband context, such manipulation would be exceedingly difficult, even with sophisticated radio equipment. An attacker would have a slightly easier time compromising the terrestrial route between an ISP ground station and a cloud-based QPEP server. However, reliably gaining access to these links would require significant effort and an attacker could probably achieve similar disruption by simply blocking traffic or degrading the connection directly. Evaluating the security of the QPEP system in the context of a high-resourced active attacker is another possible avenue for future work.

Finally, QPEP does not make use of QUIC's ability to support zero round trip (0-RTT) conversations, whereby clients may send encrypted data before completing a handshake. This is because QUIC 0-RTT requires careful consideration to ensure security against replay attacks and, as demonstrated in Chapter 5, replay attacks are particularly pernicious due to latency factors in GEO satellite broadband [224]. As discussed further in Section 6.4.3, we expect 0-RTT session initialization to offer only modest performance benefits to the QPEP architecture. That said, there is no reason, in principle, that replay protection could not be implemented as part of the QPEP protocol to facilitate 0-RTT initialization. In particular, propagating TCP sequence numbers across the QUIC link (rather than our current approach of generating new sequence numbers for each side of the tunnel) may facilitate de-duplication of replayed traffic at the TCP layer.

Ultimately, the use of QUIC in QPEP provides an intuitive mechanism to add a security layer to arbitrary TCP traffic in GEO broadband. QUIC's explicit consideration of latency in its initialization and handshake design makes it particularly attractive to QPEP's architecture. Under our threat model, we

expect QPEP's integration of QUIC to functionally eliminate the risk of low-resourced passive eavesdroppers successfully intercepting and interpreting private satellite radio transmissions.

QUIC & QPEP's Performance Characteristics

Beyond these security benefits, the use of QUIC offers notable performance advantages.

First, the initial QUIC connection can be negotiated in a single round-trip, substantially shorter than the TCP three-way handshake. When compared with alternative encrypted tunnel schemes — such as TLS-based VPNs - QUIC offers a substantial reduction in round-trip transfers (see Figure 6.4). Indeed, for previously known QUIC servers, it is possible for a client to begin transmitting data from the very first packet. While the TLS session initialization process is particularly ill-suited to satellite environments, few tunneling approaches can initialize secure channels with comparable RTT requirements to QUIC. For example, Internet Key Exchange (IKE) initialization commonly used in IPSec VPNs will generally require three or more round-trips depending on configuration and version.

Additionally, unlike TCP, QUIC does not require that all packets in a stream be processed in a particular order — removing head-of-line blocking issues and permitting heavy multiplexing. This allows QPEP to encapsulate multiple TCP flows inside a single QUIC session, reducing the number of session-initialization round-trips.

Like TCP, QUIC has built-in support for the re-transmission of lost and corrupted packets. This obviates intuitive concerns relating to UDP usage over low-reliability satellite links. Moreover, some draft proposals suggest the addition of built in forward error correction (FEC). These efforts have largely stalled due to minimal terrestrial performance gains [225]. However, satellite environments may represent a context for reviving research on QUIC-FEC.

QUIC Satellite Performance

As QUIC is a relatively recent protocol, its use in satellite environments has not been subject to much research. What does exist is largely inconclusive. Some preliminary assessments have found that QUIC facilitates a 100% increase in satellite broadband page load times compared to PEP-accelerated TCP [215]. However, others suggest that QUIC performs better, measuring up to a 50% decrease page load times [226]. Preliminary IETF discussions have lead to a number of proposed (but unimplemented) techniques for optimizing QUIC over satellite [227].

Relevant research has focused on real-world connections to HTTP2 web-servers which support the QUIC protocol. In these cases, researchers only control the client-side QUIC configuration. However, under QPEP's distributed model, it may be possible to optimize both server and client QUIC implementations for the satellite link. Much as many modern PEPs use modified TCP implementations, QPEP's architecture lends itself naturally to bespoke optimization of QUIC parameters relating to FEC, ACK decimation and congestion control.

6.4 QPEP Implementation

Reaping the theoretical benefits of QUIC as a transport alternative for satellite TCP connections raises several important engineering considerations. In this section, we focus on the specific implementation and architecture of QPEP, how it merges properties of standard VPN applications and TCP PEPs, and some of the challenges in doing so. A simplified overview of how a typical packet might be routed through a full QPEP installation to a notional web service appears in Figure 6.5.

6.4.1 System Architecture

QPEP is implemented according to a distributed PEP application architecture in order to ensure compatibility with all web services rather than only those with native QUIC support. This distributed design allows QPEP to operate transparently, converting TCP conversations into QUIC streams over the satellite hop and then




Figure 6.5: A simplified overview of typical packet flow through the QPEP system. Dotted arrows represent wireless SATCOMs traffic. Note that the QUIC initialization handshake only occurs for the first packet following a five-minute inactivity timeout.

back into TCP conversations terrestrially. The benefit is that no special software or configuration is required on individual customer devices. QPEP-encrypted traffic appears identical to normal TCP traffic. This differentiates it from application-layer commercial PEPs and ensures compatibility with IOT systems.

The practical implication of this architecture is that a QPEP deployment consists of two independent appliances (see Figure 6.3): a QPEP client on the customer side of the satellite link and a QPEP server on the internet side.

The client application can be installed as software directly on a customer's device. However, it is also designed to operate transparently if placed along the network path between a customer's device and the satellite modem. For example, an enterprise user or satellite ISP might install a QPEP client on a router within a local area network in order to encrypt and optimize internet-bound traffic from all connected devices.

The server application is similarly flexible. It can be installed by the satellite ISP on their gateway, like a traditional PEP. However, it can also be installed anywhere else on the internet, with encrypted QPEP traffic traversing the ISP gateway en-route to a cloud server or other egress point, as with a traditional VPN.

When the QPEP client launches, it opens a QUIC tunnel with an upstream QPEP server. Unlike in normal QUIC services, where idle sessions are short-lived, QPEP sets the timeout for this tunnel to a relatively long period of time (5 minutes of link inactivity). QPEP does this because QUIC session initialization requires a full round-trip over the satellite link, so by re-using recently established QUIC tunnels, QPEP can save round-trips over creating a new QUIC tunnel for each connection.

Each TCP connection which is managed by QPEP is assigned to its own unique QUIC stream within this QUIC tunnel. This allows QPEP to multiplex concurrent TCP connections and avoid creating redundant session initialization handshakes over the satellite link. As discussed Section 6.4.2, this also allows for better congestion control as losses in each stream can be handled independently.

QPEP does not naively convert every incoming TCP packet to a QUIC packet. If it did so, we would expect performance akin to that of a traditional VPN. This is because the TCP three-way handshake would still occur over the latent satellite

hop. Instead, QPEP must selectively terminate incoming TCP connections, drop spurious acknowledgments, and send only meaningful data across the satellite. This requires both the QPEP server and QPEP client to internally maintain state regarding each TCP connection.

When the QPEP client receives a TCP SYN packet, it immediately initiates a three-way handshake across the customer LAN — effectively “spoofing” the upstream destination TCP server. Upon finishing this handshake and receiving a TCP packet with payload data, it opens a new stream inside the QUIC tunnel session it established with the QPEP server at initialization. The client then strips away the TCP header information and encapsulates the payload data into a QUIC packet. A simple “QPEP header” consisting of a TCP four-tuple ($\langle src\_ip, src\_port, dst\_ip, dst\_port \rangle$) is prepended to this packet. The client maintains a local state dictionary which maps the QUIC stream identifier, this “QPEP header,” and the associated TCP socket.

When the QPEP server receives an incoming QUIC payload, it checks its own state dictionary for any sessions associated with the received QPEP packet header. If no such entry is present, it opens a fresh TCP connection to the upstream TCP server on the basis of the received QPEP header and completes a three-way handshake across the internet — effectively “spoofing” the customer’s device. It then updates its state dictionary to map this TCP session with the appropriate QPEP header and QUIC stream. From then on, each packet which the server receives in this QUIC stream will be converted into a TCP payload and then transmitted across the associated TCP socket to its destination. This same process happens in reverse for each response which comes from the internet, with the client extracting payloads sent by the server across the QUIC stream and then routing them to the appropriate TCP socket and onwards to the customer’s device.

6.4.2 Error Handling and Session Management

The main challenge with this protocol splitting approach is correctly propagating errors which occur over one of the three network segments ($Customer \leftrightarrow QPEP (client)$;

$QPEP (client) \leftrightarrow QPEP (server)$; $QPEP (server) \leftrightarrow Internet$) to the others.

Over the satellite link, session management and congestion control is implemented within the QUIC session. This involves a modified version of the popular CUBIC congestion control algorithm for responding to losses which occur over the satellite hop. Congestion control is applied on a per-stream basis, preventing any individually troublesome stream from impacting the performance of other streams in the $QPEP (client) \leftrightarrow QPEP (server)$ session. While our implementation uses CUBIC for this purpose, only modest engineering effort would be required to replace it with any other QUIC-compatible algorithm. Adopting existing delay-tolerant algorithms (such as TCP-Hybla) to the QPEP architecture may thus represent one avenue for future research.

Over the terrestrial links, this process is handed down to the host's TCP stack. QPEP's "spoofed" endpoints resolve connection issues terrestrially just as any other TCP application. By design, this makes QPEP functionally transparent to users and ensures compatibility with upstream network appliances, such as firewalls or traditional VPN software.

The more difficult case is for errors which occur in one link and have implications for the others. For example, if a TCP connection on the $QPEP (server) \leftrightarrow Internet$ link fails, the error state of that TCP socket must be propagated up from the host's TCP/IP stack to the QPEP server application. The QPEP server will then issue a message to the QPEP client across the $QPEP (client) \leftrightarrow QPEP (server)$ segment designating the associated QUIC stream for closure. Upon receiving this message, the QPEP client will remove the stream from its session mapping dictionary and terminate the appropriate TCP connection on the $Customer \leftrightarrow QPEP (client)$ network. Finally, both the server and client will close the corresponding QUIC stream.

6.4.3 Limitations

In the implementation evaluated in this chapter, QPEP only modifies TCP/IP connections. This is because our objective is to evaluate a secure alternative to traditionally unencrypted PEP appliances, which also focus exclusively on TCP

connections. Only minor engineering modifications would be required to tunnel other protocols into the QUIC stream, such as UDP and ICMP. However, it is worth noting that we expect QPEP to have only marginal performance impacts on such protocols as they do not incur the same latency penalties as TCP over the satellite hop. Nevertheless, doing so may be desirable for end-users as it would bring over-the-air encryption by default to DNS queries and other non-TCP traffic.

It is also worth noting that we have not implemented QUIC's optional zero-round trip (0-RTT) session initialization handshake. While being able to further reduce the number of costly satellite round trips is an attractive prospect, prior work on QUIC's 0-RTT raises some security concerns with respect to replay attacks [228]. The potential harms of replay attacks are especially acute in the wide-footprint and high-latency context of satellite broadband. Indeed, a satellite eavesdropper may have closer physical proximity to a given QPEP server than the satellite ISP's gateway, which could allow them to even deliver "replay" messages faster than legitimate ones. Given that QPEP relies on long-lived QUIC sessions, the benefits of 0-RTT are likely marginal at best. This is because a QUIC handshake is only required for the very first connection a QPEP client makes, with subsequent streams re-using that session. Nevertheless, consideration of 0-RTT initialization dynamics may offer a route for some further optimization in future work.

6.4.4 Availability

An open-source reference implementation of QPEP, written in Go, is available in conjunction with this thesis. Go was selected to increase accessibility without substantial performance sacrifices. To the best of our knowledge, only two non-proprietary PEPs exist [216, 229]. Both are implemented in C/C++, lack encryption capabilities, and have received only minimal development attention over the past several years. Other notable academic PEPs are either not publicly available or restricted to particular simulation tools [112, 230].

The QUIC implementation used by QPEP is based on the widely used quic-go library which roughly tracks the IETF QUIC proposal [231]. As discussed in

Section 6.6.5, minor optional modifications to the QUIC implementation can be made to optimize performance in the satellite networking environment. Future work might also consider the suitability of other QUIC implementations such as Chromium's [232].

6.5 Secure PEP Testbed

One challenge in developing and evaluating PEPs has been creating replicable simulations of system performance. While systems can be tested on live satellite networks, understanding performance under adverse conditions (e.g., poor reception quality or network congestion) or creating experiments which others can verify often requires some degree of simulation. Simulating satellite IP networks involves more than the simple injection of artificial latency, which can result in misleading and inaccurate results [233].

The OpenSAND engine, previously Platine, is a long-standing satellite network simulation environment for more faithfully replicating satellite broadband [188]. The engine supports built in attenuation and modulation emulation, replicating conditions which can have significant implications for TCP performance. OpenSAND emulates satellite networks down to the link layer, simulating low-level protocol noise mitigations and creating realistic traffic routing behaviors.

However, the OpenSAND environment is somewhat difficult to configure — requiring multiple devices and precise network conditions. This has been noted in prior work as a barrier to its use, despite its relatively high degree of accuracy when validated against real-world networks [233].

In the process of assessing QPEP's performance, we have developed a simple dockerized deployment of the OpenSAND engine specifically tailored towards replicable PEP benchmarking. Our testbed models a basic GEO satellite network consisting of a single gateway and satellite terminal (akin to the networks shown in Figure 6.1 and Figure 6.3). This testbed is open-source and publicly available in the QPEP source repository (see Footnote 1). Its intention is to simplify the

process for future researchers interested in making related contributions towards secure PEP development.

The testbed's gateway container is linked through the simulation's host machine to the broader internet, allowing a testbed user to open a web-browser and visit real websites as if they were using the simulated satellite link. We also connect the gateway container to a simulated LAN environment with a workstation containing several network benchmarking tools. A similar LAN environment is connected to the satellite terminal, replicating a satellite customer's devices.

On the satellite container, we include packet capture tools for real-time monitoring of simulated over-the-air transmissions. This allows for immediate verification that secure PEPs are not leaking sensitive data in clear-text.

Finally, pre-configured installations of QPEP, OpenVPN, and PEPsal are installed for both the gateway and satellite terminal networks. A set of example python scripts are provided to orchestrate the environment and run the experiments presented in this chapter. These scripts are designed as modular benchmarks which can also be adapted to future secure PEP proposals to facilitate direct and replicable comparisons with QPEP in future work.

6.6 Evaluating QPEP

In this section, we present an evaluation of the QPEP approach and its impact on the performance of TCP-based traffic within our testbed environment.²

No comparable encrypted satellite PEP is publicly available. As such, we selected PEPsal, one of the only open-source unencrypted PEPs, and OpenVPN, a popular VPN product without specific satellite optimizations, to provide some context to measurements made [216, 217]. Future work including commercial and proprietary PEPs might be of merit, although these are not readily available to researchers. It is worth noting that the particular VPN product (e.g., OpenVPN vs. PPTP vs.

²In this chapter, QPEP is evaluated only through simulation. While this has benefits for reproducibility, we originally intended to supplement these experiments with validation in a real-world VSAT network. Unfortunately, due to restrictions during the global coronavirus pandemic, this has not been possible. When real-world benchmarks are available, they will be added to QPEP's public source code repository (see Footnote 1).

IPSec) is unlikely to have meaningful impact on performance benchmarks inasmuch as all hide the true TCP headers of the customer's connection from ISP PEPs.

6.6.1 Experimental Setup

First, we consider preliminary results under ideal OpenSAND network conditions. Next, we present QPEP with various adverse network situations, assessing its performance in the presence of high rates of packet loss and under variable delay conditions such as those in LEO constellations. Finally, we briefly consider how performance modifications to the QUIC protocol itself may impact QPEP's behavior.

Unless otherwise noted, the OpenSAND network is configured to use the DVB-S2 protocol with GSE encapsulation for forward-link communications and DVB-RCS2 with RLE encapsulation for the return link. The clear-sky SNR is set to 20 dB and Adaptive Coding and Modulation (ACM) is used at the physical layer to provide quasi error free (QEF) communications at this SNR level. A constant speed-of-light delay of 125 ms is used from both the satellite terminal and the satellite gateway to the satellite (resulting in a 500 ms RTT). The forward-link carrier frequency is allocated 50.0 MHz of bandwidth with a roll-off factor of 0.25 and the return-link is allocated approximately 7.4 MHz of bandwidth. These bandwidth values are well within the simulation capabilities of the machines used to run the scenarios, reducing the risk of artificial network caps due to hardware limitations. While simulations were run on multiple hosts for efficiency reasons, comparisons made within a given experiment (e.g., all measurements shown in a single figure) were conducted on the same physical host.

Our configuration is intended to represent the characteristics of a typical GEO satellite broadband network. We also briefly touch on an alternative LEO network configuration in Section 6.6.3, which demonstrates the performance of QPEP in a situation where latency is variable, depending on the geographic location of the end user and the corresponding satellites and ground stations. The testbed supports arbitrary delay and bandwidth models which may be useful in future

work considering more esoteric network designs, such as those involving space-to-space routing or polar orbits.

Simulations of QPEP are configured with a QPEP server sitting local to the satellite gateway network and listening for incoming QUIC tunnel connections. The QPEP client is hosted on the satellite terminal and listens transparently into all incoming TCP connections. The QPEP server is configured to accept up to 40,000 concurrent streams from a single host — substantially higher than quic-go's default of 100. This is to enable compatibility with concurrent download benchmarks.

OpenVPN simulations are deployed similarly to QPEP, with an OpenVPN client connected to the satellite terminal and an OpenVPN server connected to the satellite gateway. OpenVPN is configured to leverage a UDP tunnel as this is expected to perform better in the satellite environment.

PEPsal is evaluated under two different configurations — a distributed installation and an integrated installation. Evaluations of distributed PEPsal are implemented with a PEPsal endpoint transparently listening to all incoming TCP traffic on both the satellite gateway and the satellite terminal. In integrated PEPsal, a PEPsal endpoint listens to incoming TCP traffic on the satellite terminal but no endpoint is installed on the satellite gateway.

Diagrams summarizing these configurations can be found in Appendix C.

6.6.2 Baseline Performance

An initial comparative assessment of goodput can be made through the use of Iperf which attempts to provide consistent performance evaluations of network speed. For these benchmarks, an Iperf server is hosted on the satellite gateway network and is used to transfer data to an Iperf client connected to the satellite terminal. For each tool, one-hundred iterations of Iperf are run at data transfer sizes ranging from 0.5 to 10 MB in 250 KB intervals. Varying the volume of data transferred provides insights into the extent to which results are influenced by session initialization time. We would expect smaller transfers to demonstrate larger susceptibility to latent TCP handshakes as a proportion of total transfer time while larger transfers should




Figure 6.6: Goodput Comparison by Iperf Transfer Size. The shaded zones represent standard deviation across 100 simulation runs at each file size. Note that QPEP performs well for small transfers and matches the performance of the unencrypted distributed PEP for larger transfers. Meanwhile, the traditional VPN, integrated PEP, and unencrypted satellite connection all perform relatively poorly throughout.

be more heavily influenced by congestion control and total available bandwidth. The results of these experiments are summarized in Figure 6.6.

We see that QPEP is capable of making significantly greater use of bandwidth for small to moderate-sized downloads than any of the evaluated alternatives, even, surprisingly, the unencrypted PEPs. This makes sense as QPEP is able to send data along with the stream initialization packets, allowing very small transfers to be completed in a single round-trip. As shown in Figure 6.6 this has a large effect on the measured goodput for small transfers, but diminishes at larger transfer sizes until QPEP approaches the performance of unencrypted distributed PEPs.

Integrated PEPsal offers little advantage here as it is constrained by head-of-line blocking over the satellite hop and the majority of download traffic originates on the un-optimized route from the gateway to the user. Distributed PEPsal performs much better as it is able to optimize both directions of the satellite

conversation. However, it lacks QPEP's ability to encapsulate concurrent streams and to make use of the first few handshake packets for data delivery. Finally, as expected, OpenVPN performs much worse than QPEP, essentially matching, or slightly underperforming, an un-optimized satellite link.

This benchmark, while meaningful, is somewhat misleading. Iperf provides one important measure of goodput but the scenario it evaluates is not representative of real-world behavior. Specifically, opening a connection to a port, ramping it up to maximum speed, and then maintaining that speed for many file transfers is not how most web services operate. PEPs were explicitly invented to optimize web-browsing and visits to text and image-based services. Even if QPEP were well suited to encrypting certain types of file transfers, its adoption would likely hinge on its performance for web-browsing tasks.




Figure 6.7: ECDF Comparison of PLTs over Alexa Top 20. Note that QPEP shows significantly faster PLTs than traditional VPNs and marginally better PLTs compared to unencrypted PEPs. Each line represents 2,000 simulations with a connection timeout set to three minutes.

A more realistic sense can be found through the evaluation of the time it takes

to visit actual websites. Unlike Iperf, web-browsing consists of the transfer of many small files (e.g., embedded images or style-sheets) over multiple TCP sessions. Often, these files can be hosted on a variety of servers. This makes web traffic more sensitive to latency effects.

Experimentally measuring page load times (PLTs) is an imprecise art. For our simulations, we used the open-source tool Browsertime [234]. Browsertime reports PLT as the number of elapsed milliseconds between the “navigationStart” and the “load” event of the browser’s navigation timing API (as defined in [235]). This roughly translates to the amount of time between a user hitting the enter key in the browser’s navigation bar and the moment when all page resources, including the DOM, images, and stylesheets, are loaded.

To conduct these experiments, we connected our simulated satellite gateway to a real terrestrial broadband network. This naturally induces measurement variability depending on network conditions at measurement time. To reduce this variability, we conducted 100 connections with each tool to each of the top 20 distinct domains listed by Alexa Internet Inc [236]. Between each visit, the browser (a headless version of Firefox) and the DNS cache were reset. Any page loads which took more than 3 minutes were terminated as timeouts. The results of these PLT measurements are summarized by means of an Empirical Cumulative Distribution Function (ECDF) in Figure 6.7.

This page load time comparison shows that QPEP is able to encrypt realistic web browsing traffic without undermining the performance users have come to expect from status quo unencrypted PEPs. QPEP’s median page load time (PLT) across the Alexa Top 20 is 13.77 seconds. This is roughly 54% faster than distributed PEPsal’s 30.16 seconds and integrated PEPsal’s 30.5 seconds. It makes sense that both integrated and distributed PEPsal perform similarly here as PLTs are dominated by large numbers of client-initiated TCP handshakes for various web resources. In terms of mean PLTs, which are more heavily influenced by “worst-case” long-running connections, QPEP still significantly outperforms the traditional insecure PEP,

with a mean PLT of approximately 25.80 seconds compared to distributed PEPsal's 37.61 second average and integrated PEPsal's 40.70 second average.

The most important benchmark comparison, however, is between QPEP and the status quo options for end-to-end web traffic encryption. In this case, we find that QPEP more than halves median PLTs when compared to OpenVPN's encryption, achieving 72% faster page loads than an OpenVPN-encapsulated connection's 49.42 second median PLT. In terms of mean PLTs, QPEP still roughly halves OpenVPN's mean PLT of 50.01 seconds. As expected, we further observe that OpenVPN roughly matches, or slightly under-performs, a basic unencrypted and unoptimized satellite link.

The relative disadvantage of using a traditional VPN for over-the-air encryption in GEO broadband is clear when considering this PLT metric. QPEP is functionally the same from a security perspective (eavesdroppers cannot interpret intercepted traffic), but significantly more performant by design. The surprising additional outcome that QPEP achieves significantly lower PLTs than established and architecturally similar insecure PEP appliances suggests that QUIC is particularly well-suited for the satellite tunneling use-case.

6.6.3 Performance Under Adverse Conditions

While these basic evaluations present a compelling case for the use of QPEP in a typical GEO environment, satellite networks can exhibit many atypical characteristics. Packet loss, rain-fade, and orbit altitudes can all have significant performance implications. As such, we have elected to evaluate the relative performance of QPEP under some of these conditions.

Intuitively, packet loss and rain fade conditions are significant threats to encrypted tunneling PEPs like QPEP. Loss of critical packets related to the key exchange process or session initialization could impose heavy additional round-trip costs not observed in clear-sky conditions. In a tunneling PEP, severe packet loss can even cause the tunnel between the PEP client and server to timeout or otherwise break. However, at mild loss levels, PEPs are expected to improve

network performance by mitigating the impact of TCP congestion-control restarts as discussed in Section 6.2.2.

Given these requirements, a series of simulations were run to assess QPEP's performance under adverse network conditions. For these experiments, losses are expressed in the form of "Packet Loss Rates" (PLR) between the satellite and the customer's satellite terminal. This represents the probability that any given DVB-S encapsulated packet is irrecoverably corrupted in transmission. Measuring "typical" PLRs in satellite networks is a deceptively complex task as definitions of both "packet" and "loss" are closely tied to the specific process by which IP transmissions are framed and fragmented by satellite ISP equipment [237]. A common worst-case upper bound often referenced in satellite broadband standards is 1×10^{-3} , but real world conditions run the gamut from "quasi error free" conditions (where nearly all packet errors are corrected by the DVB-S link layer) to rates upwards of 1×10^{-2} [238–240].

Satellite networks leverage a variety of techniques to mitigate packet losses. For example, DVB-S forward error correction (FEC) can resolve modest bit errors at the link layer, resulting in error-free IP-layer transmissions. Many modern networks now implement adaptive modulation and coding (ACM) schemes at the physical layer, allowing the network to intelligently trade-off bandwidth in favor of reliability under adverse conditions. In our experiments, these network specific particulars are abstracted away to IP packet loss rates, which allows for direct focus on the final performance implications of a problematic link. However, future work may benefit from considering other dimensions of satellite link resilience, especially bandwidth variations. That said, it is not obvious that these lower-layer noise responses, to the extent that they manifest in the same ultimate PLR, would have significant impact on the relative performance of the tools evaluated in our simulations.

In these adverse condition simulations, we first measured Iperf goodput at 20 PLR levels distributed logarithmically from 1×10^{-9} (packet loss is very rare) to 1 (all packets are lost). At each of these levels, 25 IPerf simulations were conducted




Figure 6.8: Iperf Performance in Lossy Environments. The shaded intervals represent a standard deviation in measurements across 25 simulation runs for each of 20 PLRs. Note that QPEP performance degrades rapidly in the presence of high PLRs, although it always meets or exceeds the performance of the only other encrypted tool (OpenVPN).

to transfer a file measuring 2 MB in size, amounting to a total of 2500 simulation runs. Our results are summarized in Figure 6.8.

As expected, we find that QPEP suffers at higher rates of packet loss. This makes sense as QUIC was not designed with lossy links in mind and, in particular, the loss of key cryptographic handshake packets during initialization can impose substantial RTT penalties. That said, QPEP outperforms distributed PEPsal at modest levels of packet loss and would be well suited to networks with strong signal-to-noise ratio (SNR) or physical-layer mitigations against packet loss. Generally though, this initial Iperf metric, suggests that QUIC’s cubic congestion control mechanism is not as robust to PLR as PEPsal’s TCP-Hybla based approach. Future work which adapts TCP-Hybla to the QUIC protocol may prove one avenue to maintain QPEP’s performance edge under such conditions.

It is worth noting that, regardless of PLR, QPEP consistently meets or exceeds




Figure 6.9: Mean PLT of NASA.gov Homepage at Increasing PLRs. Lines which are lower and to the right demonstrate better PLTs at higher PLRs. Note that QPEP performs better here than in the IPerf case. This makes sense as connections are relatively short-lived and some errors may be resolved by the browser (e.g., by re-issuing failed requests).

the performance of OpenVPN as an encryption tool. This suggests that, from the perspective of a security-conscious user, QPEP is net-beneficial compared to traditional VPN encryption.

Of course, as mentioned in Section 6.6.2, this Iperf benchmark only tells part of the performance story. In many cases, the short-lived data connections of web-browsing are likely more resilient to packet loss. To assess the impact of attenuation on page load times, a series of simulations were run measuring the average PLT of the NASA.gov homepage over fifty visits at each PLR interval (Figure 6.9).

Here, QPEP performs better, meeting or exceeding the performance of distributed PEPsal and substantially exceeding the performance of OpenVPN-based encryption throughout. This suggests that the goodput issues QPEP encounters at high PLRs may not necessarily translate to meaningful performance reduction for real web-browsing traffic, as QPEP’s ability to rapidly deliver small images

and text files over the latent satellite connection may counteract more error-prone delivery of larger transmissions.

In short, this preliminary look at packet loss effects suggests that QPEP is a better alternative than status quo VPN encryption under adverse conditions and performs reasonably well compared to insecure PEPs at low to moderate PLRs. However, our findings suggest that future work optimizing QUIC's response to packet loss could offer significant improvements, especially for file transfer operations.

6.6.4 Performance in LEO

While this chapter has focused on GEO networks and performance under constant speed-of-light delays, some proposed “next-generation” satellite networks focus on the use of low earth orbit (LEO) to reduce transmission latency. While GEO broadband is likely to remain relevant for the foreseeable future due to its wide coverage and heavy industry adoption, it is worth considering QPEP's performance in future LEO systems as well. Unlike in GEO, latency from LEO can vary substantially due to the shifting relative locations of satellites and the geographic position of the customer. Additionally, as LEO is much closer to the Earth's surface (approximately 2,000 km), speed of light latency effects are reduced.

To emulate a LEO system, we implement an OpenSAND simulation model which replicates observed delay characteristics from a satellite terminal in the Atlantic Ocean connecting through the Iridium LEO constellation to a gateway in London [241]. In this particular network, one-way delay varies from as low as 25 ms to as high as 140 ms, depending on the time of transmission and the route a packet must take through the constellation. The same PLT benchmark from Section 6.6.2 was repeated in this environment. The results of these experiments can be found in Figure 6.10.

As expected, the performance benefits of PEPs are much less pronounced in LEO networks and VPNs represent a more viable encryption option. QPEP still generally outperforms OpenVPN in this context, with a median PLT of 8.3 seconds compared to OpenVPN's 14.2 seconds. However, this is counteracted by the fact




Figure 6.10: ECDF of Alexa Top 20 PLTs in Iridium Simulation. As expected, QPEP offers very little benefit in this lower-latency environment. However, it also imposes smaller overheads than traditional unencrypted PEPs.

that all three PEPs seem to struggle with more complex/slow-loading pages where the added overhead of connection splitting is not always worth the benefits. It is worth noting, however, that relative to both PEPsal architectures, QPEP does appear to impose less overhead costs. Taken together, these measurements suggest that QPEP would be an adequate mechanism for providing encryption in LEO constellations but, unlike in GEO networks, the performance gains over more established VPN options are, at best, marginal.

6.6.5 QUIC Optimizations

One of the principal theoretical advantages of a distributed PEP configuration is the ability to adopt non-standard and environmentally tailored protocols over the satellite hop. In this section, we consider a demonstrative example as to how such optimizations might be identified and incorporated into QPEP.

One common strategy for improving the performance of TCP over satellite

links is ACK decimation — the process of combining many ACK messages into a single transmission at regular intervals. Unlike TCP, the QUIC protocol is not ACK-clocked which diminishes the impact of ACK decimation on goodput [242]. Nevertheless, QUIC leverages ACKs for loss detection and QUIC ACK messages are relatively large compared to in TCP contexts. This means that excessive acknowledgments can potentially congest asymmetric links [243].

We conducted a set of initial experiments to determine if QUIC ACK decimation ratios had any impact on QPEP’s measured goodput. In the default QUIC implementation, this ratio is set to 2 ACK eliciting packets per ACK for the first 100 packets, and 10:1 thereafter. Due to long satellite RTT’s however, we observed in practice that the vast majority of ACKs were triggered by the QUIC implementation’s default 25 ms ACK timeout window rather than decimation. In order to measure the effect of decimation in isolation, this timeout window was increased substantially to 8,000 ms and ACK decimation was set to begin after the 4th packet over the QUIC link. As a result, these experiments are not directly comparable to those which appear elsewhere in the chapter. We further selected the Iperf benchmark as, based on Section 6.6.3, it is more sensitive to packet loss effects that are directly relevant to ACK decimation.

In these experiments, 100 Iperf benchmarks were conducted for 5 Mb transfers at each of 30 decimation ratios. These ranged from 1:1 to 30:1. Additionally, we conducted the evaluations at three different PLRs (error-free, 1×10^{-6} , and 1×10^{-4}). The results are summarized in Figure 6.11.

We observe a few relevant trends in these results. The first is that extremely low ACK decimation ratios (e.g., 1:2) perform poorly. This makes sense as large portions of bandwidth are tied up with ACK messages at these levels. Higher ratios offer some benefit, with the default ratio of 10:1 roughly doubling goodput compared to “worst-case” 1:1 ratio. However, the benefit of ACK decimation is more limited and less consistent in lossier environments, as denoted by the relatively small goodput increase and high variance observed in our experiments. Finally, we observe that at a certain point, additional decimation has little to no effect.




Figure 6.11: Iperf Goodput vs ACK Decimation Ratio.

This may be the result of timeouts again gaining dominance. We found that, in practice, increasing the minimum ACK timeout much beyond the value used in our experiments (8 seconds), led to link instability. Specifically, unrecognized packet losses caused the Iperf client to perceive its connection to the server broken, causing it to terminate prematurely.

The high variance of these preliminary experimental results makes it difficult to definitively pinpoint an optimal ACK decimation ratio. However, one clear takeaway is that increasing the minimum ACK timeout period from 25 ms allows for better exploitation of the QUIC’s ACK decimation feature in the presence of high-latency networks. We observed a roughly 25% increase in mean clear-sky goodput (from 19.25 Mbps to 25 Mbps) as a result of doing so, even when the ACK decimation ratio itself remained at its default of 10:1. Finding an ideal ACK decimation ratio for the satellite use-case, and potentially setting it dynamically in response to noise and traffic characteristics, represents a possible avenue for further performance tuning in future work.

The ACK decimation ratio is but one of many QUIC protocol constants which may be tuned to have a meaningful impact on proxy performance. Changes in congestion window parameters, congestion control algorithms, session timeouts, and multiplexing limits may all also represent avenues for further tuning. The search space for such an optimization problem is enormous and exceeds the remit of this research — especially given that default QUIC implementations already offer substantial security and performance benefits over the status quo. Nevertheless, the approach presented through this case study demonstrates how the testbed environment and benchmarks we developed for evaluating QPEP might be leveraged more broadly for protocol performance research.

6.7 Next Steps for Secure PEPs

The QPEP implementation presented here is a proof-of-concept and productive use would benefit from additional features. In Section 6.4.3 we outline a few intuitive starting points, such as support for non-TCP protocols and the implementation of 0-RTT session initialization which is robust to replay attacks. Beyond 0-RTT, other QUIC feature proposals, such as forward error correction or alternative congestion control protocols to CUBIC, may provide routes for additional performance gains.

Under our threat model, ISPs are considered completely untrusted. This means that QPEP conceals the nature of customer traffic from ISP Quality of Service optimizations. Adding additional header layers which communicate QoS relevant metadata to ISPs, while preserving privacy, may further facilitate ISP-level integration of QPEP into customer routers.

As mentioned in Section 6.1, the principal objective of this research was to develop an encryption tool which could be used to protect TCP traffic by default in satellite networks without meaningful reductions in performance. Even default QUIC implementations meet or exceed this baseline requirement without any modification. However, future work which considers the significant but surmountable engineering challenge of optimizing the performance of a QUIC tunnel over satellite represents a logical next step.

Beyond the design of secure PEPs, the testbed presented here may be useful for more general investigations of QUIC performance over satellite. Thus, although unlikely in near-term SATCOMs environments, if TCP ends up being phased out in favor of QUIC or if “TLS-everywhere” transitions from aspiration to reality, our contributions may be of enduring use for general performance research.

6.8 Summary

In this chapter, we have challenged the historical assumption that security and performance must trade off in high-latency satellite networks. The result of this assumption has been that tens of thousands of satellite customers, from individuals to corporations, continue leak sensitive data to potential wireless eavesdroppers in the status quo. By delving into the underlying causes of inadequate encryption from GEO, we isolated key physical and commercial dynamics which have prevented the adoption of terrestrial encryption tools to the SATCOM domain.

We have presented a new approach to encrypting TCP satellite communications over-the-air through the use of QPEP — a PEP/VPN hybrid which leverages the open QUIC protocol standard to provide an encrypted UDP tunnel for the satellite hop. QPEP is evaluated through replicable simulations in an open-source benchmarking test suite we developed. These tests allow for direct comparisons between PEP and satellite encryption techniques and for targeted adjustments to various physical conditions.

Through these simulations, we find that QPEP is able to provide satellite users with over-the-air encryption while reducing page load times (PLTs) by more than 70% compared to status-quo VPNs. Moreover, we find that the use of QPEP is unlikely to result in TCP performance reductions for users who already employ insecure PEP products. Indeed, under certain network conditions, we found that QPEP may offer up to 50% performance improvement over such tools while also offering significant security benefits. We present the case that future work might expect to further bolster these gains and provide demonstrative case studies of two underlying QUIC protocol implementation characteristics as a starting point.

QPEP is the first open source PEP with support for the encryption of arbitrary TCP traffic. Moreover, unlike many commercial offerings, QPEP is fully independent, allowing individuals to run their own QPEP servers without sharing sensitive metadata with ISPs or convincing their ISPs to implement costly modifications to their existing network infrastructure. This offers an actionable near-term solution for customers interested in protecting their privacy. In the longer-term, QPEP's architecture is also suited to ISP deployment on modem equipment, allowing it to serve as a drop-in replacement for proprietary TCP PEPs. QPEP is entirely software-based and compatible with existing networking equipment and protocols. This means the practical costs for a customer implementing QPEP are on the same scale as any with any other open-source VPN. That is to say, the main implementation costs are those of renting a cloud host to run a QPEP server and paying for the desired amount of bandwidth to connect that server to the internet. This contrasts substantially with many existing PEP implementations which are implemented as physical "black-box" devices along the network path.

As the next generation of satellite broadband launches, ensuring the privacy of TCP communications without sacrificing performance is more important than ever. The QPEP proof-of-concept presented here demonstrates how careful consideration of the unique physical dynamics of outer space can leverage open and verifiable standards to meet this need.

What are we going to do with all these cats?

—Bohumil Hrabal, *All My Cats* (trans. Wilson)

7

GEO Broadband Security Lessons in Context

This section began with the recognition of satellite broadband’s key role in the evolution of the modern space industry. Over the course of the section, we made direct contributions in characterizing and defending against novel threats impacting the security of millions who rely on satellite broadband.

The starting point for this process was a breadth-first effort to update decades-old research and apply it to modern satellite networks (Chapter 4). We found that many modern systems remain trivially vulnerable to attacks which have been known in academic literature for almost 15 years. However, the nature of information technology has changed over this period to dramatically increase the severity of these issues. The emergence of cyber-physical systems and internet-connected operational technology means that what was “merely” a concerning privacy issue in 2005 is now a direct safety threat to critical infrastructures.

We next decided to investigate more modern platforms and protocols to see if their eventual adoption would automatically remediate these problems. In Chapter 5, we evaluated more complex and modern maritime VSAT services using the DVB-S2 GSE protocol. The process of intercepting signals in this format was logistically more complex, requiring novel techniques for the recovery of corrupted data. However,

beyond this complexity barrier, we found no additional security protections in these networks either. Moreover, we found that the data carried in these modern networks was perhaps even more sensitive than that which we encountered in Chapter 4. We posited several threat vectors, both active and passive, which could leverage this information to impair safety and privacy across the maritime sector.

Given that even modern enterprise protocols lacked adequate eavesdropping protections, we surmised that there may be causes which go beyond mere negligence or neglect on the part of ISPs. Chapter 6 draws on our first-hand experience disclosing these vulnerabilities to satellite customers. In doing so, we validated our supposition of distinct physical characteristics and “pain points” which frustrate traditional security approaches.

Based on this analysis, we developed a new tool, QPEP, which helps overcome these barriers. QPEP is, itself, a contribution to satellite systems security, but in building it we also developed testbed and simulation environments of broader utility. Ultimately, we found that security approaches which explicitly consider the unique requirements of satellite networks can provide actionable solutions without incurring the performance trade-offs of existing terrestrial approaches.

Beyond these direction contributions, we also identified several cultural and operational aspects of the space sector that may be informative for future researchers working on such topics.

For example, our study of the MPEG-TS MPE protocol standard in Chapter 4 gives insight into the “stickiness” of some space technologies. MPEG-TS MPE was adopted in the late 1990s, in part because it ensures backwards compatibility with MPEG video broadcasting technologies developed a decade prior. Now, more than 30 years from MPEG’s emergence, MPE is still widely used to transmit satellite internet services.

There is nothing inherently wrong with enduring standards. MPEG-TS MPE is, after all, as capable of delivering bytes to antennas today as it was two decades ago. However, the content and sensitivity of these bytes has changed dramatically as internet has evolved from novelty to necessity. MPEG-TS was never conceived

with security in mind, and that dissonance between purpose and use endures today. Security continues to be passed up to higher protocol layers, where customers and ISPs are demonstrably either unwilling or unable to guarantee it.

In short, the MPEG-TS case in Chapter 4 speaks more generally to how satellite technology evolves on different timescales and with different priorities from terrestrial analogues. This makes some sense: satellite networks as mega-infrastructure projects have natural incentives towards design conservatism. Moreover, satellite operators view security as only a small part of a much larger domain-specific optimization problem — incorporating needs which terrestrial security systems rarely account for, such as extreme transmission distances, vast signal footprints, limited on-board processing, and non-upgradable hardware.

We further see how technical practices may exhibit hereditary inclinations. Even as new protocols, such as GSE, were developed for GEO broadband, satellite ISPs continued to pass responsibility for over-the-air security along to individual customers. These new protocols provided improvements in properties which prior standards prioritized, such as bandwidth and routing, but did not reconsider the relevance of properties which were originally out of scope, such as security and encryption. That this has happened before gives particular urgency to this thesis' effort to make a case for the use of encryption in next-generation satellite networks, lest they also inherit misguided prior assumptions.

Beyond the satellite broadband context, the enduring nature of these security shortcomings also speaks to another cultural dynamic in the space technology domain. As mentioned in Chapter 6, many of the companies we responsibly disclosed security issues to were notionally aware that their communications were unencrypted. However, they had accepted these risks under the assumption that attackers would need to invest in sophisticated and costly enterprise equipment to receive these signals and that credible threat actors lacked both the means and motivation to do so. The research in this section proves that assumptions of equipment symmetry as a pre-requisite to attack are not guaranteed. Using

GSExtract, we demonstrated how imperfect equipment may be more than adequate for the purposes of sophisticated cyber-attacks.

In doing so, we have presented the case that threat models which assume the cost and complexity of space technologies are sufficient defenses against cyber-attack are unlikely to stand the test of time. When the risk of wireless eavesdropping was first accepted, it was not unreasonable to do so. It would be incredibly difficult to replicate our research if one restricted themselves to consumer-grade computing equipment which was available for purchase in 1997. However, attacker capabilities have evolved while threat models have remained static — giving rise to the harms identified in this research. Moreover, it is unlikely that these dynamics are exclusive to satellite broadband. Identifying similar assumptions of attacker incompetence in the face of technical complexity, especially with respect to older space technologies and standards, thus represents one promising avenue for uncovering further overlooked satellite security issues.

In Chapter 1, we suggested that the remit of this thesis extended beyond research into any single space technology. While our work has already made direct and meaningful contributions to help secure modern and future space systems, we aspire to further generalize these lessons. In Part III, we will distill the methodological process used in this section into a repeatable approach for cyber-physical research on space systems security. Moreover, we will put that method to the test, applying it to space security topics which are unrelated to satellite broadband.

Part III

Cyber-physical Threats Beyond the Signal

Guil: Words, words. They're all we have to go on.

(Pause.)

Ros: Shouldn't we be doing something – constructive?

Guil: What did you have in mind? ... A short, blunt human pyramid ... ?

—Tom Stoppard, *Rosencrantz and Guildenstern are Dead*

8

Through the Lens of Physicality: Putting RCMA into Practice

While the process employed in Part II led to several meaningful improvements to satellite broadband security, the methods employed were inherently exploratory and observational. One of this thesis' main goals is to provide tools which others may use to identify relevant research questions relating to satellite systems. As such, it is worth taking a moment to reflect more broadly on the contents of Part II and, with the benefit of hindsight, identify key aspects of that research which may be applicable to the domain more generally.

Broadly speaking, the findings of Part II might be grouped into four general categories.

First, we identified a number of physical properties of space-based internet services which were uniquely different from internet services in other domains. For example, the remoteness of space systems meant that radio signals from GEO reached vast swaths of the Earth's surface, encompassing areas orders of magnitude larger than even long-range cellular networks. Likewise, the remoteness of GEO platforms meant that speed-of-light limitations created abnormally high transmission latency compared to modern terrestrial broadband services — even though the traffic carried by both of these network types was often quite similar.

Second, we came to the realization that many of the security shortcomings we ran into experimentally were directly tied to these physical properties. Satellite ISPs have long relied on deep-packet inspection to optimize the traffic of their customers and make bespoke protocol modifications to adjust for latency and long-range radio communications. Meanwhile, eavesdroppers could benefit from the immense coverage area of satellite signals — intercepting wireless traffic from people in different countries or on different continents. The end result was that individual customers were unable to employ their own VPN-based encryption solutions to satellite networks, while satellite ISPs for a variety of commercial and logistical reasons did not want to change decades-old protocols to support link-layer encryption.

Third, we amassed a body of data proving that these security shortcomings were operationally relevant to real-world satellite customers. Whether in the form of admin login credentials for wind turbines, or crew member passport numbers, deeply sensitive data was being made unintentionally available to adversaries. By developing novel methods for data recovery and signal interception, we showed how inexpensive home-television equipment was sufficiently robust to raise serious safety and privacy concerns in these networks, and we shared these concerns with regulators, service providers, and customers.

Finally, we studied the interaction between physical security characteristics and state-of-the-art terrestrial security approaches. We found that, although standard VPN tools perform poorly in satellite networks, there is no technical reason that this *must* be the case. By making novel modifications to existing technologies, we were able to combine the security properties of a traditional VPN with the performance enhancements required for effective satellite broadband.

Although in Part II these findings were intertwined and scattered across several chapters, they can be distilled into four distinct research steps which might be executed sequentially. These four steps comprise the *RCMA* (Recognize, Connect, Motivate, Adapt) method which was proposed at the beginning of this thesis. Table 1.1, duplicated here for convenience, details this mapping. By *recognizing* the role of space's unique physical properties in system design, *connecting* those design

Table 8.1: Demonstrative Mapping of RCMA Method to Part II: *Threats and Defenses in Satellite Broadband*. Replicated from Table 1.1 for convenience.

Research Step	Example from Part II
Recognize physical aspects of space technologies which differ from comparable terrestrial systems.	Vast transmission distances mean that geostationary broadband has extremely high latency and broad signal footprints. (Ch. 4)
Connect these dynamics to their implications for traditional security approaches.	Broad signal footprints expose data to long-range wireless eavesdropping attacks and high latency has led to ISP adoption of PEP TCP accelerators which are incompatible with customer-operated VPN software. (Ch. 4-5).
Motivate the need for security improvements by demonstrating these impacts in a domain-specific and realistic context.	Low VPN adoption in real-world satellite broadband networks exposes sensitive data to long-range eavesdropping attacks, harming safety and security for critical infrastructure, maritime and aviation customers. (Ch. 4-5).
Adapt proven security approaches to better account for these domain-specific requirements.	We build and evaluate a novel consumer-oriented VPN/PEP hybrid which combines security with necessary TCP optimizations for satellite broadband communications. (Ch. 6).

adaptations to system security properties, *motivating* the relevance of issues in those security properties in the context of real-world users, and *adapting* proven security approaches to the unique requirements of long-range satellite communications, we contributed solutions to novel and overlooked issues in real-world space systems.

Rather than simply leaving things with the assertion that the *RCMA* approach can prove useful for future work on the topic of satellite systems security, this next section of the thesis tests that methodological hypothesis. We apply *RCMA* to two satellite cyber-security topics unrelated to satellite broadband communications. *RCMA* is unlikely to be useful for all possible security problems in space and, even if it were, a treatment of the entirety of space security is well beyond the remit of a single thesis. Our intention here, rather, is to show how the *RCMA* approach might be a useful approach to organizing research and thinking on space security topics in future work.

The two topics in this section are selected from the remaining two technical sub-domains from our taxonomy in Chapter 2: ground systems and satellite platforms.

In Chapter 9, we show how the physical complexity of resident space object (RSO) tracking creates a trust-based system which can be exploited through deception attacks on centralized space situational awareness data, a key component of satellite mission planning and operations on the ground. We relate these hypothetical threats with political and empirical motivations and develop novel threat models relevant to real-world space operations. Moreover, we present defenses to these threats which leverage recent security research in anomaly detection and adapt it to the physical dynamics of orbital motion.

In Chapter 10, we consider dynamics in space itself, looking at how the physical demands of orbital launches create permeable security boundaries aboard modern rockets. Through a security-oriented analysis of payload safety standards, we identify several safety controls which may be circumvented via cyber-mediated attack vectors and demonstrate their potential harms through dynamic simulation. Moreover, we identify minor changes to the integration certification process which can unify security and safety modeling in the launch validation process.

Unlike in Part II, where equipment and access to radio signals was relatively easy to obtain, these studies touch on topics with higher degrees of commercial or governmental sensitivity. As a result, the research here also serves as a demonstration of techniques to ground technical security research in domains where information access is restricted or otherwise constrained. Chapter 9 does this through hybrid models which combine limited real-world observational data with simulation models to help limit the number of assumptions required in our research. Chapter 10 attempts to reduce complex problems and dynamics around rocket launches to a few core features most relevant to our threat model, to apply dynamic physical simulations around that core feature set, and then to cross-check the results of those simulations with peripherally related public data sources for validation.

We will find that the RCMA approach allows us to make direct contributions to state-of-the-art knowledge in both topic areas. The work in Chapter 9 represents, to

our knowledge, the first published consideration of space situational awareness (SSA) data as a target for cyber-attackers. Moreover, we show how SSA data deception attacks are fundamentally attacks on information integrity and can thus be detected using proven anomaly classification techniques from other information security domains. Ultimately, RCMA helps draw out trust relationships in modern-day space mission design around SSA data. Moreover, it allows us to challenge the underlying causes of these dependencies, creating mechanisms for the validation and verification of third-party SSA data, even in the presence of significant resource asymmetries.

Likewise, in Chapter 10, the RCMA method uncovers previously unconsidered threat dynamics present in nearly every modern space mission. Amidst reams of safety engineering requirements, we identify a significant lacuna in cyber-security standards based on an implicit assumption of benign and informed payload developers. This allows us to reconsider secondary payload safety standards from an adversarial perspective and demonstrate one of the first space-to-space cyber-physical threat models. In doing so, we highlight specific safety controls which are particularly susceptible to malicious circumvention and thus highlight clear directions for future improvements in the payload integration process.

Viewed holistically, this section demonstrates the presence of unique threat models and security concerns across the space domain. Whether in orbit or on the ground, we find that the physical realities of space missions give rise to unusual and often overlooked information security dynamics. While it is possible to draw parallels between these problems and traditional information security networks, the research presented here makes a strong case that a naive mapping of traditional IT security practices to the space domain is insufficient for robust cyber-security. We find that the *RCMA* approach is one promising technique for identifying these gaps, validating the relevant threat models, and conceiving domain-appropriate solutions. Although the topics here are necessarily more speculative and preliminary than the deeper broadband analysis of Part II, they hopefully offer some inspiration to readers seeking additional unsolved security problems in orbit.

*Quaeritis, aetheriis quare regina deorum
sedibus huc adsim? Pro me tenet altera caelum.*

*You ask why I, queen of the gods, appear here, away
from my ethereal home? Another has taken my place
in the heavens.*

—Ovid, *Metamorphoses* 2.512

9

Deceptions and Truths in Space Situational Awareness

Space is hard. Satellites operate under constant threat from their environment, besieged by extreme thermal fluctuations and intense radiation. As orbit grows more crowded, a human-made threat has become increasingly salient: space debris. Today, more than 21,000 debris objects measuring >10 cm in diameter whiz overhead in excess of 20,000 km/h [244]. These are joined by an estimated 500,000 1-10 cm diameter particles. Colliding with any one of these objects can debilitate or even destroy a satellite.

To combat this threat, satellite operators rely on a class of data known as Space Situational Awareness (SSA). SSA describes the position, nature, and movement of space objects. Physical astrodynamic models can approximate the movement of debris objects in the short-term, but dynamic factors like solar storms, micro-debris collisions, and drag frustrate this process. As a result, SSA requires continuous updates in the form of observational measurements. Identifying and tracking minuscule objects moving at bullet-like velocities thousands of kilometers away is inordinately difficult, even with sophisticated astrometry equipment.

In this chapter, we contend that the complexity of SSA acquisition gives rise to an interesting, but largely unstudied, security dynamic. The physical reality of

space has created a situation where “ground truth” space surveillance data is beyond the means of most nation-state actors, must less commercial satellite operators. As a result, the vast majority of satellite operators must rely on shared data from one of two global SSA authorities for this mission-critical information.

We begin by characterizing a novel threat model, showing that strong incentives exist for both SSA authorities and external attackers to modify centralized repositories to abuse these trust relationships. This is used to present two distinct attack scenarios. In the first, we show how an attacker might make minute modifications to SSA data which can alter space object conjunction forecasts and mislead SSA recipients into believing their satellites are destined for collisions. In the second, we build on historical examples of space powers who have used SSA deception attacks to deceive other countries into believing that covert satellites are simply pieces of space debris.

This threat model gives rise to a fundamental question: using nothing more than an untrusted third-party’s description of RSO motion, can an SSA recipient detect deception attempts?

By framing SSA deceptions as information integrity attacks, we draw a novel connection between two intuitively distant topics: systems security research on anomaly detection and astrophysical research on resident space object (RSO) characterization. The chapter concludes by developing and evaluating proof-of-concept techniques for detecting each of the proposed attacks without the use of a single telescope.

The research method employed in this chapter tracks the *RCMA* template proposed in Chapter 1 and detailed in Chapter 8. A high level overview of this mapping and its adaptation to the SSA context is summarized in Table 9.1.

9.1 Eyes on the Sky: SSA in the Status Quo

Today, more than 2,000 satellites orbit the Earth, and this number is expected to increase by an order of magnitude over the next decade [245]. Each space mission runs the risk of adding pieces of space debris to orbit. Debris can be generated as

Table 9.1: Demonstrative Mapping of RCMA Method to SSA Deception Research.

Research Step	Example Finding
Recognize physical aspects of space technologies which differ from comparable terrestrial systems.	Robust space surveillance is difficult due to costs and geographic requirements for astrometry.
Connect these dynamics to their implications for traditional security approaches.	High surveillance costs mean almost all satellite operators must trust one of two centralized SSA authorities - limiting the ability for manual cross-verification of SSA claims.
Motivate the need for security improvements by demonstrating these impacts in a domain-specific and realistic context.	Targeted SSA deception can be used by motivated attackers, or SSA authorities, to abuse this trust and achieve strategic objectives.
Adapt proven security approaches to better account for these domain-specific requirements.	Proven anomaly detection techniques can be adapted for low-cost SSA integrity monitoring and sanity checking.

a byproduct of launch operations, weapons tests, or accidental collisions. Space debris objects threaten satellites through high-velocity collisions. Each collision can potentially cause a satellite to break up or disintegrate into thousands of additional debris objects, creating a “cascade effect” with consequences for the broader space environment [246].

The trajectory of debris objects, and even satellites, is not fully predictable. Chaotic physical forces, such as solar weather and drag, will cause deviations from projected paths over time. Space surveillance is partially about detecting these changes as they occur. From the moment a satellite launches to the day it retires, accurate and recent SSA data is critical to mission safety.

Beyond this immediate operational utility, SSA is also of significant military and strategic value. The world’s largest militaries are deeply dependent on space assets for Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR) support. This heavy reliance means space powers are highly sensitive to orbital threats. Accurate SSA allows them to monitor the behaviors of peers in orbit, detect espionage and weapons systems, and enforce diplomatic norms.




Figure 9.1: Depiction of Common Earth Orbits. Note that GEO satellites are beyond the range of GPS signals.

9.1.1 SSA in Practice

SSA measurements can be either reported or observational.

Reported measurements originate from sensors aboard satellites. The type of sensor varies depending on the object's orbit (see Figure 9.1). Satellites in Low Earth Orbit (LEO), approximately 2,000 km above the Earth's surface, receive signals from Global Position System (GPS) satellites in Medium Earth Orbit (MEO), approximately 20,000 km above the Earth's surface. Thus, LEO orbit determination is normally derived from on-board GPS data which is relayed to the Earth as telemetry messages [247]. Higher orbits, such as Geostationary Orbit (GEO), are above GNSS systems and cannot use them for positioning. Instead, they employ alternative metrics such as tracking relay satellites, star-trackers, or time-difference of arrival (TDOA) calculations at multiple earth-based antennas [248–250].

Observational measurements are required to track space debris and defunct satellites which cannot transmit telemetry directly. These measurements require sophisticated astrometry platforms. For objects in LEO, radar sensors are typically used, while, at greater distances, electro-optical telescopes are necessary [251]. A single ground station cannot reliably track objects. Instead, many observations correlated from sites distributed across the Earth.

The cost of such a system is immense. While precise numbers are scarce, the latest round of capacity upgrades for the US Space Surveillance Network is believed to have exceeded \$6 billion in procurement costs [252]. This puts in-house SSA beyond the means of not just commercial actors, but the majority of states. Even if states have sufficient resources and will (such as in the case of China or the EU), they may lack the territorial reach and diplomatic leverage to deploy radar and telescope systems across the planet.

9.1.2 SSA Repositories

The result of these cost barriers is that the majority of satellite operators rely on third-party SSA data.

By far, the dominant source of SSA is the United States Space Surveillance Network (SSN). The SSN comprises more than 20 locations, leveraging the US military's extensive network of forward deployed military installations to achieve geographic distribution. It is believed to be the only system capable of tracking smaller objects measuring 5-10 cm in LEO and 1 m in GSO [253, 254].

The closest competitor is the Russian Space Surveillance System (RSSS). Its nature and capabilities are opaque, but it is believed to consist of at least 8 sensing sites (primarily within former USSR territory) and to have a catalog about 1/3rd the size of the US SSN [253]. A nominally civilian network — the International Scientific Optical Network (ISON) — is managed by the Russian Academy of Sciences and includes some smaller academic research installations [255]. In recent years, US observers have perceived the ISON as being subsumed into the state run RSSS [256].



Figure 9.2: A *Yuan Wang* vessel used for China's SSA [257].

A handful of smaller networks exist, including the European Space Agency's (ESA) Space Surveillance and Tracking program, and systems operated by China, Canada, India, Japan, Korea, France, and Ukraine [253, 258]. Of these, China's is notable for its use of *Yuan Wang* tracking ships to overcome a lack of extra-territorial military installations by carrying mobile space surveillance systems across the ocean (Figure 9.2). Recently, commercial services promising independent civilian SSA have emerged, although at present, none offer data comparable to the SSN [259–261].

9.1.3 SSA Sharing

Most operators receive SSA data from the US SSN. The US military publicly posts SSA through `Space-Track.org` in the Two-Line Element Set (TLE) format.

This data standard was developed in the 1970s to facilitate the sharing of an object's *ephemeris* (its projected orbital path and position) using two 80-column punch cards [262]. TLEs were designed in conjunction with orbital propagation models to forecast the motion of an object through orbit. In particular, TLEs are designed for use with the Simplified General Perturbations Model (SGP4) [262].

The use of archaic data standards is ostensibly to maintain backwards compatibility. However, it offers some practical strategic benefit. As the precision of TLEs is limited, information can be disseminated more freely than if the data were more




Figure 9.3: The TLE format. The notation in parenthesis is used throughout this chapter and summarized further in Figure 9.4.

granular. For example, TLEs are sufficiently precise for satellite communication purposes and even for detecting many collision threats, but they are not precise enough for ASAT weapons targeting.

Where more precise data is required, the SSN takes a case-by-case sharing approach. This requires a formal agreement which imposes additional requirements on recipients, such as that they provide the SSN with reported telemetry measurements from their own satellites [263]. This bolsters the US military’s SSA advantage and allows them to control the flow of potentially dangerous information to untrusted states while also sharing “no questions asked” public ephemerides to prevent environmental catastrophes.

The precise accuracy of TLEs is difficult to generalize and varies depending on object location and orbit. As a rough approximation, at *epoch*, the TLE’s time of issue, it is accurate to within 1 km in any dimension; this degrades at a rate of 1-3 km per day for the first week [264]. Beyond this point, SGP4-propagated TLEs become increasingly meaningless, eventually describing impossible orbits.

The TLE format itself, along with notation used to reference orbital elements, is summarized in Figure 9.3. An overview of those elements most directly related to an orbit’s physical properties can be found in Figure 9.4. It is important to remember these values are intended as inputs to SGP4 and are not direct physical attributes.

As TLEs go out of date, Space-Track provides repository updates. While the SSN claims to screen all objects daily, daily updates are not typically provided. This is because the previous TLEs are often deemed sufficiently accurate. Space-Track




Figure 9.4: A summary of orbital elements. Adapted from [265].

also notes in its user agreement that TLE updates may be withheld for “national security reasons” or as a result of sensor anomalies. Several classified objects are also never listed in public SSN catalogs.

9.2 Threat Modeling

Space is a political domain. Modern states rely on space not just as a manifestation of national prestige, but also for operational utility. Over the past half-century, space has come to be perceived as an “ultimate high ground” for information age warfighting [266, 267]. This strategic vitality stands at odds with inherent vulnerability. Satellites are fragile devices moving at incredible speeds. A marble-sized projectile could strike a satellite with the same force as a 1-ton object falling from a height of 5 stories [268].

In the seminal days of space strategy, this physical weakness was thought to undermine the strategic utility of space itself [269]. The rise of Anti-Satellite Weapons (ASATs), which exploit this weakness, has long been prophesied to bring about the end of space power. In the first part of this section, we will look at how

enduring drivers of peace in space interact with the motivations and capabilities of cyber-attackers in the abstract. In the second part, we will apply this general analysis to the specific context of attacks targeting SSA integrity.

9.2.1 The Puzzle of Peace

Given the uncomfortable combination of high dependency and low survivability, one might expect to observe frequent attacks against critical military assets in orbit. However, despite decades of recurring prophecies of impending space war, no such conflict has broken out [270–274]. It is true that a handful of space security crises have occurred — most notably the 2007 Chinese Anti-Satellite Weapon (ASAT) test and 2008 US ASAT demonstration in response [275]. Even more recently, India demonstrated in-orbit intercept capabilities in 2019 [276]. Overall, however, the space domain has remained puzzlingly peaceful. We posit three major contributors to this enduring stability: limited accessibility, attributable norms, and environmental interdependence.

Kinetic Stabilizers

Over 60 years have passed since the first Sputnik launch and only nine countries (ten including the EU) have independent launch capabilities. Moreover, a launch program alone does not guarantee the resources and precision required to operate a meaningful ASAT capability. Given this, one possible reason space wars have not broken out is simply because only the United States has ever had the ability to fight one [277, 278]. Limited access to orbit necessarily reduces the scenarios which could plausibly escalate to ASAT usage. Only major conflicts between the handful of states with ‘space club’ membership could be considered possible flashpoints. Even then, the fragility of an attacker’s own space assets creates de-escalatory pressures due to the deterrent effect of retaliation. Since the earliest days of the space race, dominant powers have recognized this dynamic and demonstrated an inclination towards de-escalatory space strategies [279].

There also exists a long-standing normative framework favoring the peaceful use of space. The effectiveness of this regime, centered around the Outer Space Treaty (OST), is highly contentious and many have pointed out its serious legal and political shortcomings [280–282]. Nevertheless, the status quo framework has somehow supported over six decades of relative peace in orbit. Although states have occasionally pushed the boundaries of these norms, this has typically occurred through incremental legal re-interpretation, rather than outright opposition [283]. Even the most notable incidents, such as the 2007-2008 US and Chinese ASAT demonstrations, were couched in rhetoric from both the norm violators and defenders depicting space as a peaceful global commons [283]. This suggests that states perceive real costs to breaking this normative tradition and may even moderate their behaviors accordingly.

One further factor supporting this norms regime is the high degree of attributability surrounding ASAT weapons. For kinetic ASAT technology, plausible deniability and stealth are essentially impossible – at least in the context of traditional rocket-based intercept ASATs. The act of launching a rocket cannot evade detection and, if used offensively, retaliation. This imposes high diplomatic costs, particularly during peacetime, to ASAT usage and testing. Stealthier ASATs which involve maneuverable on-orbit assets, such as interceptor satellites masquerading as space debris, are likely feasible for a small number of space powers. However, the complexity and cost barriers to such technologies are significant and stealth is not necessarily guaranteed. In Chapter 10, we will explore the possibility of masquerading debris in further depth and automated techniques which may be used to detect such objects.

Finally, a third stabilizing force relates to the orbital debris consequences of ASATs. China’s 2007 ASAT demonstration was the largest debris-generating event in history, as the targeted satellite dissipated into thousands of dangerous debris particles [284]. Since debris particles are indiscriminate and unpredictable, they often threaten the attacker’s own space assets [278]. This is compounded by the Kessler syndrome, a phenomenon whereby orbital debris “breeds” as large pieces

of debris collide and disintegrate. As space debris remains in orbit for hundreds of years, the cascade effect of an ASAT attack can constrain the attacker's long-term use of space [285]. Given that any state with kinetic ASAT capabilities will likely also operate satellites of their own, they are necessarily exposed to this collateral damage threat.

Cyber-Physical Destabilizers

The overall effect of cyber-attacks vis-à-vis this strategic stability in space is not well understood. Cyber-weapons in space are often characterized as one tool among many in the growing ASAT arsenal [286, 287]. However, we contend that Cyber-ASATs pose unique strategic threats in the domain on account of their accessibility, difficulty to deter, and environmental indifference.

First, Cyber-attack capabilities are far more widespread than orbital launch technology. In 2017, a former deputy director of the National Security Agency estimated that “well over 100” countries could harm the United States with offensive cyber-capabilities [288]. Of course, mere possession of cyber-capabilities does not guarantee that these capabilities can be used against satellites. Nevertheless, this suggests that, for many actors, digital attacks are far more feasible than the creation of national space weapons programs. This calculus is further bolstered by the fact that cyber-attack capacities which could threaten satellites may apply to other unrelated systems. Moreover, while the idea of terrorist cells developing orbital spaceflight programs appears almost comically absurd, even non-state actors have demonstrated sophisticated cyber-capabilities [15, 289].

Additionally, international norms influencing cyber-combat are both younger and weaker than their space parallels. This is complicated by the plausible deniability and low attribution risk of well-planned cyber-attacks. There has been a great deal of recent debate over the ultimate attributability and deterrability of sophisticated cyber-operations [290–292]. However, few on either side would contend that cyber-attacks are as attributable as the launch of an orbital rocket from sovereign territory. A kinetic ASAT would be noticed and credibly attributed within minutes, but the

average data breach evades detection for 200 days, even for critical systems [293]. A Cyber-ASAT could lie dormant on target systems for years before triggering at a critical moment.

Finally, Cyber-ASATs undermine the ecological dynamics constraining space weaponization. Actors with Cyber-ASAT capabilities may have significantly less strategic dependence on the space environment than the major space-faring powers. As such, the deterrent effect of collateral damage through space debris would be reduced. Although debris in space can have negative commercial effects on almost all countries, in times of war this may be an acceptable cost for smaller nations with asymmetric weaknesses. Cyber-ASATs also raise the new specter of non-destructive ASATs. For example, an exploit which disables or reduces the lifetime of a targeted satellite (e.g., by wasting fuel) could prove environmentally palatable even to states with exposure to space debris.

9.2.2 SSA: Terrestrial Target, Celestial Implications

In short, while sophisticated attackers and state-actors have shown historical constraint with respect to attacks on space systems, there are no guarantees that they will continue to behave this way as the feasibility of digitally-mediated counterspace operations grows. In this chapter, our focus is on the specific use of SSA data as a target and mechanism for such attacks. The trust dynamics and operation of SSA make it a particularly interesting case to consider for a few reasons.

First, SSA repository owners have many environmental incentives to disseminate data freely. This is due to the aforementioned “cascade effect” risk, whereby debris from a satellite collision can threaten the SSA operator’s own satellites. Game-theoretic studies of SSA have demonstrated that these sharing schemes benefit all stakeholders [294]. Intuitively this makes sense: the US gains little by concealing SSA data from Russian satellite operators and causing a collision which threatens both countries. This dynamic has given rise to transnational collaboration in SSA sharing, even between states which are otherwise disinclined towards cooperation.

However, strong incentives also exist to withhold SSA data. Accurate SSA regarding military and intelligence satellites can enable adversaries to target these systems via ASAT weapons or “shadow” them with sophisticated eavesdropping platforms [295]. This is more than an abstract theoretical risk. At least four major space powers (United States, Russia, China, and India) have demonstrated ASAT missile capabilities and three (United States, Russia, and China), have demonstrated dual-use proximity operations technology. As kinetic counterspace capabilities increase, SSA operators may feel pressure to conceal the precise location of key satellites, or even their existence altogether.

This creates a dilemma. On the one hand, environmental concerns support the sharing of accurate and complete SSA; a classified satellite is no less at risk of collisions than an unclassified one. On the other hand, this same data can be abused by adversaries for counter-space or counter-intelligence operations. Moreover, if an SSA operator is caught withholding or modifying shared data, this choice can cause unwanted attention or reputational harm.

At present, the ephemerides of tens of thousands of objects are physically unknowable to anyone other than the US military. Foreign governments, including major space powers like China and Russia, face an unpalatable reality of total reliance on the goodwill of a foreign military for the safety of their own space missions. Historically, SSA sharing has been a well-regarded function of space diplomacy, and international SSA partnerships have proven resilient to broader geopolitical tensions. In recent years, however, both China and Russia have grown increasingly skeptical of the US military’s benevolence in the SSA domain [296, 297]. These states, and others, have begun investing in expanded domestic capabilities. However, the inordinate expense and complexity of such efforts means that, for the foreseeable future, SSA-sharing will remain common practice by necessity if not by choice.

9.3 Scenario: Targeting SSA Projections

Given that most actors lack the capability to independently verify SSA claims with comparable degrees of accuracy to data provided by the US SSN, one attack




Figure 9.5: A national overview of the SSA data flow and potential targets.

scenario worth consideration would be the impact of even very small changes to the data values stored in this centralized repository.

In practical terms, the mechanism of compromise is not particularly relevant. A cyber-attacker might gain access to such repositories through Stuxnet-esque attacks against sensors, direct compromise of centralized databases, modification of data stored at the flight controller’s operation center, exploitation of third-party SSA aggregation services, or alteration of data in transit (Figure 9.5). Some components of this infrastructure (such as radar sensors or encrypted connections) might require high degrees of sophistication to attack while other components (such as SSA-sharing APIs) may be within the means of most cyber-adversaries.

Once the attacker can gain access, their objective is to covertly affect satellite operator behavior. For example, they might manipulate the SSA data to make a near-miss between a debris object and a targeted satellite appear as a collision. This would cause the victim to undertake collision avoidance maneuvers, shortening the satellite’s lifetime through fuel wastage. The reverse attack could also be executed, where an attacker conceals a projected collision and destroys the targeted satellite, all without launching a single rocket.

In essence, SSA exploitation elevates simple integrity compromises into Cyber-ASAT capabilities. Furthermore, the fuel-wastage attack scenario does not threaten

collateral debris damage. As such, an attack against SSA data meets all three design requirements hypothesized in Section 9.2.1.

9.3.1 Experimental Design and Assumptions

We elected to assess the technical feasibility of attacks on SSA repositories through simulations with a commercial spaceflight planning tool [298].

The simulated attacker's overall objective was to cause an arbitrary satellite in Low Earth Orbit to take unnecessary collision-avoidance maneuvers over the next 72-hours (the current SSN emergency notification threshold). We assumed that our attacker wished to be stealthy and that significant modification of SSA data (such as the creation of new debris objects) would be detected. Finally, we granted that the attacker had already obtained the ability to modify data through traditional cyber-exploitation techniques (e.g., malware installed on the SSA web servers).

The simulations themselves were built using real-world data from the US SSN. Projections were propagated with SGP4.

9.3.2 Attack Method

Our proposed attack consists of three stages: acquisition, perturbation, and generation. In the acquisition phase, five 'near-miss' debris objects are selected as candidates for potential tampering. In the perturbation phase, the SSA data describing these objects are strategically altered to artificially cause a collision projection. Finally, in the generation stage, these alterations are merged with authentic data to create a falsified TLE entry for insertion into the SSA repository.

Acquisition Stage

To begin, an attacker must provide accurate TLEs characterizing a victim satellite's orbit and any debris objects to be considered. This information is readily available online. Our attack tool then automatically synchronizes these TLEs to a common starting epoch. From this epoch, the debris objects and victim satellite are propagated to project their locations over a simulated 72-hour period subdivided into 10-second intervals.

Dec 14 2018 02:40:54.718000000 UTC
 Target: Earth
 Source: Earth(164° RA, 31° Dec, 22809 km Radius)
 FOV: 45°




Figure 9.6: The three-step debris filter. Debris object 87848 has just entered a 1000km sphere centered on the victim satellite.

```
1 C:\dev\tle_attack\venv\Scripts\python.exe C:/dev/
  tle_attack/attack.py
2 Searching for targets
3 Propagating legitimate estimates for 72 hours (typical
  runtime ~100seconds)
4 Debris Object 89146 passes within 9.70km around 28467.
  4583333333
5 Debris Object 81683 passes within 12.32km around 28466.
  6250000000
6 Debris Object 82637 passes within 19.95km around 28468.
  9166666667
7 Debris Object 81096 passes within 27.39km around 28468.
  5833333333
8 Debris Object 87235 passes within 93.86km around 28467.
  7083333333
```

Figure 9.7: Typical acquisition stage output.

At each interval, a three-step filter is employed to remove irrelevant debris objects (Figure 9.6). First, we select only debris objects currently inside the victim satellite's orbit plane (represented by a 100 km deep cylinder centered at the Earth's core and oriented along the victim's orbit). Second, we remove debris with altitudes outside a range bounded by the victim satellite's perigee (lowest orbital altitude) and apogee (highest orbital altitude). Third, we remove debris objects more than 1000 km away from the victim satellite in any direction.

For any debris which survive this filtering, we calculate the time and distance

TLE Field	Maximum Alteration	TLE Precision
Orbital Inclination	$\pm .1$ degrees	.0001 degrees
Right Ascension of the Ascending Node	$\pm .1$ degrees	.0001 degrees
Eccentricity	$\pm .01$	.0000001
Argument of Perigee	$\pm .1$ degrees	.0001 degrees

Table 9.2: Modified TLE Fields and Tampering Boundaries

of closest approach to the victim over a full orbital period. Ultimately, the five objects which pass closest over the whole 72-hour window are selected (as in Figure 9.7). TLE data for these objects is passed on to the perturbation stage along with times of their closest approaches.

Perturbation Stage

In the perturbation stage, TLEs of the five selected debris objects are altered with the goal of reducing the projected nearest pass distance to the target to less than 1 km. This is based on Air Force Space Command guidance that TLEs can be considered accurate to approximately 1 km of precision. Any object which passes within this range could thus trigger an anticipated conjunction.

In order to reduce the risk of detection, two further constraints are imposed. First, only four TLE fields (along with the TLE checksum) are subject to modification. Moreover, these fields are altered within certain boundaries (detailed in Table 9.2). To our knowledge, no study has been done as to what extent, if any, satellite operators vet SSA data for anomalies. As such, these boundaries were selected arbitrarily based on the overall precision of the TLE format (also detailed in Table 9.2). Decreasing these bounds lowers the chance of detection but increases computational complexity.

SGP4, like most orbital projection models, is complex and the overall effect of any given modification over a 72-hour window is non-trivial. However, we can greatly reduce this complexity by recognizing that there is no need to find the optimal perturbation set, but rather only an adequate set to cause a collision.

This realization allows us to employ a rudimentary genetic algorithm where we treat the TLE fields themselves as genetic features. Our model's fitness is simply the

```

10 Launching attack on TLE data
11 ***** Running GA for 89146 *****
12 gen nevals   avg      std      min      max
13 0      200      8.71705 0.516543  7.56435 9.89076
14 1      104      7.98663 0.415876  6.01983 9.95535
15 2      118      7.49821 0.51147   6.01983 9.39338
16 3      116      6.62903 0.574535  4.87107 8.70708
17 4      127      5.94082 0.474319  4.11844 7.63035
18 5      132      5.12766 0.64398   2.82309 8.6329
19 6      120      4.27379 0.573196  2.48353 6.74056
20 7      118      3.52179 0.664061  2.21946 6.82699
21 8      120      2.75998 0.605173  1.26693 6.50113
22 9      105      2.29535 0.410936  1.2321  4.37685
23 Search Completed on generation: 10
24 Malicious TLE for object 89146 with pass distance of 0.
    5720459504

```

Figure 9.8: Typical perturbation stage output. In this case, a set of modifications was detected that caused debris object 89146 to pass within 600m of the victim satellite.

minimization of nearest pass distance and our initial population size is arbitrarily set to 200 individuals. Over a span of up to 40 generations, each individual is used to generate a fake TLE and propagated for the 3-hour period surrounding the debris object's closest approach (Figure 9.8). Once a sub-1 km pass is found, this result is passed along to the generation stage.

Our naïve genetic algorithm may be further optimized. However, the operational benefit of finding a pass within 10 meters versus a pass within 900 is minimal since both fall within the collision detection radius. Further, given that an attacker has hours, if not days, to calculate these modifications, computational efficiency is far from vital.

Generation Stage

In the generation stage, the results of the five genetic algorithm runs may be compared using two further metrics:

- The proximity of the projected pass caused by a malicious TLE
- The overall magnitude of modifications introduced into a malicious TLE

The first metric is useful for an attacker who wishes to have the highest likelihood of causing a satellite maneuver. The second metric would be more desirable for

```

1 89146U 00000AAA 18347.88483769 .00000000 00000-0 10326-3 0 9999
2 89146 098.0408 311.5309 0132000 353.5856 290.3549 14.41709923258234

```

Figure 9.9: A typical original TLE.

```

1 89146U 00000AAA 18347.88483769 .00000000 00000-0 10326-3 0 9999
2 89146 098.1129 311.4806 0163674 353.6118 290.3549 14.41709923258239

```

Figure 9.10: A typical malicious TLE.

attackers seeking to minimize the risk of detection. An attacker can also ignore these metrics and simply select the first valid attack found to minimize search time.

Once a malicious TLE parameter set has been found, its modifications are merged with data from the original debris TLE (as in Figure 9.9). The result of this process is a new TLE which can be inserted into the SSA database by an attacker as required, completing the attack (Figure 9.10).

9.3.3 Attack Simulation and Results

To test this approach experimentally, we simulated attacks against each of 111 satellites in the Iridium constellation. Iridium is a commercial communications service with over one million satellite customers [299]. The network’s largest customer is the US Defense Information Systems Agency [300]. For our debris field, we selected 529 objects from **Space-Track.org**’s “Well-Tracked Analyst Objects of Unknown Origin” dataset [301]. Prior to launching our attack, none of the Iridium satellites were projected to pass within 1 km of these debris objects over a 72-hour window.

In order to simulate attacks against many satellites quickly, we enforced no optimizations in the “generation” phase. This means our experiment represents the ‘worst-case’ scenario for our method in terms of pass distance and stealth.

Our technique successfully generated collision events for more than 93% of the Iridium constellation. On average, it took about 12 genetic generations to find a valid attack and the total attack run time for each object averaged a little over 6 minutes on consumer-grade hardware.

Although we accepted any pass under 1 km, the mean pass distance of our attack parameters was around 600 m and the minimum only 2 m. No obvious




Figure 9.11: Original and malicious pass distances from the Iridium attack simulation.

correlation between original pass distance and malicious pass distance was observed (Figure 9.11). This suggests more restrictive boundaries and more demanding proximity requirements are obtainable using this general approach.

Our findings demonstrate that, once an attacker has compromised the integrity of an SSA repository, elevating this to ASAT capability is quite feasible. With consumer-grade hardware and a minimally optimized attack method, we falsified collision projections for over 100 real-world satellites used by the world's largest militaries.

9.4 Scenario: Classification Deceptions and Defenses

While forecast tampering attacks are interesting due to their high potential harms, the scenario is inherently quite speculative. The feasibility of the attack is highly dependent on the specific collision modeling, propagation, and maneuvering criteria used by a given space operator. As a result, we have elected to follow this preliminary threat model with a more rigorously defined and empirically grounded attack

scenario. The benefit in focusing on such a case is that it is possible to better consider real-world defensive mitigations in context.

In this second scenario, the “attacker” is an SSA repository owner who wishes to deceive a third party SSA recipient (the “defender”) into believing that a spy satellite is a piece of space debris. In doing so, the attacker seeks to bypass the aforementioned safety/privacy tradeoff: sharing the location of the object accurately to avoid collisions, but also preventing adversaries from learning the true nature of a strategically important space asset.

This has several benefits for the attacker when compared with merely omitting the secret satellite from the SSA repository altogether. By sharing accurate data indicating that there is *some* object in that orbit, an attacker can issue conjunction alerts to protect that object without arousing suspicion. Moreover, while other SSA operators may be able to determine the presence of an object in that orbit, determining the object’s purpose is a much more complex task requiring significant manual effort and sophisticated equipment. Against a backdrop of tens of thousands of other debris objects, defenders must prioritize and are more likely to allocate limited resources to studying objects explicitly designated as satellites.

For an attacker, success would allow them to prevent defenders from altering their behaviors during overpasses — allowing for the collection of imagery intelligence (IMINT) data. Alternatively, it allows the attacker to protect key communications missions from ASAT weapons, or to block competitor access to a high-value orbit without negotiating appropriate rights.

For a defender, the primary objective is to identify which objects in an SSA repository are maliciously misdesignated. Under our threat model, we assume the defender is a nation-state actor with limited SSA capabilities relative to the attackers. For example, states like China and Japan have only small SSA networks. By flagging the most suspicious objects, they could optimize limited resources.

9.4.1 Empirical Support

The core conceit of this attack — using space debris as a disguise for critical space systems — is nearly as old as spaceflight itself. Throughout the Cold War, both the US and USSR dedicated significant resources to the interception and monitoring of space debris under the assumption that it might include disguised hardware used for transferring mission data to Earth [302]. For example, the top secret 1959 US spy satellite “CORONA” ejected film canisters disguised as orbital debris fragments to transfer IMINT to Earth [302].

Much more recently, Russia was credibly accused of deceiving the international community by claiming a piece of orbital debris has been generated from a Rokot-Briz launch on May 9th, 2014 when, in fact, the object in question was a nano-satellite [303]. At least in unclassified contexts, this deception worked, with the US military tracking object *2014-28E* as a piece of space debris for more than six months. In November 2014, the object began to engage in orbit-altering maneuvers, revealing its true nature as a satellite. It is worth noting that the majority of nano-satellites never engage in maneuvers and, if *2014-28E* was a typical “CubeSat,” it would have evaded detection for much longer.

In 2015, Oleg Maidanovich, head of Russian space command, reported that they had identified a cluster of espionage satellites which were masquerading as space debris objects [304]. Maidanovich declined to provide further detail on either the objects or how they were detected, but Russian press coverage heavily implied US involvement. Regardless, this demonstrates that Russian military officials perceive this threat model as a plausible attack vector.

Unofficial statements from individuals within the defense community further support an abstract interest in developing nano-satellites which are small enough to be only trackable by the SSN [305]. While comments from the US National Reconnaissance Office (NRO) on their use of nano-satellites are characteristically vague, intelligence experts have argued that the ability to evade detection by foreign SSA networks is one driving motivation [305].

Finally, in 2020, the US Department of Defense and Department of Commerce began negotiations on a series of regulations restricting the imaging of space objects. Of particular relevance are the requirements that space-to-space photographs of satellites and debris have a maximum resolution of 50 cm and a blanket ban on photographing any objects not listed in the Space-Track.org repository [306]. Many nano-satellite platforms are smaller than 50 cm in any dimension. For example, CubeSat form-factors start at 10 cm³.

In short, recent developments point to a growing interest in concealing operational satellites as pieces of space debris as evidenced by the behaviors of both of the world's leading SSA authorities (the US and Russia). For states which rely on data from either power, detecting SSA deceptions will become increasingly relevant.

9.4.2 Experimental Design

Given the lack of public specifics on historical SSA attacks, defense evaluation hinges on the development of realistic simulations. We seek to minimize the assumptions required by focusing on a foundational threat model: lying about a space object's purpose.

A traditional conceptualization of this challenge would revolve around improving physical sensors and measurement techniques. However, we take the unconventional approach of treating it as an attack on digital information integrity. This gives rise to the defensive strategies outlined in Section 9.4.4 and lets us leverage existing work on anomaly detection to both build and evaluate our system.

Treating this as an integrity attack helps us incorporate real-world data. Rather than attempting to emulate the launch of spy satellites through inherently imperfect physical simulations, we can instead “tamper” with the contents of existing public ephemerides. This allows us to ensure that our evaluation and simulations account for the diverse breadth of satellite and debris characteristics.

That said, there is at least one critical assumption required. We must assume that the real-world ephemeris data used is reasonably trustworthy. The physical reality of SSA collection means that, to have any real-world data, we must start

with the assumption that the US military is not already trying to deceive us at scale. The reasonableness of this depends on personal political perspectives.

However, we need not have perfect faith in US SSN data. Even if this “ground truth” information includes a small number of maliciously concealed satellites, they are highly unlikely to represent a meaningful proportion of the overall dataset due to the reputational and environmental dynamics discussed in Section 9.2.2. Given that the ability for defenders to determine *anything* about the nature of an object from its public ephemeris alone would represent a substantial improvement, a theoretically perfect system is far from necessary. Nevertheless, it is worth remembering that, when we discuss our model’s accuracy at recognizing satellites/debris objects, this is, in fact, shorthand for its accuracy at predicting whether the US military would publicly label an object as such.

Data Description

Our source of SSA data is the United State’s SSN catalog. We consider one year (March 1, 2019 to March 1, 2020) of public SSA. This amounts to approximately 2 million TLEs describing around 19 thousand objects. Roughly 40% are satellites and 60% pieces of orbital debris. However, satellite ephemerides tend to be updated at slightly greater frequencies, meaning that the dataset is roughly balanced with respect to individual TLE classes (see Figure 9.12).

A descriptive summary of individual orbital elements appears in Table 9.3. We have excluded the second time derivative of mean motion ($\ddot{n}$) from analysis as the vast majority (>99%) of records report this value as zero. Neither $\ddot{n}$ nor $\dot{n}$ are used by SGP4 but are legacy holdovers from older propagators [307].

Note that several values are distributed across the full range of the TLE format. This holds for many angular elements (Ω , ω , M , and, to a lesser extent, i). This makes sense as orbits are dictated by mission requirements and variation in angular elements encapsulates their spread around the Earth. For other features, the physical requirements of maintaining a stable orbit manifest clearer boundaries on the parameter space. For example, n , roughly a proxy for satellite velocity and




Figure 9.12: The distribution of SSA entries by object type. The width of each plot is proportional to the number of objects which have been updated at a given annual rate.

Table 9.3: Distribution of Orbital Element Values.

Element	Mean	Std. Dev.	Min	Median	Max
$\dot{n}$	0.00004141	0.00104152	-0.11202822	0.00000097	0.99999999
B^*	0.0003716	0.0231546	-0.9934500	0.0000486	20.7840000
i	72.8193	31.5599	0.0000	82.7970	144.6450
Ω	177.8059	108.2238	0.0000	175.6950	359.9998
e	0.0563491	0.1644054	0.0000002	0.0039360	0.9184180
ω	175.9580	104.4194	0.0006	172.2482	359.9998
M	183.0258	107.9505	0.0004	186.4745	360.0000
n	11.773592	4.846328	0.037454	14.026936	16.474780
rev	37,138	28,991	0	31,266	99,999

Note: The number of significant figures is representative of the precision of the TLE data format. The only exception is B^* , which has variable precision.

altitude, has an upper-bound of around 17 revolutions per day. Although the TLE format could allow for the representations of satellites with greater velocities (up to 200,000 km/h in LEO) such claims are physically implausible. For this experiment, we assume that the data is presently trustworthy. That is to say, the US military

has not, over the one-year period concerned, engaged in SSA deceptions. The reasonableness of this assumption depends largely on personal political perspectives. However, even if the US military does presently engage in SSA classification attacks, they are unlikely to do so for a statistically meaningful proportion of the SSN repository. Thus, while treating SSN data as “ground truth” might lead to minor accuracy discrepancies, it is unlikely to cause radical performance deviations.

A second source of SSA, such as Russia or China’s, would have been desirable. Unfortunately, such data is not generally available. Commercial SSA is similarly difficult to access, with models revolving around direct contracts with national militaries or agreements with defense-industrial companies [308]. To our knowledge, the US military is the only entity which publicly shares a robust SSA catalog, although some organizations will repackage and redistribute this data with additional processing.

Fortunately, this SSA is representative of real-world practice. It is likely that nearly all space-faring nations at least consider the SSN as a factor in their SSA processes. This is particularly true for smaller objects, such as 10 cm CubeSats and small debris particles, which are believed to be only trackable using SSN’s technology. The main variation between defenders is thus less in terms of whether they must trust US data and more on the extent of that dependence.

Attack Implementation

In the SSN, debris objects are designated by a naming convention whereby the suffix “DEB” is appended to the name of their originating mission. In the simplest sense, we can replicate an object deception attack by appending the “DEB” label to a satellite’s descriptor. It is important, however, to do so for all TLEs belonging to an object from the moment of launch onwards, or defense models may receive prior information which biases performance. Additionally, TLE parameters such as the object’s ID and name are trivially alterable by an attacker. We thus assume that the attacker optimally alters all such metadata elements.

As mentioned in Section 9.2.2, the attacker has strong incentives *not* to alter the orbital elements themselves. This is because accurate positional information




Figure 9.13: A photograph of a 20 cm satellite by the Thai National Space Objects Observation [310]. (License: CC-BY 4.0)

is necessary to issue conjunction alerts which protect the concealed satellite from orbital collisions.

9.4.3 Signature Extraction Approaches

Differentiating between debris and satellites on the basis of their motion alone is a complex problem largely neglected in prior work. In the status quo, RSO classification is handled by SSA operators who leverage billions of dollars in labor and equipment to perform this task. Even with high-end astrometry equipment, both satellites and debris objects appear as little more than tiny dots, making determining their nature a challenge (see Figure 9.13). This may explain why the US SSN failed to identify *2014-28E* following its launch. When observation is required, the process can entail many measurements in order to build composite signatures from physical attributes, such as light reflections [309].

Fortunately, such effort is only rarely required. SSA operators typically can rely on reported data from satellite owners and launch operators. Following a launch,




Figure 9.14: Distribution of key orbital elements by object type. † The B^* and e plots are “zoomed in” around the 10th-90th percentile values to improve legibility due to heavy clustering in both elements.

they collaborate with owners to hunt down and make radio contact with platforms. Any objects not identified in this process can be considered debris.

On Physical Signatures

An initial, but misleading, intuition might be to use maneuvers, such as the one which ultimately revealed *2014-28E* (see Section 9.4.1), as a signature to identify satellites. However, many small satellites never maneuver. Nano-satellites reach orbit by “hitchhiking” through vehicle-sharing agreements. At some point in the launch sequence, they are pushed away from the rocket — normally using a simple spring-powered mechanical plate [311]. From this point, they spend the entirety of their lifespan adrift. Even when larger satellites *do* have maneuvering systems, they use them only sparingly. Moreover, space debris objects themselves can appear to “maneuver” as the result of space weather effects.

Still, there is intuitive reason to expect differentiation between debris and satellite orbits. While space missions are diverse and their orbits vary, all satellite orbits are fundamentally *selected* by sentient individuals while debris orbits are *incidental* to this selection process. The selection criteria for space missions is also generally guided by a few common factors. For example, communications constellations will

attempt to maximize the time at which certain points on the Earth's surface have line-of-sight to one or more satellites. Imagery and remote sensing systems will prefer orbits which offer sensor coverage of specific regions.

Even more fundamentally, cost and convenience influences the selection of orbits. Some missions seek to maximize longevity and minimize drag. Others may be required by regulators to de-orbit after some time and may prefer orbits which naturally cause this. Finally, small satellite payloads are often forced into orbits based on the requirements of primary payload they “ride-share” with.

Debris faces none of these constraints. While debris is made up of materials that, at one point, belong to a “designed” mission, its position is the result of chaotic physical dynamics.

It is tempting to imagine a hypothetical checklist of physical characteristics which separate the signatures of “designed” and “incidental” orbits. However, the diversity of both spacecraft missions and debris objects increases the complexity of such a task substantially. State-of-the-art methods for space object classification and shape determination thus continue to rely on photometric and other observed characteristics [312, 313].

In our threat model, the fundamental challenge for defenders is their lack of access to this direct physical knowledge. Instead, they must leverage a restrictive set of public TLE parameters, all of which exhibit substantial overlaps across the population of objects (see Figure 9.14). While we can observe certain distribution biases in favor of a particular class, such as the many near-zero n values belonging to GEO communications satellites, these distributions hint at the inherently multidimensional nature of our defender's problem. It makes sense that such a task is non-trivial as it is unlikely that states would continue to invest billions of dollars in laser/optical/radar RSO classification systems if the matter could be trivially resolved by, for example, considering the altitude of an orbit or the separation of a given object from others around it.

One feature in Figure 9.14 appears promising: *revs*. This is misleading. As *revs* is simply the count of revolutions an object has made since it first appeared

in the database, the feature lacks explanatory power for newer objects. The clustering observed is thus reflective of increased satellite launches in recent years coupled with a handful of major debris generating collisions (e.g., the 2007 Chinese ASAT demonstration). While a defender might use *revs* as a components of their classification system, for example to exclude objects from these events, it alone is insufficient for differentiation.

Why Use Machine Learning?

An alternative to manually devising a complex rules-based system for orbit differentiation would be the use of machine learning classifiers. The intuition here is that classical machine learning methods may allow us to effectively grapple with the astrodynamic complexity and multidimensionality of RSO classification. The focus of research to date has been on improving the accuracy of orbital forecasting or improving the RSO classification value of sensor data [312, 314–316]. To our knowledge, no attempt has been made to tackle the RSO task with a feature set as limited as the one available to defenders in our threat model.

Nevertheless, this research offers valuable insights for our own. For example, Furfaro et al. found that the light curves generated in simulated telescope images were differentiable using a convolutional neural network (CNN) [317]. Using photographs from one Russian SSA telescope, they managed to classify three object classes (Rockets, Satellites, Debris) with roughly 77-85% accuracy. Similarly, Liu et al. proposed an ontology-based classification model which combines a holistic rules-based approach and classical machine learning [318]. This model is used to identify space mission purposes (e.g., commercial vs remote sensing). Again, they rely on supplemental data not available in our threat model, such as telescope brightness, radar cross section sizes, power supply, object mass, and object ownership records. They conclude with the finding that their approach matches the performance of a random forest (80-90% accuracy), but with reduced training times.

These studies suggest that machine learning can solve RSO classification tasks that have proven too costly or difficult using manually-tailored physical models.

However, prior work assumes access to data which our defender is unable to trust. One area where solutions to this issue may be found is the systems security domain. There, researchers have leveraged machine learning techniques to detect implausible relationships between elements, even in untrusted data.

For example, decision trees have been used to identify fishing vessels pretending to be vessels of different types on the basis of their location and motion information as reported in Automatic Identification System (AIS) messages [319]. The underlying concept here — detecting incongruities between tampered object labels and position data concerning those objects — is quite similar to our own task. Likewise, decision trees have been used to detect covert channel communications disguised as multimedia applications and various network identifier spoofing threats [320, 321].

In short, existing systems security research makes the case that machine learning approaches can help draw out complex interrelationships between immutable physical characteristics. These relationships can be used to assess the plausibility of related, but more mutable, claims in the presence of an attacker.

9.4.4 Defense Implementation and Evaluation

In designing our defensive system, it is important to highlight that machine learning is not, itself, the focus of our research. Rather, it is a tool for evaluating a more strategically relevant hypothesis: that even basic positional data about a space object contains signatures which may expose concealed satellites. That said, one of our stated objectives is to facilitate future work in the domain. To this end, we start by briefly highlighting feature-engineering challenges and pitfalls that arise from the unique nature of SSA data.

Initially, one might treat this as a straightforward binary classification problem with conveniently pre-labeled data. Running a trivial k -nearest neighbors classifier (KNN) trained on randomly sampled TLEs detects satellites astonishingly well ($f1$ -score: 0.92, $precision$: 0.90, $recall$: 0.94). However, these results are deceptive. This is because there are spatio-temporal interrelationships between TLE entries which “leak” information.




Figure 9.15: Demonstrative KNN performance comparisons. Note that a naive model performs quite well, while a model trained under realistic conditions does not.

For example, a naive classifier can over-fit if training data contains prior (or future) TLEs belonging to an object in the test data. If we restrict the evaluation data only to *objects* not included in the test data, performance plummets (Figure 9.15). Moreover, information about related objects, such as the future location of other satellites from the same launch, can be used to foresee clustering patterns unavailable to a defender who lacks time-traveling capabilities. After restricting training sets to TLEs before a certain epoch, we find our initially promising classifier now performs only superficially better than random guessing (*f1-score*: 0.58, *precision*: 0.63, *recall*: 0.55).

Baseline Performance

In this chapter, we will focus on decision-tree classifiers as they have proven viable in prior work on both RSO and related anomaly detection tasks (see Section 9.4.3). While there are a plethora of machine-learning approaches that could be employed, applying them here serves little purpose other than academic diversion. If decision trees can detect object-type signatures in TLE-data that is sufficient to prove the existence of such signatures, even if another technique (e.g., neural networks) could also do so.

Table 9.4: Decision Tree Performance Comparison.

	Decision Tree	Bagged Tree	Random Forest	Histogram Boosted
Accuracy	0.89	0.91	0.93*	0.92
Precision	0.91	0.94*	0.94	0.91
Recall	0.87	0.86	0.92	0.93*
F1-Score	0.89	0.90	0.93*	0.92
TPR	0.87	0.87	0.92	0.93*
TNR	0.91	0.95*	0.94	0.92
Area Under ROC	0.89	0.91	0.93*	0.92
Train Time (s)	16.09*	30.33	81.77	19.56

* denotes the column with the best value for a metric. The target for true-positive rate (TPR) and true-negative rate (TNR) is “Satellite.”

On the basis of these intuitions, we tested four popular decision-tree based classifiers against our attack: a CART decision tree [322], a bagged meta-estimator using “random patches” [323], a basic random forest [324], and a histogram-based gradient boosting model [325].

These classifiers are trained using TLEs which have been appropriately segregated on the basis of object identifier and observation time. The training set consists of ~800,000 TLEs describing ~15,000 space objects. The test set consists of ~200,000 TLEs describing ~4,000 space objects. The elements n , B^* , Ω , ω , $revs$, $\dot{n}$, i , e and M are used to predict one of two labels: “Debris” or “Satellite.”

The relative performance of these classifiers is summarized in Table 9.4. We find that all four models far-exceed random-guessing in terms of their predictive ability to determine an object’s type on the basis of its TLE. A histogram boosted model offers the greatest performance to training-time ratio. However, the practical difference compared to a basic random forest model is marginal, with the random forest slightly outperforming in some metrics (such as TNR) while slightly under-performing in others (TPR). As a few minutes of additional training time has essentially no impact on testing our hypothesis, we have elected to focus on the random forest approach as it is widely used and well understood in both aerospace and security communities.

The dominant features leveraged by the random forest classifier are summarized in Figure 9.16 in terms of permutation importance [326]. These features make sense




Figure 9.16: Random Forest Feature Importances.

in light of our reasoning on the plausible physical differences that could be used to distinguish debris objects from satellites. i , e , and n roughly express the shape of an object's orbit and its altitude. While Ω , ω and M also play some role in this regard, Figure 9.14 suggests distribution across the full range of these values, while i , e , and n sit within tighter bounds. Additionally, the importance of B^* aligns well with the supposition that debris orbits tend towards different shapes and altitudes. We find that, contrary to our initial intuition, $revs$ is a relatively weak indicator. Indeed, excluding it entirely has no meaningful impact.

Final Model

A balanced random forest classifier does not quite meet our defender's requirements. After all, our defender doesn't want to sort isolated TLEs into categories, but rather to determine the nature of the underlying space objects behind them.

We can better serve this need by implementing a voting mechanism which aggregates label probabilities across all of an object's TLEs and applies the most favored label from multiple observations. This improves the Random Forest's performance slightly: TPR increases from 0.92 to 0.93, TNR from 0.94 to 0.95, and area under ROC from 0.93 to 0.94. One might expect a defender could further improve on this by incorporating history as part of the feature space

directly. However, grouping the previous 20 TLEs of a given object into a single feature-set actually results in worse performance due to increased dimensionality (*f1-score*: 0.91).

The voting model we use offers an additional benefit by providing a simple mechanism for incorporating defender astrometry capabilities into the classifier's tuning. By decreasing the voting threshold required for classification, a defender can improve their chances of detecting a hidden satellite at the cost of an increase in false positives. For example, reducing the voting threshold from a balanced 50% to 25% increases TPR to 98% with an FPR of 11%. For mid-sized SSA powers, such as China and Russia, these numbers likely represent manageable false positive rates.

9.4.5 Defense Evaluation

While the machine-learning accuracy metrics tell us about system performance in pure numerical terms, it is difficult to understand what a 93% TPR means for defenders in terms of required astrometry capabilities and investment.

One way to contextualize these values is by simulating a hypothetical attack scenario several thousand times and seeing how the defense performs. For this scenario, our attacker is assumed to launch two covert spy satellites over the course of a single year, these TLEs for these two malicious objects are mingled with a set of 398 legitimate debris objects, all drawn from data previously unfamiliar to our classifier. This total number of 400 objects is roughly equivalent to the number of "debris" entries added to the Space-Track catalog in 2019 and represents a full year of catalog updates. In total, we preformed this experiment 10,000 times with randomly selected sets of attack and debris objects.

The defender's task is to identify the spy satellites. However, we can also use this scenario to better understand the relationship between defender capabilities and attack detection rates by considering the voting-threshold discussed in Section 9.4.4 as a second independent variable — running the 10,000 simulated attacks at a range of voting-threshold intervals between 0% (nearly all objects are flagged as suspicious) and 100% (almost no objects are flagged).




Figure 9.17: Defense Performance over 10,000 Simulations. At high voting thresholds, a defender receives fewer annual alerts, but may miss the spy satellites.

The results are summarized in Figure 9.17. They suggest that even a small investment in SSA verification capabilities can substantially improve a defender’s likelihood of detecting attacks under realistic conditions. For example, a defender with the ability to measure 50 RSOs each year can set their voting probability threshold to 22%, allowing them to detect at least one spy satellite in 100% of the simulations and both spy satellites with a 98% success rate. Even a defender with the ability to only vet one object per month can detect at least one spy satellite 96% of the time and both 70% of the time. In a “worst-case” scenario where the defender has no astrometry capabilities, the defender’s best guess of the two most suspicious objects includes at least one disguised satellite 82% of the time and both 35% of the time.

2014-28E

As mentioned in Section 9.4.1, the closest empirical example of this attack is the case of NORAD object *2014-28E*. We can use the *2014-28E* incident as a demonstrative test of our model by excluding it from the random forest’s training data. This is a particularly interesting case as it was launched with the conscious intention of evading detection and “behaving” like a piece of space debris. Moreover, the attacker (Russia) and the victim (the US SSN) are the world’s two dominant SSA powers.

Table 9.5: “Analyst Objects” Most Like Satellites

NORAD Catalog Identifier	Satellite Classification Confidence
85412U	0.980
87934U	0.962
84006U	0.960
87730U	0.957
82939U	0.945
89177U	0.928
82891U	0.911
85410U	0.910
84353U	0.903
85205U	0.895

We find that our model flags *2014-28E* as a disguised satellite from the very first TLE entry in Space-Track.org - six months faster than existing monitoring techniques did. It does so at a $> 91\%$ confidence, which equates to 2.55 total anomaly detection alerts per year under the conditions simulated in Section 9.4.5. This suggests that essentially any implementation of our system — whether or not it assumes defender access to astrometry equipment — would have identified Russia’s claim that *2014-28E* was space debris as dubious.

Analyst Objects

As a matter of interest, the US SSN maintains a secondary list of objects outside the main catalog. These “Analyst Objects of Unknown Origin” are tracked by the SSN, but the US military has been unable to determine their origin. One would expect the vast majority to be orbital debris and, indeed, the SSN states that “Analyst objects that meet well-tracked criteria are generally debris objects, as is the majority of the space catalog” [327]. By applying the defensive classifier to these TLEs, it may be possible to determine which are most like concealed satellites.

We test this experimentally on a dataset of approximately 23,000 ephemerides relating to 519 analyst objects over a one-year period. At a 50% voting confidence threshold, 43 objects are flagged as potential satellites. Our expected false positive rate at this level is 5.7%. Increasing the threshold to 89% provides a list of the

ten analyst objects most like satellites (Table 9.5). The false positive rate at this threshold was only 0.4%, so one would expect, in absence of deception attacks, only two objects to appear on such a list.

This is, of course, not conclusive proof that the US military either accidentally or deliberately miscategorizes operational satellites as analyst objects. The astrometry equipment necessary to ascertain the nature of these objects is well beyond our means as researchers. However, these ten entries represent an intuitive starting point for states interested in evaluating such a possibility.

9.4.6 Scenario Discussion

Under our threat model, attackers initially appeared incredibly powerful, with full control over access to “ground truth” information regarding the nature of objects in orbit. However, a closer analysis shows that strong environmental and strategic considerations bound the extent to which attackers are able or willing to modify SSA claims. Acknowledging the existence of these constraints has allowed us to present an approach by which SSA recipients can significantly increase their chances of detecting satellites masquerading as debris objects.

Implications for SSA Operators

Given these findings, one might expect attackers to alter their space missions to improve stealthiness. For example, they might launch spy satellites into orbits designed to emulate debris objects as a sort of “adversarial attack.” While this is certainly possible, it is not without cost. Deploying a satellite which exhibits such behavior would likely require increasing its size and weight for maneuvering hardware — decreasing its physical stealth. Similarly, the selected stealthy orbit may be commercially/strategically flawed, shortening mission lifespan or coverage of key surface locations. Even where this factors can be mitigated, the effort of doing so can increase mission design costs and time.

A more promising tactic for attackers would be to modify the contents of SSA *data* while keeping the mission unaltered. Again, this is non-trivial as attackers

wish to keep the forecasts generated by propagation models which use this SSA data sufficiently accurate to prevent orbital collisions. Even minute modifications to orbital elements can result in changes on the orders of hundreds or thousands of kilometers in magnitude over multi-day forecast windows. As such, striking a balance between adversarial attack effectiveness and TLE usability represents a difficult, but not impossible, option for improving these attacks in the future.

In the short term, the best option for SSA operators who own stealth satellites may be simply not reporting them at all. This is far from ideal since, if a foreign state *does* detect an unlisted space object, its very omission can serve as an indicator of its nature.

Implications for SSA Recipients

For SSA-recipients who are contemplating the need for upgrades to domestic SSA capabilities, our research shows that even small investments in astrometry can be leveraged effectively when coupled with our anomaly detection model. States do not need to go toe-to-toe with the US military on SSA capabilities in order to impose effective constraints on the ability of third parties to abuse their trust in shared SSA. Likewise, SSA-sharing need not be treated as a zero-sum game. “Trust but verify” systems are possible which allow states to catch deception attempts while also permitting them to reap commercial and diplomatic benefits from continued international SSA cooperation.

Implications for Security Researchers

This experiment represents one of the first attempts to identify and simulate a credible threat to SSA data. We present the case that systems security thinking can bring valuable and novel solutions to cross-disciplinary problems in seemingly distant domains. However, our research also puts forward several questions that may be of interest in future work.

Our work deliberately prioritizes an attack with real-world empirical examples; however, there may be interest in anomaly-detection based defenses to previously unseen attacks on SSA data, such as the one presented in Section 9.3. Similarly,

research on attacks which target the CMOS camera systems or radar sensors used in SSA data collection, although quite expensive to demonstrate on realistic hardware, could answer important strategic questions regarding the capabilities of advanced persistent threat (APT) actors to harm space surveillance missions. Finally, a variant threat model which considers adversarial machine learning techniques vis-a-vis environmental trade-offs could be of significant utility to spy satellite operators.

9.5 Summary

In this chapter, we began with a broad recognition that a few key physical properties of space systems — such as their high velocities and low visibility — have had significant political and strategic implications. Specifically, we isolated three key features derived from these physical properties which have historically acted as barriers to ASAT weapons development and use: limited accessibility, attributable norms, and environmental interdependence. Likewise, these factors have driven trans-national cooperation on space surveillance and SSA data.

However, when connecting these dynamics with an adversarial cyber-security perspective, we found that many of the resultant trust relationships were intuitively exploitable as single points of failure in key information flows.

To better motivate and impact this theoretical security weakness, we posited two attack vectors which target SSA trust relationships. First, we proposed a model in which attackers approximate Cyber-ASAT capabilities by altering the orbital motion projects of RSOs to meet pre-defined attack criteria through minute data tampering attacks. Second, we proposed a simpler and historically proven attack where SSA originators seek to deceive other states as to the nature of their spy satellites. Both of these attack vectors were demonstrated in the form of orbital motion simulations which leverage real-world SSA data.

Finally, we presented initial work demonstrating how existing state-of-the-art techniques in classical machine learning and anomaly detection might be adapted to the space domain. Through careful feature-engineering of SSA data, we developed demonstrative machine learning models to detect SSA classification deceptions

without the use of a single telescope. These models were contextualized with respect to astrometry hardware requirements, showing that defenders can reasonably expect to use them to detect the vast majority of SSA classification attempts.

These results have broad direct implications. Our experiments show that even the most basic and essential information about an object's motion — less than 140 characters of TLE-data — can be reliably leveraged to detect disguised satellites. In the status quo, space powers are making critical decisions regarding investments in domestic SSA capabilities and next-generation stealthy nano-satellite platforms. Many of assumptions underpinning these policy actions are far from certain.

Our findings suggest that SSA data represents a likely battleground for future digital-mediated counterspace operations. The lack of “ground truth” information, heavy centralization of ephemerides, and the critical operational importance of these data streams makes them an attractive target for a wide range of attackers. To our knowledge, this research offers the first observation of the security-critical nature of SSA data integrity and the threat of deliberate SSA deception. Application of the RCMA method helped us to uncover and contribute solutions relevant to a novel and domain-specific problem area affecting satellite ground systems security.

*So that notwithstanding all these seeming impossibilities,
tis likely enough, that there may be a meanes invented of
journeying to the Moone; And how happy shall they be, that
are first succesfull in this attempt?*

—John Wilkins, *The Discovery of a World in the Moone*

10

Big Rockets, Small Satellites, and Cyber-Trust

The emergence of small, low-cost secondary satellite payloads, referred to as “CubeSats,” has underpinned a revolution in modern space mission design. This has, in turn, reshaped the satellite launch market. Where, in the past, rockets carried hardware belonging to a single nation-state or a handful of domestic organizations, today a single launch vehicle may take satellites belonging to dozens of foreign entities on a shared ride to the stars. In this chapter, we consider how these trends intersect with the evolving domain of space cyber-security.

We take an interdisciplinary approach, starting with an analysis of the global CubeSat launch market and relevant interstate political dynamics. This motivates a novel threat model, leveraging CubeSat payloads as cyber-physical attack vectors against launch operations. We isolate five key CubeSat safety standards which may constrain cyber-adversaries but find that most operate under trust assumptions which are vulnerable to malicious circumvention.

Rather than restricting ourselves to high-level strategic threat modeling, we cultivate a baseline intuition for the implications of such malicious safety violations through dynamic physical simulations of a space-to-space radio frequency interference (RFI) attack scenario. The results of these simulations suggest that,

Table 10.1: Demonstrative Mapping of RCMA Method to Adversarial Secondary Payload Research.

Research Step	Example Finding
Recognize physical aspects of space technologies which differ from comparable terrestrial systems.	High costs of orbital launch mean that many operators must share the same rockets in order to access orbit.
Connect these dynamics to their implications for traditional security approaches.	Secondary payloads can come from actors with disparate geopolitical interests and cyber-security norms. Their close physical proximity to the LV and other payloads gives rise to opportunities for violating trust which are not addressed by existing safety standards.
Motivate the need for security improvements by demonstrating these impacts in a domain-specific and realistic context.	Covert violations of CubeSat safety standards can leverage inexpensive COTS hardware in severe radio frequency interference attacks against flight systems.
Adapt proven security approaches to better account for these domain-specific requirements.	Modest revisions of safety standards to include third party verification of operational claims and adversarial models for risk assessment can significantly reduce attacker capabilities.

even limited to standard CubeSat components, attackers have wide physical margins within which to cause sustained intentional degradation to safety-critical communications during launch.

This research makes several contributions — presenting a novel analysis at the intersection between “launch diplomacy,” hardware safety, and cyber-security. It represents one of the first attempts to consider the cyber-security properties of space launches and, to our knowledge, the first publication to consider space-to-space cyber-warfare operations from secondary payloads as a threat vector. Methodologically, this study demonstrates how policy analysis, model-based engineering methods, and system security techniques can combine to provide cross-domain insights into emerging threats. Finally, the case study which makes up the latter portion of the chapter serves as a cautionary example of how safety engineering controls, which assume probabilistic physical effects as the main source of system failure, are not necessarily robust to intelligent and strategic adversaries.

Table 10.2: Example Per-Launch Costs and Capabilities of Modern LVs.

Vehicle	Approx. Launch Cost (USD)	Approx. Mass-to-Orbit (t)
Ariane 5 (ESA)	\$150 million	10 (GTO) - 20 (LEO)
Delta IV (NASA)	\$300 million	14 (GTO) - 29 (LEO)
Falcon 9 (SpaceX)	\$60-100 million	8 (GTO) - 23 (LEO)

NB: GTO = Geosynchronous Transfer Orbit, LEO = Low Earth Orbit

Throughout, this chapter follows the *RCMA* research method proposed in Chapter 8. A mapping of each of the four method stages to this study can be found in Table 10.1.

10.1 Background: The Practice (and Politics) of Sharing Rockets

Orbital access is expensive. Even with state-of-the-art technology, single rocket launches can exceed hundreds of millions of dollars (see Table 10.2). To overcome this barrier, satellite owners engage in ridesharing, purchasing excess capacity on someone else's launch vehicle (LV) for a secondary payload.

Ridesharing practice has co-evolved with a satellite design template, referred to as CubeSats [311]. CubeSats are small and lightweight, with the smallest size (1 CubeSat Unit or 1U) fitting inside a 10 cm³ area and weighing approximately 1.3 kg. For missions which require large components, multiple 1U cubes can be combined. For example, a 30x10x10 cm payload weighing around 4 kg would be referred to as a 3U CubeSat.

Compared to traditional satellites, CubeSats are small and cheap, with complete mission costs ranging from the tens of thousands to low millions of dollars [328]. Ready-made CubeSat platforms can be purchased online for as little as €25,000, although most missions will require some additional customization [329]. This has made CubeSats the platform of choice for many space startups and research missions.

The standard shape and mass of CubeSats allows for easy integration to LVs via standardized deployers, thereby creating a sort of commodity market for global CubeSat launch capacity. The dominant deployer type is the “P-Pod” (see Figure

I) [311]. A P-Pod is essentially an aluminum box with a door on one end and a spring on the other. When the door's latch is released by the LV's flight computer, the spring ejects up to 3U of CubeSats into space at 1-2 m/s velocity. Other deployer types tend to follow similar design principles [330, 331].

The global rocket launch market to deliver such payloads is consolidated into a handful of major players. Between 2016 and 2019, 90% of the estimated \$US 29 billion spent on launch services went to one of seven space powers: United States, European Union, China, Russia, Japan, India, and New Zealand [332]. The content of these missions, on the other hand, is highly internationalized. For example, the European Space Agency (ESA) Vega SSMS mission in 2020 delivered a total of 53 satellites to Low Earth Orbit (LEO) [333]. These included platforms for the Thai military, a Russian nuclear physics institute, an Estonian university, and a Facebook subsidiary. In total, 21 customers from 13 countries shared the same journey to the stars.

10.1.1 Space Diplomacy and Ridesharing

These multi-state missions occur against a complex geopolitical backdrop. LVs have been longstanding subjects of tension due to their dual-use potential; other than the direction they face, and the logo painted on their side, there is little differentiating an LV from an intercontinental ballistic missile (ICBM). Indeed, both the US and Russia regularly repurpose retired ICBMs for space launches, and responses to North Korea's domestic space program have been inextricably linked to arms-control concerns [334–336].

The tensions do not stop at the atmosphere's edge. Major military powers rely heavily on space for battlefield communications and operations. As satellites are physically fragile, there is significant fear of attacks on space assets in future conflicts [337]. As noted in Chapter 9, states have strong structural incentives to engage in cyber-attacks due to domain specific advantages compared to other counter-space operations.

However, there have also been many indications of interstate cooperation. Throughout the Cold War, significant efforts were made by both the US and USSR to cooperate on space launches, giving rise to the Apollo-Soyuz Test Project (ASTP). It has been argued that “track II diplomacy” resulting from interpersonal relationships cultivated during ASTP gave rise to broader diplomatic gains, such as strategic arms control agreements and the demilitarization of Russia’s launch sector [338]. In a more modern context, launch collaboration for the International Space Station (ISS) was one of the few aspects of the US-Russia bilateral relationship to survive the diplomatic fallout of Russia’s invasion of Crimea in 2014 [339].

Some classical realists treat this sort of cooperation with skepticism. For example, Wang’s review of US-EU space cooperation argues that the US used LV ride-sharing as a tool to undermine and weaken European rocketry development efforts [340]. Likewise, Chalecki contends that the ASTP was little more than a guise for US and Soviet military intelligence to spy on each other [341].

In short, satellite ridesharing is as much a geopolitical matter as a technical one. Ridesharing offers direct economic benefits, but it also redirects huge sums of money into foreign aerospace industries and provides political leverage to LV operators that maybe unpalatable to some satellite owners.

10.2 Threat Models for CubeSat Integration

In this context, we can surmise several motivations for cyber-attackers to target launches. A launch failure could prevent or delay the deployment of key space assets. Moreover, commercial actors may see benefit in harming the reputation of key competitors. For example, this was briefly investigated as a possible cause of a 2016 SpaceX rocket explosion [342]. The prestige and economic importance of space programs may also make them attractive targets for hostile states — as Russian officials suggested following a string of rocket failures in the early 2010s [343, 344].

For this case study, we focus on threats involving the compromise of an inexpensive CubeSat secondary payload. We propose four reasons CubeSats may represent attractive targets:

- Heavy use of commercial-off-the-shelf (COTS) components allows attackers to develop exploits on representative hardware or software. This contrasts with larger platforms which tend to rely on bespoke components.
- The COTS supply-chain can be compromised — for example, through a backdoor in an open-source software library or the online sale of a malicious sensor. The high number of CubeSats per LV increases the odds of a backdoored product ending up attached to an LV of interest to an attacker.
- While large satellites and LVs are typically built by nation-states and defense contractors, CubeSats frequently come from start-ups or universities. These organizations are comparatively permeable to digital compromise, insider threats, sabotage, and social engineering.
- CubeSats are inexpensive. Combined with the pseudo-commodity market for CubeSat launch slots, a proxy corporation or state-sponsored university could afford many attempts at building and launching a CubeSat with malicious flight software that abuses trusted/approved COTS components to cause harm.

To date, little prior technical research exists on CubeSat cyber-security — in large part due to their low capabilities and small size. To quote one CubeSat developer: “What’s the worst that could happen? [...] With no propulsion and no pointing control, it’s very likely that you couldn’t do anything other than turn the camera off” [345]. CubeSat manufacturers have lobbied against cyber-security standards, contending that they pose “an excessive and unnecessary burden, and a major potential mission-reliability risk” [345]. The effect of this mentality is that CubeSats tend to forgo security to meet aggressive cost and schedule requirements. Additionally, in a high-level review of CubeSat security practices, Ingols and Skowyra note that CubeSat developers will often “conflate reliability engineering with security engineering” [346].

This is an important point, because while security risks are frequently dismissed, attackers may still struggle to cause meaningful harm after successfully compromising a CubeSat. CubeSats represent many organizations’ first space mission and, as a

result, fail often. Roughly 50% of CubeSats suffer “infant mortality,” failing within six months, and one in five are “dead on arrival,” never making contact with Earth at all [347, 348]. Launch providers are thus keenly aware of the risks of strapping unreliable novice hardware, however small, to a cylinder full of rocket fuel. This has given rise to extensive controls designed to limit the mechanical and electrical risk a CubeSat can pose to the LV. In Section 10.3, we will consider these safety controls and their implications for an intelligent cyber-adversary.

10.3 Adversarial Analysis of Launch Safety Controls

CubeSat safety requirements can vary substantially and revolve around a series of mission-specific Interface Control Documents (ICDs) provided by the mission integrator. These requirements are complex and certification is non-trivial; NASA recommends 18 months of time for certification and licensing [349]. In this chapter, we focus on two dominant standard documents (among myriad) for CubeSat missions: the *CubeSat Design Specification, REV 13* (CDS) and the *Air Force Space Command Manual 91-710, Volume 3* (AFSPCMAN) [350, 351].

10.3.1 CubeSat Design Specification (CDS)

The CDS focuses mostly on the physical properties which may impact a CubeSat’s ability to deploy smoothly from a P-Pod. Beyond this, it imposes three broad categories of controls which appear to constrain cyber-adversaries.

First, CDS requires deployment switches, small pins on CubeSat rails which are depressed while the CubeSat sits in its P-Pod [350, Sec. 3.3]. These electrically isolate the CubeSat’s flight computer from power during launch to prevent a CubeSat from deploying hardware in the P-Pod. They also prevent attackers from launching software-based attacks prior to deployment. Second, CDS prohibits CubeSats from transmitting radio signals until 45 minutes have elapsed from deployment, although the CubeSat may boot up and perform other tasks in that time [350, Sec. 3.4]. This mitigates the risk of both unintentional and malicious radio frequency interference

(RFI). Third, CDS typically limits stored chemical energy to 100 Watt-Hours [350, Sec. 3.1]. This limits the available power for a cyber-attacker seeking direct physical effects — such as deliberate overheating of key components.

The controls are normally verified by three mechanisms [349]. Battery characteristics are outlined in a battery report which details specific part numbers and modifications. Radio and electrical interrupts are summarized in an electrical report containing circuit diagrams. Finally, inhibits are verified during a Day in the Life (DITL) test. In a DITL, the CubeSat runs through a simulated separation and a timer is used to verify that no premature transmissions take place. The DITL is typically conducted by the CubeSat developer in their own lab [352, 353].

10.3.2 Air Force Space Command Manual 91-710 (AFSPC-MAN)

AFSPCMAN consists of more than 200 pages of requirements for launch operations, the primary purpose of which is range safety. The objective of range safety is to protect persons, vehicles, and structures from harm and ensure that rockets adhere to intended trajectories. Range safety violations can result in the initiation of a self-destruction system, known as a Flight Termination System (FTS), which is designed to ensure that a launch vehicle combusts fully prior to colliding with the Earth's surface.

The primary AFSPCMAN burden for CubeSat developers is the provision of a Missile System Prelaunch Safety Package (MSPSP), prepared by the CubeSat developer [351, p. 214]. It consists of a detailed description, including schematics and functional diagrams, of the payload and relevant hazards.

The most obviously applicable portion of AFSPCMAN to cyber-security is the portion on “Computer Systems and Software” [351, p. 200]. Software security requirements are derived from Software Criticality Indexes (SwCIs) specified in MIL-STD-882E [354]. A synthesis of these requirements can be found in Table 10.3. In most cases, CubeSat software falls in the range of SwCI 4-5, with DITL testing meeting validation burdens. The only additional software safety hurdle

Table 10.3: Overview of AFSPCMAN Software Safety Standards.

Software Control Category	Severity Level of Safety Failure			
	Catastrophic (e.g., loss of life, >\$10M damages)	Critical (e.g., hospitalization of 3+ personnel, > \$1M damages)	Marginal (e.g., injury causing lost workdays, >\$100K damages)	Negligible (e.g. minor injury, <\$100K damages)
<i>Autonomous</i>	SwCI 1 (Code Review)	SwCI 1 (Code Review)	SwCI 3 (Architecture Review)	SwCI 4 (Safety - Specific Testing)
<i>Semi-Autonomous</i>	SwCI 1 (Code Review)	SwCI 2 (Design Review)	SwCI 3 (Architecture Review)	SwCI 4 (Safety - Specific Testing)
<i>Redundant Fault Tolerant</i>	SwCI 2 (Design Review)	SwCI 3 (Architecture Review)	SwCI 4 (Safety-Specific Testing)	SwCI 4 (Safety - Specific Testing)
<i>Influential / Informational</i>	SwCI 3 (Architecture Review)	SwCI 4 (Safety-Specific Testing)	SwCI 4 (Safety-Specific Testing)	SwCI 4 (Safety - Specific Testing)
<i>No Safety Impact</i>	SwCI 5 (No Analysis)	SwCI 5 (No Analysis)	SwCI 5 (No Analysis)	SwCI 5 (No Analysis)

Note: The controls in this table are synthesized from multiple tables in MIL-STD-882E and controls in AFSPCMAN 91-703v3 [351, 354]. All controls which apply to lower severity Software Criticality Indexes (SwCI) apply to high severity indexes cumulatively. For example, SwCI 1 software is subject to: Code Review, Design Review, Architecture Review, and Safety-Specific Testing.

imposed is likely a descriptive overview of computing hardware components and software logic [355].

Beyond software safety, the MSPSP also imposes requirements to mitigate the risk of electromagnetic interference. A CubeSat developer typically must provide a transmitter survey, which lists all radio transmitters and their fundamental characteristics. This includes an outline of frequency ranges, bandwidth, and deployed and maximum power delivery to a given antenna [349]. Range Safety may require verification of emission characteristics through measurements conducted by an approved representative [351, p. 43]. However, in practice, CubeSat missions may be able to avoid the costs and scrutiny of such assessment through the use of RF power inhibits which comply with CDS [355]. If frequency analysis is required, the main purpose is to ensure that payload emissions do not broadcast on key frequencies outlined in the LV's specification. These frequencies are often listed in public documentation and typically consist of telemetry and FTS modules [356, 357].

10.3.3 Adversarial Analysis

Initially, these controls appear to severely constrain an attacker's capabilities. However, their implementation assumes an informed and benign CubeSat developer

who shares the launch integrator's desire for a successful mission.

Under our adversarial model, this shared priority does not exist. CubeSat developers may be unaware of or complicit in efforts to circumvent controls. As large parts of the certification process are self-reported, violating controls is often little more than a matter of ticking an incorrect box or writing down inaccurate numbers on a form. Attackers can strategically evade only a small subset of the hundreds of standards, maximizing potential harm while minimizing detectability.

For example, Ingols and Skowyra note that CubeSats spend the months between completion and launch being passed around different storage facilities and may be subject to post-certification tampering via social engineering vectors [346]. A sophisticated attacker may make minor software modifications to devices during this time with little risk of detection. Even more severely, if the CubeSat developer misrepresents DITL results or electrical diagrams, there is no clear mechanism for detecting this; CubeSats are too fragile to tear apart for manual inspection.

In Table 10.4 we present a demonstrative analysis of five selected controls from CDS and AFSPCMAN under adversarial conditions. For each, we note the source of verification authority and deception exposure to both insiders and outsiders. Taken together, this analysis suggests that the primary techniques used for CubeSat safety — such as redundancy, certification, and documentation — provide only weak defense against malicious intent.

This analysis suggests that many of the controls which help ensure safety during the CubeSat integration process are not robust to an intelligent adversary. For example, requiring triple-redundant radio inhibits (CDS 3.4) dramatically reduces the risk from equipment failure. However, there is little difference from the perspective of a malicious CubeSat developer lying once in their electrical report versus lying thrice. Even absent insider access, the lack of software and supply-chain auditing processes provides ample opportunity for cyber-attackers to circumvent key safety requirements.

Table 10.4: Adversarial Circumvention Analysis for Selected Safety Controls.

Safety Control	Primary Reference	Responsible for Verification	Likely Vulnerability to Malicious Outsider	Likely Vulnerability to Malicious Insider
Deployment switches prevent power-on in deployer	CDS 3.3	CubeSat Developer (DITL, Electrical Diagrams)	Low <i>CubeSat developer would likely detect unauthorized power draw during DITL.</i>	High <i>CubeSat developer could forge documentation and DITL results.</i>
Software timers prevent RF transmission for 45 minutes	CDS 3.4	CubeSat Developer (DITL)	Moderate to High <i>Otherwise trivial modifications to code may necessitate special effort to evade DITL detection.</i>	High <i>CubeSat developer could forge DITL results or program DITL behavior to differ from launch.</i>
Battery power limitation	CDS 3.1	CubeSat Developer (Battery Report, MSPSP)	Low <i>Malicious vendor could misrepresent battery specs but targeting is logistically complex.</i>	Low to Moderate <i>Weight and physical properties act as limits on plausible extent of deception.</i>
Software Safety Guidance	AFSPCMAN A2.2.4.14	CubeSat Developer (MSPSP)	High <i>Software, especially third-party libraries, is unlikely to be audited beyond cursory summary in MSPSP.</i>	High <i>CubeSat developer will likely only need to provide easily falsified summary information on software operations and design.</i>
RF Emission Compatibility	AFSPCMAN A2.2.4.10.2, Launch Vehicle User's Guide	CubeSat Developer (MSPSP) Range Safety (EMF testing)	Low to Moderate <i>Malicious vendor could backdoor telemetry hardware. If a software defined radio (SDR) is used, attacker may modify configuration through code.</i>	Moderate to High <i>Absent independent EMF testing, CubeSat developer can lie. Otherwise, they may modify code to change behavior under test conditions.</i>

10.4 Attack Simulation and Evaluation

Given common perceptions that a CubeSat's low capabilities mean that, even in the event of full compromise, it cannot pose a physical threat, it is worth considering the specific technical implications of malicious safety control violations. To do this, we will replicate a hypothetical attack scenario through dynamic physical simulation. The intent is not to completely model the behavior of LVs and satellites but rather to evaluate the general plausibility of harm from compromised CubeSat hardware during launch.

Our hypothetical threat scenario focuses on GPS interference attacks for three

reasons. First, RFI attacks are intuitively bolstered by physical proximity — one of the main boons from compromising a secondary payload. Second, what limited public information is available on LV FTS hardware makes it clear that GPS is a key data source [358]. Finally, due to US commercial radio licensing regulations, there is a relative abundance of technical data regarding representative radio hardware, helping to better ground our simulations [359].

10.4.1 Scenario Overview

The compromised CubeSat in our simulation is summarized in Table IV. It consists of a notional 3U commercial payload, weighing 4 kg and scheduled for launch on a SpaceX Falcon Heavy. The mission sequence is loosely modeled on that of the STP-2 launch. STP-2 is selected as an example of a mission which deployed CubeSats en route to delivery of the primary payload. This emerging practice offers commercial and logistical benefits, but also raises the risks from compromise as CubeSats are deployed while the primary payload and substantial fuel quantities remain in the LV.

Our attacker is derived from the insider model in the rightmost column of Table 10.4. They are a malicious state sponsored business who has built a CubeSat with the express purpose of circumventing key safety controls. To reduce scrutiny, the attacker is restricted to standard CubeSat components. There are two relevant hardware modules used in the attack, both belonging to the CubeSat's Telemetry, Tracking and Control (TT&C) subsystem.

First, the CubeSat leverages a software defined radio (SDR) transceiver. Specifically, we have modeled our simulation around the 1U μ SDR-C from Space Micro [360]. An SDR permits the attacker to dynamically alter radio transmission parameters, including carrier frequencies, using undisclosed software logic. SDRs are commonly used in CubeSats and the presence of an on-board SDR alone would be unlikely to arouse suspicion. Additionally, the attacker has selected an antenna with undisclosed operability in the 1.1-1.6 GHz range as well as the allocated TT&C band. This can be achieved with a customized deployable antenna, a multi-band module, or

an ultra-wideband offering [361–363]. This frequency range is selected due to its potential to cause interference with GPS reception.

The attacker has also inserted malicious programming logic with the intention of circumventing two safety controls from Table 10.4. First, the attacker will begin RF transmission immediately after separation from the P-Pod, violating the 45-minute silence mandate. Second, the attacker will transmit on frequencies prohibited by AFSPCMAN A2.2.4.10.2 and the Falcon User’s Guide [356]. To evade detection during lab certification and DITL tests, this malicious logic will check the measurements of on-board sensors (e.g., a thermometer) and only trigger the attack when conditions match LEO.

The attacker’s goal is to introduce radio-frequency interference (RFI) of sufficient magnitude to trigger a range safety incident on the LV. For example, if positional telemetry data is unavailable or indicates a rocket has strayed from its intended trajectory, this can lead to a mission abort.

This is particularly relevant for the Falcon Heavy, as it is one of the first LVs to include a fully autonomous flight termination system (AFTS) [356, p. 8]. This AFTS can automatically self-destruct the launch vehicle without human approval if sensors show deviation from approved mission parameters. Although the precise AFTS specifications are, unsurprisingly, restricted, NASA documents confirm GPS observations as a key decision metric for termination [358].

10.4.2 Experimental Design and Assumptions

The primary purpose of these simulations is to determine the plausible limits of CubeSat hardware to emit RF which causes sustained degradation to GPS reception. In practice, many relevant dynamics are mission-dependent, such as antenna directionality, GPS satellite locations, and precise launch trajectories. Here, we focus on a “worst case” scenario based on typical GPS signal characteristics, idealized isotropic antennas, and assumption of equivalent receiver gain across legitimate and illegitimate transmission sources.

Model Parameters

According to the Falcon User's Guide, the launch vehicle contains GPS receivers which operate in the L1 signal band (1574.2 MHz) [356]. To determine the necessary jammer characteristics to cause disruption to these signals, we must approximate the strength of legitimate signals at the receiver. The GPS specification only provides information regarding the Earth's surface, but we can derive a more accurate value for LEO. One method for doing so is presented in [364], suggesting an approximate received power of around -120 dBm in dynamic simulation. This is fairly close to the value predicted by a simple Free Space Path Loss (FSPL) model on the basis of the public GPS L1 link budget — with minor modification to account for LEO conditions (see Equations 10.1, 10.2, and 10.3) [365].

Letting:

$$FSPL(\text{dB}) = -10 * \log_{10} \left[\left(\frac{4\pi d}{\lambda} \right)^2 \right] \quad (10.1)$$

$$P_{rcvr}(\text{dBm}) = EIRP_{tx} + FSPL - 30 \quad (10.2)$$

Where:

d = distance from transmitter $\approx 19,000$ meters (depends on orbit/time)

λ = wavelength ≈ 0.19 meters

$EIRP$ = effective isotropic radiated power ≈ 26.5 dBW

$$\begin{aligned} P_{rcvr}(\text{dBm}) &= 26.5 - 10 * \log_{10} \left[\left(\frac{4\pi 19000}{0.19} \right)^2 \right] - 30 \\ &= -125.48 \text{ dBm} \end{aligned} \quad (10.3)$$

We can supplement this theoretical analysis with experimental data from the US Department of Transportation (DOT) [366]. Through anionic chamber measurements evaluating the threat of interference from cellular LTE towers (at 1530 MHz) on LEO GPS reception, DOT calculated a receiver threshold of -73 dBm for near-band interference on two NASA platforms [366, p. 110]. As our attacker can jam directly in the L1 band, rather than the adjacent LTE frequencies, we can reasonably assume equivalent or greater interference at this threshold.

Simulation Process

Our physical simulation consists of two sub-components — an astrodynamics model for CubeSat separation and an RF interference model. In the astrodynamics model, we replicate the separation of a CubeSat from a P-Pod deployer into LEO. This is implemented in FreeFlyer, a commercial space mission planning tool [298]. The CubeSat ejects from the launch vehicle through a contra-velocity maneuver at 2 m/s as is typical for a 4 kg CubeSat [367]. The CubeSat and launch vehicle are propagated for a 2-hour period following separation, and a separation vector is calculated between the two objects at regular one-minute intervals.

These separation vectors are then leveraged in RF interference simulations. We replicate RF dynamics using MATLAB's Antenna Toolbox, a commercial communications system simulation and development toolkit [368]. Two transmitters are modeled: an L1 GPS transmitter based on the aforementioned P_{rcvr} characteristics and a CubeSat jammer with varying EIRPs from 1-10W. A GPS receiver is replicated on board the rocket. Antenna positions are derived based on the separation vectors calculated in the astrodynamics model and used to compute signal-to-interference-plus-noise ratios (SINR) and $P_{jammer \text{ at } rcvr} (dBm)$ at regular one-minute intervals.

Under benign conditions ($P_{jammer \text{ at } rcvr} (dBm) = 0$), our model computes: $P_{gps \text{ at } rcvr} (dBm) = -125.48$, and $SINR(dB) = -21.41$. These values align with our analysis in Section 10.4.2 and prior work, suggesting reasonable fidelity [364, 369].

10.4.3 Results and Evaluation

Figure 10.1 summarizes the output of our astrodynamics model. Note that the separation vector of magnitude does not increase linearly. This is a result of the relative orbital motion of the CubeSat and LV, both of which are in LEO at time of deployment. In our threat model, the attacker does not adhere to the 45-minute radio silence window mandated by the CDS. This means that they can jam immediately after separation and at close proximity to the LV.

Incorporating these results into the interference model shows that the attacker is capable of degrading GPS signal quality (see Figure 10.2a). As expected, the





Figure 10.1: CubeSat Separation from LV over Time.



(a) SINR at LV Receiver During Attack

(b) SINR at LV Receiver During Attack

Figure 10.2: Attack Simulation Results.

attack is most effective at high levels and during the first few minutes follow separation. Using the aforementioned DOT near-band threshold of -73 dBm gives a conservative estimate of between 20-40 minutes of disruption depending on the attacker's amplifier power (see Figure 10.2b).

To further validate these bounds, we can convert SINR to Carrier-to-Noise-plus-Interference Density ratio C/N_{O+I} , assuming a typical front-end bandwidth (BW) of 4e6 Hz and applying the conversion method presented in [369] and in Equation 10.4.

$$C/N_{O+I} (dB - Hz) = SINR + 10 * \log_{10} (BW) \quad (10.4)$$

Standard GPS L1 receivers function at C/N_{Os} between 35-55 dB-Hz, with complete loss of signal acquisition below 28 dB-Hz — although this can vary

depending on specific hardware conditions [370]. This suggests that our attacker can have a severe impact on GPS quality, keeping C/N_{O+I} below 28 dB-Hz for upwards of 45 minutes at low EIRPs and throughout the simulated period at higher EIRPs (SINR ≤ -38 dB). It may be prudent to assume that GPS receivers on LVs have access to the wider 20.46 MHz P(Y) frequencies restricted for military use. If this were the case, an attacker would be weaker, but could still expect between 30-60 minutes of successful disruption (SINR ≤ -45 dB).

In short, these results suggest it is physically plausible for COTS CubeSat hardware to introduce meaningful disruptions to LV GPS reception on the scale of tens of minutes to several hours depending on mission hardware. While operationalizing such an attack would take significant effort, the low cost and accessibility of CubeSat hardware and launch capacity make it well within the means of state sponsored attackers. Moreover, the reputational risk of attack failure or attribution is limited as key forensic evidence of the attack would be trapped 1,000 m in the sky.

10.4.4 Mitigations and Future Work

The scenario considered here is but one of many possible manifestations of our threat model. The underlying vulnerability proposed here has less to do with GPS reception than with the implicit trust dynamics in secondary payload integration. One promising avenue for future work might thus be to build on this adversarial analysis to identify other technical attack vectors of interest (e.g., premature hardware deployment to jam P-Pod deployers).

Our own RFI scenario also leaves room for future work. Due to limited public information, we could not account for the specific AFTS design. AFTS systems may already have a variety of undocumented defenses, such as leveraging multi-constellation GNSS data, elevating the importance of accelerometer readings in the case of GNSS anomalies, or employing various jamming resistance techniques [371]. To the extent that such mitigations are not implemented, they also represent feasible technical steps towards mitigating the attacks proposed here.

At a high level, our research suggests ample opportunity for future work in adjusting trust models around CubeSat integration policies. This is complex as CubeSats are built under aggressive timeline and budgetary constraints. However, certain properties — such as the validity of hardware interrupts, operational frequencies of RF hardware, or behavior during DITL testing — may be of sufficient importance to merit the added cost of third-party validation. Launch operators may consider offering expedited certification routes for certain pre-approved COTS components, such as antennas which lack capabilities in sensitive frequencies, to reduce compliance costs. Similarly, they may consider allowing developers to gain trust over time, easing the pathways to large-scale CubeSat deployments while still mitigating the risks from naive or fraudulent first-time developers.

In short, a comprehensive review of the existing integration certification process from an adversarial perspective is beyond the scope of this thesis but represents an intuitive next step for launch operators and regulators concerned about potential harms from compromised or malicious third-party payloads.

10.5 Summary

In this chapter, we have presented the case that strong political and strategic motivations exist for attacks targeting space launch missions. Moreover, we present, to our knowledge, the first cyber-physical threat model targeting LVs through a secondary payload.

While existing CubeSat safety standards employed in the integration and certification process initially appear to constrain cyber-adversaries, we find that unverified trust assumptions underpin the real-world practice of this safety qualification process. When considered in the context of a sufficiently motivated malicious cyber-adversary, many safety protections appear trivially circumventable.

The implications of this are evaluated experimentally through physical simulations of a novel space-to-space radio interference attack scenario targeting a modern LV. Our results demonstrate that inexpensive CubeSat hardware has sufficient physical capabilities to potentially threaten the reliability of key safety metrics

during launch. We further considered how future work might identify related attacks against other launch systems and isolated steps towards mitigating both this specific attack and others of this nature.

We have shown how operational and physical adaptations made by satellite operators have created security tensions in the launch safety domain. In a perfect world, satellite operators might remain in full control of an exclusive launch vehicle to deploy their systems — limiting exposure to threats during this critical mission stage. In practice, this is only financially viable for large nation-state operators willing to expend hundreds of millions of dollars on launch operations.

By applying the RCMA method, we identified novel threat models that take advantage of this physical reality to undermine the security of space missions. Additionally, we presented key improvements which may be made by adapting safety regulations to an adversarial context. In doing so, we have demonstrated how RCMA might be applied to cross-disciplinary research which touches on both the operational policy and technical dimensions of space missions.

For hundreds of satellite operators, transnational launch collaboration has brought space closer than it has ever been. It offers access for start-ups, universities, and states who would otherwise be unable to reach orbit. Moreover, it fosters key links for communication and diplomacy between scientists and engineers in otherwise deeply sensitive domains. However, trust is a keystone component of sustained cooperation. Ensuring security against both cyber and physical risks will be critical to reaping sustained benefits from globalized launch services.

Part IV

Conclusion

The answers, when they came, were vague and noncommittal. In a fury he whipped the machine [...], crying “The truth now, out with it, you blasted old digital computer!”

—Stanisław Lem, *Fables for Robots* (trans. Kandel)

11

Conclusion: What’s Next in Space Cybersecurity?

Contents

11.1 Research Summary	242
11.2 Summary of Key Contributions	245
11.2.1 Methodological Contributions	245
11.2.2 Direct Contributions	246
11.3 Future Work in Space Cyber-Security	248
11.4 Final Remarks	251

We began with a declaration: *space is changing*. In the time it took to conduct the research contained in this thesis, the number of operational satellites in orbit has more than doubled [245]. It now appears near-certain that the coming decade will see more satellites launched than the previous six combined.

Diving headfirst into this next era of human spaceflight, we can no longer afford to rely on “security through obscurity” as the primary defense against cyber-exploitation of space systems. The field sits at a critical inflection point. Design decisions made in today’s system will shape the safety and security of orbit for years to come.

The primary objective for this thesis was to meet an emergent need for open, evidenced, and technical research on space systems security. Throughout this

work, we have demonstrated how status quo practices (often enshrouded in layers of governmental and commercial classification) have struggled to keep abreast of growing attacker capabilities and system complexity. The security of many keystone aspects of modern space missions have never been assessed. The contributions in this thesis represent a tangible first step towards rectifying some of these shortcomings.

11.1 Research Summary

The historical analysis at the core of Chapter 2 laid the initial foundations for our own research effort. Through our systematization of a sparse and disparate body of existing work on satellite cyber-security coupled with our own archival research, we learned that satellites have been the target of digitally mediated attacks from both state and non-state actors for decades. We further proposed a taxonomy of satellite security questions into four sub-domains: radio communications security, ground systems security, space platform security, and mission operations security. In each domain, we synthesized cross-disciplinary research from fields as diverse as international relations and aerospace engineering to isolate key threats and unsolved security challenges. Over the course of this thesis, we touched on each of these sub-domains, with a particular focus on the first three, and demonstrated how systems security research may contribute to their development.

In Part II, we took a close look at the communications domain and modern satellite broadband services. Through real-world experiments, we demonstrated that services from geostationary orbit (GEO) were leaking sensitive data belonging to a wide range of customers, including some of the world's largest corporations and critical infrastructure providers. Moreover, we demonstrated novel attack vectors that might give cyber-adversaries the ability to alter these communications or otherwise cause active disruption to networked systems.

Investigating these issues further, we determined that their underlying causes were intrinsically connected with the physical characteristics of long-range satellite communications. Security practices used by terrestrial internet customers, such as end-to-end VPN encryption, were not designed for compatibility with




Figure 11.1: The RCMA method for applied space systems security research (repeated from Chapter 1).

the traffic optimizations used in satellite broadband services. As a result, we observed cross-industry failures to adopt otherwise long-established best practices for communications security due to the perceived performance costs of doing so.

To remediate this situation, we proposed and implemented a novel VPN-hybrid which combined traditional VPN security models with the performance optimizations required for satellite networking. Further, we developed a replicable method for evaluating both our system and future proposals through an open-source simulation testbed. In assessing our proposal, we demonstrated its ability not only to obviate historical security/performance trade-offs in the domain, but to even bring modest performance improvements over existing *insecure* protocols.

While Part II makes up the bulk of this thesis and a substantial contribution in its own right to the security of satellite broadband services, we aspired to broader contributions as well. By stepping back and reflecting on the research approach we used to explore satellite broadband security, we hypothesized its applicability to other

dynamics in space mission security. The result of this reflection was the four-step RCMA method, detailed initially in Chapter 1 and replicated here in Figure 11.1.

Rather than end this thesis with a theoretical assertion of the utility of this approach, in Part III we tested its viability through experimental analysis in two topics unrelated to broadband communications.

The first of these looked at space situational awareness (SSA) data used by satellite ground systems for navigational and monitoring purposes. We found that physical astrometric challenges in resident space object (RSO) monitoring had led to heavy centralization of navigational data in two military-operated databases (one belonging the United States and the other to Russia). The result was that many state and non-state operators were left in the unpalatable position of reliance on the goodwill of foreign militaries for critical data. We demonstrated how attackers, or repository owners themselves, may tamper with this shared data to deceive recipients and proposed a number of political and strategic motivations they may have for doing so. Finally, we adapted state-of-the-art anomaly detection approaches used in systems security to demonstrate how even low-resourced space situational awareness data recipients might detect such deceptions.

The second study related to modern rocket launches which reduce overall costs of orbital access by integrating dozens of satellites on the same launch vehicle (LV). In analyzing the safety requirements applied to these rideshare payloads, we identified a number of regulations which appeared trivially circumventable by a motivated cyber-adversary who compromised an inexpensive or malicious secondary payload. Through dynamic simulation models, supplemented with real-world data from related work, we demonstrated how attackers might abuse this circumvention to cause intentional radio frequency interference during rocket launches with potentially severe implications for the safety and security of the overall mission. Finally, we identified small modifications which could be made that would dramatically reduce the ability of attackers to leverage these sorts of exploits in the future.

In both analyses, our identification of key physical challenges in astrometry and rocketry allowed us to build threat models around trust relationships which had

evolved in response to these engineering hurdles. In considering these threat models in the context of operational risk for modern space missions, we demonstrated the need for improvements over status quo security practices and the importance of adversarial modeling approaches to space mission design. Finally, we drew on techniques which had been proven in other information security domains and adapted these strategies to the unique requirements of satellite operators.

Together, these studies suggest that the issues identified in Part II are but a small portion of the unexplored research topics and security challenges related to space systems security. They demonstrate that the unique physical aspects of space systems give rise to a wide range of novel technical effects with direct implications for cyber-security. Further, they show how our physically-grounded RCMA method can help uncover topics of interest in the domain well beyond those directly considered in this thesis.

11.2 Summary of Key Contributions

This thesis makes a number of contributions. The substantive technical research which comprises Parts II and III identifies a number of previously unknown or under-studied threats to modern satellite missions. Each of these threats are verified experimentally or through simulation, and novel defenses are proposed to address problems relevant to hundreds or thousands of status-quo space missions. Meanwhile, this thesis also offers contributions in support of future research, systematizing existing knowledge and presenting methods for studying an historically inhospitable topic.

11.2.1 Methodological Contributions

The broadest methodological contribution of this thesis is the RCMA (Recognize, Connect, Motivate, Adapt) method for cyber-physical space security research. This research process is designed to help identify a sequence of relevant questions that support space systems security researchers. Some of the stages, such as *Recognize* and *Motivate*, benefit from aerospace engineering perspectives and methods. Their

inclusion helps shape interdisciplinary security contributions into operational and practical relevance for real-world space missions. Other stages, such as *Connect* and *Adapt*, leverage traditional systems security techniques to evidence hypothetical threat models or adapt proven defensive tactics to newly identified risks. Although this research method is hardly a radical break from any other hypothesis testing model, its conscious decision to interweave cross-discipline perspectives is useful in its ability to uncover and address topics that may be missed or under-served by pure aerospace or computer science methods.

Beyond this general research process, this thesis also contributes methodological guidance for others interested in studying this domain. In Chapter 2, we outline the lack of prior technical work on space systems security and many of the financial and logistical barriers which have prevented security contributions in the past. Through the technical studies in this thesis, we demonstrate how consumer-grade hardware (Chapters 4-5), dynamic computer simulation (Chapters 6, 9, & 10), and public data-sets (Chapter 9) can be combined into a low-cost evidentiary foundation for meaningful technical research on space systems security. In doing so, we offer a model for others unsure of how to start contributing to this topic in the face of restricted information and limited technical accessibility. These techniques and the pitfalls we encountered in grappling with them throughout this thesis offer a starting point for future work on space systems security.

11.2.2 Direct Contributions

In addition to the broader methodological contributions, each chapter in this thesis included a number of direct technical outputs relating to threat models in satellite systems security. The most developed of these contributions appeared in Part II, where the relative accessibility and affordability of representative hardware allowed for in-depth study of real-world satellite broadband networks. Meanwhile, the most exotic appeared in Part III, where we considered topics that have received little to no prior research attention due to various economic and logistical barriers.

In the context of satellite broadband security, we identified a real-world threat impacting tens of thousands of satellite broadband customers and millions who rely on the services those customers provide. Our threat model revealed previously unknown risks to critical sectors, ranging from companies collectively responsible for more than 40% of the world's cargo shipping capacity to critical infrastructure providers of electricity. A direct contribution of this research has been heightened awareness across the satellite broadband market of the potential for long-range eavesdropping threats. This thesis has directly led to emergency threat intelligence notifications issued through the United States Federal Bureau of Investigation (FBI) and to substantial coverage by popular press outlets such as *Forbes* and *Ars Technica*. The end result is that both satellite ISPs and customers are more aware than ever before of the importance of robust information security practices and the need for encryption in satellite communications.

Achieving this substantial change in threat model was made possible through a reframing of prior assumptions regarding attacker capabilities. By developing *GSExtract* in Chapter 5, we demonstrated how intelligent forensic reconstruction of corrupted data captures could allow attackers to achieve reliable eavesdropping capabilities with inexpensive consumer grade hardware, decreasing the cost of executing such attacks by several orders of magnitude. This packet reconstruction technique may be of use for other protocols, both in satellite communications and in any domain where receiver equipment costs represent an assumed defense against eavesdropping threats. Following an extended responsible disclosure campaign, *GSExtract* itself is now publicly available on GitHub for other researchers interested in adapting its techniques and contributions.

Additionally, this thesis has contributed new tactics for defending against these long-range eavesdropping attack vectors. In Chapter 6, we demonstrated how existing security techniques might be combined with satellite performance enhancing proxies to create a secure optimization tunnel for satellite broadband. The resultant tool, *QPEP*, is now publicly available for others to modify and benchmark. Moreover, we developed a replicable testbed environment for conducting such benchmarks

and detailed techniques which leverage that testbed to make apples-to-apples comparisons between security protocols for these networks.

Beyond satellite broadband security, we have devised several previously unconsidered attacks targeting space situational awareness (SSA) data in satellite ground stations and threat models which leverage these attacks to achieve various objectives relevant to both state and non-state actors. Through a combination of simulations and real-world data, we have further developed and demonstrated techniques which satellite operators may use to bolster the trustworthiness of their SSA data feeds without relying on prohibitively expensive and inaccessible astrometry equipment. The end result of this contribution is that both SSA operators and data-dependents can have a deeper understanding of implicit trust assumptions in information sharing, the risks that can result from violations of those assumptions, and tactics which may be deployed in practice to detect or dissuade such violations.

Finally, we have presented what is, to our knowledge, the first public technical model of a space-to-space cyber-physical threat and the first public technical model of cyber-mediated threats to space launch missions. In the course of reviewing CubeSat integration safety standards, we identify a number of controls which are robust against physical misfortune but not robust against intelligent adversarial circumvention. The result of this contribution is a characterization of key weak-points in the launch integration process which can be bolstered through minor policy and operational changes to better ensure the safety and security of shared launch operations. The ultimate benefit of this contribution is that satellite launch operators can have better visibility into a novel category of risk which is largely overlooked under otherwise extensive status-quo risk assessment frameworks.

11.3 Future Work in Space Cyber-Security

As a core objective of this thesis was to lay the groundwork for future research in an historically neglected topic, it is unsurprising that there remain many avenues for future work. For example, we have proposed several demonstrative research topics in each sub-domain included in Chapter 2, only a handful of which are

explored in this thesis. It is also worth noting that each chapter of this thesis outlines several narrower, specific improvement opportunities which logically follow from its core contributions. However, this section will highlight some of the more notable directions for future work from across this thesis:

LEO Broadband Security As this concluding chapter was being written, SpaceX onboarded its first public beta testers for access to the next-generation *Starlink* constellation. Unlike the internet services considered in Part II, these next generation constellations aspire to provide ubiquitous broadband services from Low Earth Orbit (LEO). While GEO broadband is likely to remain relevant for decades to come — especially for industrial customers and IOT applications — these LEO constellations bring with them entire new categories of security concerns and opportunities. In proposed mega-constellations, latency is dictated more by routing decisions than transmission distances, potentially allowing for the use of traditional VPNs or reducing the importance of performance enhancing proxies (PEPs). However, mega-constellations also require significantly more complex routing protocols as satellites move along their orbital paths. Securely balancing authentication, encryption, routing, and performance in an ever-shifting mesh network of thousands of nodes is a complex task at potentially unprecedented scale. As these constellations move from theory to practice, and more details emerge regarding their design and capabilities, there will be need for focused security research to ensure that the protocols employed adequately protect the security and privacy of their customers.

Alternate Communications Systems Our research has focused on one specific subset of satellite data services: broadband from DVB-S based telecommunications providers. However, there are numerous other satellite data applications, such as the Inmarsat BGAN protocol used for international maritime communications, that may be vulnerable to related threat models. Understanding how these proprietary and esoteric services interact with adversarial threats, and how they may be best defended, is a challenging but potentially fruitful area for further work.

Satellite Control Protocols There is significant interest in ensuring the security of satellite telemetry, tracking, and command (TT&C) protocols. Due to the physical visibility of satellites, vulnerabilities in TT&C systems may be used by attackers to hijack, disable, or otherwise harm satellites from vast distances. Today, research on this topic is difficult to conduct responsibly due to the high costs of space systems and low incentives for industry participants to collaborate with academic or independent security researchers on these topics. However, as the nature of space missions continues to evolve and COTS hardware becomes more available to researchers, work on secure TT&C systems represents a promising direction.

QPEP Validation and Development As a result of the global coronavirus pandemic and subsequent restrictions on lab access and travel, we elected to evaluate QPEP through network simulations. The simulation approach offers many advantages, such as the ability to dynamically alter network conditions and to provide replicable benchmarking scripts in support of future research. However, developing and evaluating QPEP as an overlay on real-world networks and, in particular, tailoring QPEP to multi-user environments for ISP deployment represents a short-term and direct evolution of the contributions in this thesis. In the longer term, QPEP's proof of concept may be bolstered through the addition of novel protocol modifications, such as QUIC-layer forward error correction (FEC) to improve performance in lossy networks.

SSA Decentralization In the status quo, Space Situational Awareness (SSA) data is highly centralized within the apparatus of the United States Space Surveillance Network (SSN) and a handful of smaller state-operated repositories. In this thesis, we have shown how this centralization leads to potentially exploitable trust dynamics. As space becomes both more crowded and, potentially, more multi-polar, there may be interest in revising these relationships. One strategy for doing so might be cost-sharing between many small state and non-state operators, allowing them to combine limited astrometry budgets into the foundations of a full SSA catalog. While such a system offers some clear political and technical benefits, it also raises

its own security challenges. Future research might consider techniques by which parties who do not trust each other can credibly share data on the motion and location of resident space objects (RSOs).

Payload Security Certification As shown in Chapter 10, one of the key reliability and safety checkpoints in satellite development occurs at the payload integration stage. The current system has few substantive information security controls, especially in the case of smaller secondary payloads. In this thesis, we focused on existing safety controls and their susceptibility to an adversarial threat model. However, the payload integration and certification stage may also represent an ideal opportunity for broader cyber-security verification of space systems. Future work determining which tests and criteria might be incorporated into existing integration standards to bolster the security of space systems could be of significant value to stakeholders across the space community. That said, the payload certification process is already arduous and costly. Determining appropriate tradeoffs between mission capabilities, costs, and security and distilling these into actionable policy recommendations would represent a significant future contribution to space systems security.

11.4 Final Remarks

The possibilities of outer space are nothing short of magical. Whether in the form of protecting life from natural disasters, enabling global transport and logistics systems, or bolstering our understanding of the universe, space touches the lives of billions. Our collective relationship with both outer space and cyberspace is still in its infancy, and the decisions we make now will shape it for decades to come. As we step forwards into the next era of space exploration and development, ensuring that these systems remain secure and trustworthy will be critical to reaping the benefits of these technologies.

Over the course of this thesis, we have seen how status quo practices fall short of this goal. We have isolated clear gaps between potential threats to these systems

and their defenses against cyber-exploitation. Moreover, we have contributed to the reduction of this gap through our consideration of several unsolved problems specific to the domain. Whether in the context of satellite communications, space surveillance, or rocket launches, this thesis offers new insights on the nature of cyber-risk for modern space missions and techniques for its mitigation.

This is just a first step. There remain hundreds, if not thousands, of unstudied security relationships in modern space missions. Moreover, the industry is entering a period of rapid change which will surely raise still more novel security challenges. Throughout this work, we have discussed techniques and methods that might be of use in exploring these topics. Sustained research will be vital to the responsible development of orbit, and this thesis offers itself as a launchpad for others ready to explore the intersection of outer space and cyberspace.

Part V

Appendices and References

Appendices



Chronology of Significant Satellite Hacking Incidents

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
1962	Jamming	Government	Hypothetical	Commercial	United States
In a 1962 congressional hearing on the first American commercial satellite company, the prospect of signal jamming and potential satellite hijacking was suggested as a possible threat to low-altitude satellite missions. <i>Primary/Contemporary References:</i> [45]					
1972	Jamming	Government	Soviet Union	Multiple	Multiple
A UN proposal by the Soviet Union is raised suggesting that states have an intrinsic right to jam satellite signals in their territories via technical means. <i>Primary/Contemporary References:</i> [46]					
1986	Signal Hijacking	Insider	United States	Commercial	United States
An industry insider injected video and audio into an HBO television broadcast in Florida. Interestingly, this attack may have been inspired by a fictional article which appeared the previous year in a satellite television enthusiast magazine about an individual who hijacked HBO signals in protest of new scrambling policies. <i>Primary/Contemporary References:</i> [47, 372] <i>Secondary References:</i> [2, 24]					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
1986	Jamming	Government	United States	Ground (Accidental)	United States
In 1986 a garage door company discovered that communications satellites which were directed towards Reagan's vacation home in California were jamming terrestrial garage door openers more than 200 miles away. <i>Primary/Contemporary References:</i> [373]					
1986	Eavesdropping	Government	Indonesia	Commercial	United States
In 1986 the government of Indonesia was accused by an American satellite imaging firm of using large satellite receivers to intercept earth observation images without subscribing to the service. <i>Primary/Contemporary References:</i> [50]					
1987	Signal Hijacking	Individual	United States	Commercial	United States
In 1987, Thomas Haynie, an employee of the Christian Broadcasting Network, hijacked satellite transmissions from the Playboy Channel and replaced them with static text from the bible. <i>Primary/Contemporary References:</i> [49] <i>Secondary References:</i> [2]					
1987	Groundstation	Individual	Germany	Gov. Military	United States
In 1987, a group of youths in West Germany managed to compromise top secret networks belonging to NASA and other major space agencies. These networks provided at least the ability to find secret information about space missions and potentially information which could have compromised these missions. <i>Primary/Contemporary References:</i> [51]					
1993	Cryptographic	Individual	United Kingdom	Commercial	United Kingdom
A group of hackers distributed BSkyB satellite channels through a decoder-card sharing scheme across an apartment complex. <i>Primary/Contemporary References:</i> [56]					
1993	Jamming	Government	Indonesia	Commercial	Tonga
Satellite operators from Indonesia and Tonga threatened each other over the proposed Tonga Gorizont 17 satellite in GEO above New Guinea. The orbital slot was under contention due to potential interference as both states threatened to jam the other's transmissions from the orbit. This is the first public record of a state threatening digital counterspace operations against another state's assets.					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Primary/Contemporary References: [374]</i>					
1994	Jamming	Government	Iran	Commercial	Multiple
The Iranian government was suspected of jamming foreign television programs from Arab-Sat and Asia-Sat platforms during ongoing debate over banning the domestic use of satellite dishes altogether. <i>Primary/Contemporary References: [52]</i>					
1994	Cryptographic	Individual	United States	Commercial	United States
Gregory Manzer was sentenced on charges of creating and distributing technology to break the VideoCipher encryption technology used by HBO and ESPN satellite channels. <i>Primary/Contemporary References: [57]</i>					
1996	Jamming	Government	Indonesia	Commercial	Hong Kong
In 1996 the Indonesian government used a communications satellite called Palapa B1 to jam signals from a Hong Kong/British satellite Apstar-1A which was leased by Tonga - making good on threats from three years prior. <i>Primary/Contemporary References: N/A Secondary References: [2, 3, 54]</i>					
1996	Jamming	Government	Turkey	Commercial	United States
The Turkish Government is believed to have jammed broadcasts originating from MED-TV, a Kurdish nationalist satellite television station operating on an Eutelsat satellite. The jamming campaign continued sporadically between 1996 and 1999. Turkish authorities claimed MED-TV was "Terrorist Television" and incited acts of violence. British authorities ultimately terminated the MED-TV transponder license in 1999. <i>Primary/Contemporary References: [375, 376] Secondary References: [54, 377]</i>					
1997	Groundstation	Government	Russia	Gov. Scientific	United States
Hackers in 1997 successfully compromised Goddard Space Flight Center computers capable of satellite command and control. Later investigation linked this incident to Russia-government associated hackers, although full verification of this claim cannot be made without access to classified investigation reports. <i>Primary/Contemporary References: [32] Secondary References: [2]</i>					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
1998	Payload Damage A cyber-intrusion at Goddard Space Flight Center possibly caused the German-US ROSAT telescope to face the sun and burn its optical sensors. <i>Primary/Contemporary References:</i> [32] <i>Secondary References:</i> [33]	Unknown	Unknown	Gov. Scientific	United States
1998	Jamming A Moscow-based company began selling a \$4000 portable jammer capable of disabling GPS signals over a 200km radius. <i>Primary/Contemporary References:</i> [55]	Commercial	Russia	Multiple	Multiple
1998	Groundstation Individual In 1998, a hacker group called “Masters of Downloading” claimed to have stolen classified software that provided sensitive information and limited control over military satellites include GPS systems. The pentagon acknowledged a minor breach but contended that the hackers exaggerated their capabilities. <i>Primary/Contemporary References:</i> [378] <i>Secondary References:</i> [2]	Individual	United States	Gov. Military	United States
1999	TT&C In 1999 hackers claimed to have hijacked a British military satellite’s control systems and to have demanded ransom from the British government. However, the British military strongly disputed these claims. <i>Primary/Contemporary References:</i> [379] <i>Secondary References:</i> [2]	Individual	Unknown	Gov. Military	United Kingdom
1999	Jamming Russian government admits jamming satellite phone networks in Chechnya to prevent communications among separatists. <i>Primary/Contemporary References:</i> [36]	Government	Russia	Commercial	Russia
1999	Groundstation Individual A teenager going by the handle “cOmrade” plead guilty to charges of compromising NASA computer systems that support the International Space Station. The intrusions occurred in 1999. <i>Primary/Contemporary References:</i> [128]	Individual	United States	Gov. Manned	Multiple
2000	Jamming During a 2000 tank competition to demonstrate tanks for sale to the Greek military, French forces used ground-based GPS jammers to cause navigation problems during US and British entries.	Government	France	Navigational	Multiple

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Secondary References: [2, 380]</i>					
2000	Jamming	Government	Iran	Commercial	France
Iran was accused of using jamming devices impacting Turkish territory to interfere with Eutelsat-based opposition broadcasts. <i>Primary/Contemporary References: [63]</i>					
2000	Groundstation	Individual	United States	Gov. Scientific	United States
Jason Dikeman was charged in 2000 of gaining unauthorized access to systems which control NASA satellites. <i>Primary/Contemporary References: [381]</i>					
2000	Groundstation	Unknown	Unknown	Gov. Military	United States
In 2000 unknown hackers stole software from a US defense contracting company which enables ground stations to send commands to satellites. <i>Primary/Contemporary References: [127]</i>					
2001	Groundstation	Individual	United Kingdom	Gov. Scientific	United States
UK based hacker Gary McKinnon was indicted on charges of compromising 16 NASA computer systems. McKinnon claimed to have been looking for evidence of a cover-up relating to extra-terrestrial intelligence and unidentified flying objects. While there is not evidence that McKinnon compromised systems related to satellite control, it represents an early high-profile attack against a space agency with the intent of stealing space-mission data. Subsequent coverage has focused on matters of extradition and human rights for cyber-crime. <i>Primary/Contemporary References: [382] Secondary References: [3, 383]</i>					
2002	Groundstation	Individual	Venezuela	Gov. Scientific	United States
A Venezuelan hacker using the pseudonym “RaFa” provided a reporter at Computer World copies of a PowerPoint documents detailing the design of NASA launch vehicle Cobra and other sensitive engineering information. Later, Rafael Aponte was sentenced and extradited for compromises and defacement of US military information systems conducted under the same pseudonym but charges for the NASA compromise were never pressed. Some sources associate this compromise with the 2002 Marshal Space Flight Center Intrusions, although this attribution is disputed. <i>Primary/Contemporary References: [384] Secondary References: [32]</i>					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2002	Signal Hijacking Falun Gong transmitted protest videos from Taipei over official Chinese Central Television satellite broadcasts. <i>Primary/Contemporary References:</i> [58]	Dissident	Taiwan	Gov. Media	China
2002	Groundstation An attacker compromised computers at Marshall Space Flight Center, stealing intellectual property related to launch vehicle design. This attack has since been tenuously attributed to China, although more contemporaneous sources associate it with the RaFa intrusions. <i>Primary/Contemporary References:</i> [384] <i>Secondary References:</i> [2, 32]	Government	China	Gov. Scientific	United States
2002	Jamming Several sources assert that in 2002, a poorly installed CCTV camera in the town of Douglas, Isle of Mann caused interference with GPS signals over a 1 km area. We were unable to find a primary source for this claim, but it is a commonly referenced example of accidental GPS interference. <i>Secondary References:</i> [3, 385, 386]	Individual	United Kingdom	Navigational	United States
2002	Eavesdropping John Locker, a satellite eavesdropper, reported the ability to intercept images from NATO surveillance aircraft. NATO respondents claimed that the images did not contain sensitive information, but media reports claimed that they revealed sensitive details regarding the capabilities and location of classified vehicles. <i>Primary/Contemporary References:</i> [387] <i>Secondary References:</i> [3]	Individual	United Kingdom	Gov. Military	United States
2003	Signal Jamming US government broadcasts in favor of regime change in Iran were jammed by attacks on the Telstar-12 satellites by the Iranian government. <i>Primary/Contemporary References:</i> [388] <i>Secondary References:</i> [389]	Government	Iran	Commercial	United States

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2003	Jamming	Government	Cuba	Commercial	United States
In 2003 the Cuban government was accused of deliberately jamming US signals for the Voice of America station which were being broadcast to Iran, perhaps on behalf of the Iranian government and with communications gear supplied by China. This has been associated with the Iranian jamming of the Telstar-12 incident by some secondary sources. <i>Primary/Contemporary References:</i> [64] <i>Secondary References:</i> [3]					
2003	Signal Hijacking	Dissident	Taiwan	Commercial	China
Falun Gong again transmitted protest media across an AsiaSat transponder in 2003 to interrupt CCTV coverage of the Zhenzhou V space mission. <i>Primary/Contemporary References:</i> [59]					
2004	Signal Hijacking	Dissident	Taiwan	Commercial	China
Falun Gong again transmitted protest media across an AsiaSat transponder. <i>Primary/Contemporary References:</i> [60]					
2005	Groundstation	Government	China	Gov. Manned	United States
Windows malware installed in Kennedy space center's vehicle assembly building sent information about the space shuttle to computers in Taiwan. While this may have been espionage to mimic shuttle technology, investigators also believe information that could threaten the shuttle was exfiltrated. Weak attribution to the PLA has been made. <i>Secondary References:</i> [2, 32]					
2005	Jamming	Government	Libya	Commercial	United States
In 2005 the Libyan government was accused of jamming telecommunications satellites which impacted both European television stations and government communications. <i>Primary/Contemporary References:</i> [390] <i>Secondary References:</i> [2]					
2006	Jamming	Government	Libya	Commercial	United Arab Emirates
In 2006 the Libyan government was accused of jamming satellite telephone frequencies in order to combat the use of satphones by smugglers.					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Primary/Contemporary References: [391] Secondary References: [2]</i>					
2006	Signal Hijacking	Government	Israel	Commercial	Lebanon
Israeli forces in 2006 hijacked the Hezbollah-associated satellite television channels to air threatening anti-Hezbollah messages. <i>Primary/Contemporary References: [392] Secondary References: [2]</i>					
2006	Groundstation	Unknown	Unknown	Gov. Scientific	United States
A purported 2006 phishing incident targeting NASA employees led to the leak of NASA budgetary documents detailing satellite investment priorities. We were unable to find primary source information regarding this breach, but several prior surveys have cited it as example of IP theft attacks. <i>Secondary References: [2, 3, 32]</i>					
2006	Sensor Disruption	Government	China	Gov. Military	United States
China beamed a ground-based laser at sensors on a US spy satellite. Very little information about the incident and its effects is public. <i>Primary/Contemporary References: [35]</i>					
2007	Signal Hijacking	Terrorist	Sri Lanka	Commercial	United States
Tamil rebels may have hijacked an Intelsat satellite signal to broadcast propaganda. The rebels claim they had purchased access to the satellite but Intelsat disputes this. The incident went on for more than 2 years. <i>Primary/Contemporary References: [61] Secondary References: [41]</i>					
2007	Groundstation	Government	China	Gov. Scientific	United States
The ground station analysis process for Earth Observation Data at Goddard Space Flight center was compromised by attackers believed to be associated with the Chinese state according to secondary sources. No primary source coverage of this incident could be found, but it is cited in several surveys as an instance of state sponsored espionage. <i>Secondary References: [2, 3, 32]</i>					
2007	TT&C	Government	China	Gov. Scientific	United States
Two NASA satellites in 2007 and 2008 suffered major disruption attacks. Initial reporting suggested that these were just jamming attacks but later reports suggest ground station control takeover and accuse China. <i>Primary/Contemporary References: [65] Secondary References: [33, 41]</i>					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2008	Payload Damage An astronaut is believed to have introduced a virus to ISS windows-XP computers by bringing a compromised laptop on board. More recent reports suggest the virus was brought aboard by Russian cosmonauts, but it is unlikely to have been done deliberately. <i>Primary/Contemporary References:</i> [40] <i>Secondary References:</i> [33, 66]	Insider	Russia	Gov. Manned	Multiple
2008	Groundstation Attackers were reported as having used a Trojan horse installed on devices at NASA's Johnson Space Center to compromise communications to the international space station and disrupt some services on-board. It is unclear if the attack was targeted or coincidental. <i>Secondary References:</i> [2, 3, 30]	Unknown	Unknown	Gov. Manned	United States
2009	Groundstation In March 2009, an Italian hacker compromised several NASA systems, including systems used to control NASA's Deep Space Network and control systems in Goddard Space Flight Center. NASA claims that no critical harm was posed to space missions. <i>Primary/Contemporary References:</i> [42]	Individual	Italy	Gov. Scientific	United States
2009	Eavesdropping Iraqi insurgents intercepted unencrypted video streams via satellite links using a commercial software product called SkyGrabber. <i>Primary/Contemporary References:</i> [62] <i>Secondary References:</i> [41]	Terrorist	Iraq	Gov. Military	United States
2009	Eavesdropping A 2009 Blackhat presentation demonstrates the ability to intercept live video feeds from DVB-S signals, including sensitive military and media feeds by modifying existing satellite hardware. <i>Primary/Contemporary References:</i> [137]	Researcher	United Kingdom	Commercial	United Kingdom
2009	Signal Hijacking In 2009 almost 40 individuals in Brazil were arrested on charges of hijacking UHF frequencies belonging to US Naval satellites for personal usage. UHF transponder hijacking is believed to be widely used by criminal organizations and individuals seeking free long-range communications services in remote parts of the country.	Individual	Brazil	Gov. Military	United States

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Primary/Contemporary References: [393] Secondary References: [2]</i>					
2009	Jamming	Government	Egypt	Commercial	United Kingdom
In May of 2009 the Al-Hiwar satellite station broadcast from the United Kingdom was jammed. No culprit has been conclusively identified but the Egyptian government is strongly suspected. <i>Primary/Contemporary References: [68]</i>					
2010	Groundstation	Individual	China	Gov. Scientific	United States
A Chinese hacker was arrested on charges of stealing export-controlled data from NASA computer systems. The hacker was arrested by Chinese authorities with supporting evidence provided by the United States. It represents one of the first cooperative law enforcement actions regarding government systems compromise between the two states. <i>Primary/Contemporary References: [42]</i>					
2010	Eavesdropping	Researcher	Spain	Commercial	Spain
A 2010 Blackhat presentation demonstrates the ability to intercept live internet feeds from DVB-S signals using general purpose equipment <i>Primary/Contemporary References: [82]</i>					
2010	Groundstation	Accidental	United States	Navigational	United States
In 2010 an Air Force update to GPS ground control stations resulted in multi-day outages effecting as many as 10,000 military GPS devices. <i>Primary/Contemporary References: [394] Secondary References: [33]</i>					
2010	Jamming	Government	Iran	Commercial	France
A series of jamming incidents around the 31st anniversary of the Islamic Revolution in Iran jammed broadcasts from international satellite television channels on a Eutelsat satellite. The Iranian government is suspected of instigating the attacks. <i>Primary/Contemporary References: [69] Secondary References: [2]</i>					
2010	Jamming	Government	Jordan	Commercial	United Arab Emirates
In 2010 Jordan was accused of jamming Al-Jazeera satellite television feeds, including some which broadcast the World Cup. <i>Primary/Contemporary References: [70]</i>					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2010	Jamming	Government	North Korea	Navigational	South Korea
North Korea has attempted to disrupt South Korean GPS navigational signals through jamming attacks starting in 2010 and continuing thereafter. <i>Primary/Contemporary References:</i> [395]					
2010	Groundstation	Government	United States	Gov. Scientific	United States
An Office of the Inspector General for NASA audit found that e-waste systems prepared for resale relating to the Space Shuttle missions retained sensitive data which was not correctly deleted, including export controlled information. Similar sensitive information was found on hard drives in dumpster outside a NASA facility. <i>Primary/Contemporary References:</i> [42]					
2011	Groundstation	Government	United States	Gov. Scientific	United States
The Office of the Inspector General for NASA issued a report indicating that critical vulnerabilities were found in at least six systems which could be used by a remote attacker to control or debilitate ongoing satellite missions. <i>Primary/Contemporary References:</i> [42]					
2011	Groundstation	Unknown	China	Gov. Scientific	United States
Attackers in 2011 gained administrative control of computer systems in the NASA Jet Propulsion Laboratory using previously stolen credentials. The attack was later attributed to China. <i>Primary/Contemporary References:</i> [90] <i>Secondary References:</i> [14]					
2011	Jamming	Government	Bahrain	Commercial	France
A Bahraini opposition station called LuaLua TV was jammed within 5 hours of its first broadcast over a Eutelsat transponder, likely by the Bahraini government. <i>Primary/Contemporary References:</i> [71] <i>Secondary References:</i> [2]					
2011	Jamming	Government	Ethiopia	Commercial	United States
Ethiopian Satellite Television - an anti-regime satellite television channel - was jammed by the Ethiopian government in 2010 (and several times thereafter). Some have suggested that the equipment and technology for these attacks was provided by Chinese government officials. <i>Primary/Contemporary References:</i> [72] <i>Secondary References:</i> [2]					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2011	Jamming	Government	Libya	Commercial	United Arab Emirates
In 2011 the Libyan government again jammed satellite telephone frequencies in order to combat the use of satphones by smugglers. <i>Primary/Contemporary References:</i> [73] <i>Secondary References:</i> [2]					
2011	Jamming	Government	Saudi Arabia	Commercial	Iran
Iran has accused Saudi Arabia of jamming its state run satellite television networks starting in 2011. <i>Primary/Contemporary References:</i> [74]					
2011	Groundstation	Unknown	Unknown	Gov. Scientific	United States
In 2011 a laptop containing command and control algorithms used for the operation of the International Space Station was stolen. The laptop was unencrypted, but it is unclear if the attacker specifically targeted NASA information. <i>Primary/Contemporary References:</i> [42] <i>Secondary References:</i> [3]					
2011	Groundstation	Unknown	China	Gov. Scientific	United States
Chinese hackers are suspected of having compromised accounts of privileged users at the Jet Propulsion Laboratory which provided attackers with full access to devices on the network. <i>Primary/Contemporary References:</i> [42]					
2012	Groundstation	Individual	Romania	Gov. Scientific	United States
In February 2012, NASA's Inspector General pressed charges against a Romanian national for intrusions into Jet Propulsion Laboratory computer systems to steal information regarding a scientific sensor for space missions. <i>Primary/Contemporary References:</i> [42]					
2012	Groundstation	Individual	Romania	Gov. Scientific	United States
In January 2012, the Romanian government arrested a 20-year-old hacker who had compromised both NASA and Romanian government information systems. Other than a low-impact denial of service, this had no lasting repercussions. <i>Primary/Contemporary References:</i> [42]					
2012	Jamming	Government	Eritrea	Commercial	France
An Eritrean opposition satellite radio channel called Radio Erena was jammed by the Eritrean government in 2012.					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Primary/Contemporary References: [75]</i>					
2012	Jamming	Government	Ethiopia	Commercial	United Arab Emirates
The Ethiopian government is suspected of jamming Eritrean satellite communications signals on ARABSAT platforms starting in 2012 (and several times thereafter).					
<i>Primary/Contemporary References: [76] Secondary References: [2]</i>					
2012	Jamming	Government	Syria	Commercial	France
Eutelsat was targeted by jamming signals believed to originate in Syria.					
<i>Primary/Contemporary References: [77]</i>					
2012	Jamming	Government	North Korea	Gov. Mili- tary	South Korea
North Korea is believed to have jammed South Korean military communications satellites starting in 2012.					
<i>Primary/Contemporary References: [80]</i>					
2012	Cryptographic	Researcher	Germany	Multiple	Multiple
German researchers published a paper detailing the ability to decrypt voice communications over many satellite phones implementing the common GMR-1 and GMR-2 encryption algorithms.					
<i>Primary/Contemporary References: [84]</i>					
2013	Spoofing	Researcher	United States	Navigational	United States
University of Texas at Austin researchers demonstrated the ability to leverage GPS spoofing to redirect an \$80 million yacht remotely.					
<i>Primary/Contemporary References: [396] Secondary References: [33]</i>					
2013	Jamming	Government	Azerbaijan	Commercial	Turkey
The Azerbaijani government was found by the USA to be deliberately jamming opposition satellite television stations on Turksat platforms.					
<i>Primary/Contemporary References: [78]</i>					
2013	Jamming	Government	Egypt	Commercial	Qatar
The Egyptian government was accused of jamming Al Jazeera satellite broadcasts during instability in 2013.					
<i>Primary/Contemporary References: [397]</i>					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2013	Jamming	Individual	United States	Navigational	United States
A limousine driver in New Jersey had installed a GPS jammer in his vehicle to prevent his employer from tracking the vehicle. The jammer caused interference with navigational systems at a nearby airport. <i>Primary/Contemporary References:</i> [398] <i>Secondary References:</i> [3]					
2014	Groundstation Audit		United States	Gov. Scientific	United States
US department of commerce office of the inspector general found more than 9,000 high risk issues in the Joint Polar Satellite System (NOAA) ground stations <i>Primary/Contemporary References:</i> [91]					
2014	Groundstation Researcher		United States	Commercial	United States
A presentation at Defcon in 2014 found severe vulnerabilities - such as hard-coded passcodes) in 10 SATCOM terminals. Some of these are remotely exploitable but many require physical or at least logical access to the devices. <i>Primary/Contemporary References:</i> [93]					
2014	Jamming	Dissident	Thailand	Commercial	Thailand
Thailand government television stations were repeatedly jammed in 2014 during a series of government protestors. No culprit was identified but it is believed to have been the protestors. <i>Primary/Contemporary References:</i> [399]					
2014	Jamming	Unknown	Egypt	Commercial	Saudi Arabia
In 2014 a comedy broadcast in Egypt was deliberately jammed with interference from two stations in Cairo. It is unclear who is responsible. <i>Primary/Contemporary References:</i> [400]					
2014	Jamming	Government	Libya	Commercial	United Arab Emirates
Libya is believed to have jammed a dozen channels by Dubai-headquartered MBC. <i>Primary/Contemporary References:</i> [401]					
2014	Groundstation	Government	China	Gov. Military	Germany
Hackers are accused of having compromised computer systems at DLR with spyware that may have been able to implicate the security of critical space missions and missile technologies. Initial attribution suggests Chinese attackers, but the evidence is uncertain.					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Primary/Contemporary References: [402] Secondary References: [3]</i>					
2014	Eavesdropping	Government	Russia	Multiple	Multiple
The Russian satellite Lurch, launched in 2014, is suspected of hovering close to other communications satellites in order to intercept signals <i>Primary/Contemporary References: [86]</i>					
2014	Groundstation	Government	China	Gov. Mili- tary	United States
In 2014 CrowdStrike released a report indicating that Chinese government-affiliated hackers targeted information about satellite control systems and successfully compromised some sensitive space networks. Few additional details are available. <i>Primary/Contemporary References: [89]</i>					
2014	Signal Hijacking	Dissident	Palestine	Commercial	Israel
Hamis briefly successfully compromised Israeli Channel 10 satellite television broadcasts and transmitted a message threatening Gaza residents. <i>Primary/Contemporary References: [39]</i>					
2014	Groundstation	Government	China	Government - Weather	United States
Chinese hackers, believed to be associated with the Chinese government, compromised a sensitive network related to NOAA weather satellites and caused a brief network outage during incident response. It is unclear what systems were compromised or what ability the hackers had. <i>Primary/Contemporary References: [88]</i>					
2014	Jamming	Government	Ethiopia	Commercial	Saudi Arabia
Television broadcasts from the ARABSAT platform were jammed by an attacker in Ethiopia, potentially associated with the Ethiopian state which has a history of similar jamming attacks targeting Eritrean broadcasts. However, some sources have conjectured that the incident was accidental as ARABSAT does not broadcast to either country. <i>Primary/Contemporary References: [403] Secondary References: [3]</i>					
2015	Signal Injection	Criminal	Russia	Commercial	Multiple
Russian-government affiliated group Turla was found to use satellite internet signals to exfiltrate data from malware infections with minimum traceability. Evidence of this method was found in malware dating back to 2007.					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
<i>Primary/Contemporary References: [94]</i>					
2015	Eavesdropping	Researcher	United States	Commercial	United States
	2015 Blackhat demonstration indicated practical ability to spoof devices on the Globalstar network and intercept simplex data messages intended for other devices. Globalstar contended that they simply provide hardware and that encryption was the job of their clients based on mission need. <i>Primary/Contemporary References: [83]</i>				
2015	Eavesdropping	Researcher	Germany	Commercial	United States
	Security researchers demonstrated the ability to intercept and interpret communications over the Iridium LEO network using a software defined radio. <i>Primary/Contemporary References: [404] Secondary References: [3]</i>				
2015	Groundstation	Individual	United Kingdom	Gov. Military	United States
	A British individual was arrested on charges related to compromising pentagon satellite communications systems. The hacker posted threats online claiming to have the ability to “control” satellites but the Pentagon has not confirmed the extent of the intrusion. <i>Primary/Contemporary References: [405]</i>				
2016	Jamming	Government	North Korea	Navigational	South Korea
	In April 2016, North Korea resumed the 2012 (and occasionally thereafter) jamming campaign against South Korean GPS signals. Russia is suspected (but not proven) to have provided the jamming equipment. <i>Primary/Contemporary References: [406]</i>				
2016	Jamming	Government	Russia	Commercial	Ukraine
	Media Group Ukraine’s broadcast of a 2016 football match was targeted by a malicious jamming attack. No attribution for the attack has been made but Russia is highly suspected. <i>Primary/Contemporary References: [407]</i>				
2016	Signal Hijacking	Dissident	Palestine	Commercial	Israel
	Hamas again compromised satellite transmissions, this time of the popular Israeli TV show Big Brother, and replaced them with propaganda films. <i>Primary/Contemporary References: [81]</i>				

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2016	Signal Hijacking An individual hacker or group of hackers in Saudi Arabia hijacked Israeli news satellite feeds in protest of an Israeli bill restricting the volume of calls to prayer (muezzin bill). They replaced the media feed with the call to prayer and text threatening punishment from God. <i>Primary/Contemporary References:</i> [408]	Individual	Saudi Arabia	Commercial	Israel
2016	Cryptographic Cryptographic researchers in China present a realtime attack against the GMR-2 encryption algorithms used by many satellite phones, updated prior research from Germany. <i>Primary/Contemporary References:</i> [409]	Researcher	China	Commercial	Multiple
2017	Groundstation The Chinese Thrip espionage group was found by Symantec in 2017 to have attempted to infect computers which monitor and control satellites. <i>Primary/Contemporary References:</i> [87]	Government	China	Commercial	United States
2017	Groundstation A French security researcher on twitter claimed to have compromised a Cobham VSAT terminal on a naval vessel over the internet using a default username and password combination. <i>Primary/Contemporary References:</i> [410]	Researcher	France	Commercial	United Kingdom
2018	Jamming Israel is suspected of having initiated a jamming attack against Syrian satellite television stations in retaliation for an attack on an Israeli jet flying over Syrian territory. <i>Primary/Contemporary References:</i> [79]	Government	Israel	Commercial	Syria
2018	Jamming Russia is accused of having jammed GPS signals across Norway and Finland to disrupt ongoing NATO war games in the region. The jamming attacks also impacted commercial aviation systems. <i>Primary/Contemporary References:</i> [411]	Government	Russia	Navigational	NATO

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2018	Groundstation	Unknown	Unknown	Gov. Scientific	United States
A Raspberry Pi microcomputer attached to Jet Propulsion Laboratory systems was compromised and used by attackers to further access other JPL systems, including systems which control the Deep Space Network radio systems and systems which might allow for malicious control of ongoing space missions. <i>Primary/Contemporary References:</i> [31, 43]					
2018	Groundstation	Unknown	Unknown	Gov. Scientific	United States
An advanced persistent threat attacker was found to have compromised Jet Propulsion Laboratory mission networks and to have maintained access to the systems for nearly a year prior to detection in April 2018. They would have had the capability to disable critical space communications systems and were found to have exfiltrated export regulated and sensitive information. <i>Primary/Contemporary References:</i> [31]					
2018	Groundstation	Researcher	United States	Multiple	United States
An updated version of IOActive research presented in Blackhat 2014 found that VSAT stations could be used to find GPS coordinates of military installations and potentially weaponized to cause interference. The attacks again focused on VSAT terminal firmware. <i>Primary/Contemporary References:</i> [92]					
2018	Groundstation	Researcher	Germany	Commercial	Unknown
A security researcher demonstrated the ability to compromise maritime satellite terminals over the internet and use them to send NMEA messages to cause harm to operational technology aboard yachts at sea. <i>Primary/Contemporary References:</i> [412]					
2019	Groundstation	Unknown	Unknown	Gov. Scientific	United States
A zero day attack in specialized satellite operations software was compromised on a server belonging to the Jet Propulsion Laboratory. Attackers had the ability to upload control instructions to the spacecraft. <i>Primary/Contemporary References:</i> [31]					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2019	Groundstation	Government	United States	Gov. Military	Iran
Media sources have asserted that a series of unexpected failures of Iranian rocket launches in 2019 were the result of a US-government sabotage effort involving either supply-chain compromises or cyber-attacks on Iranian launch vehicles. No official confirmation of these conjectures has been made by either state. <i>Primary/Contemporary References:</i> [413, 414]					
2019	Groundstation	Government	Iran	Gov. Military	United Arab Emirates
Media sources have asserted that an unexpected launch failure of an military spy satellite belonging to the UAE may have arisen from an Iranian cyber-attack. No official confirmation of these conjectures have been made by either state. <i>Primary/Contemporary References:</i> [414]					
2019	Misc.	Individual	United States	Commercial	United States
US Astronaut Anne McClain was accused of illicitly accessing the bank account of a former partner whilst living aboard the International Space Station. This was widely reported as the first crime accusation against a person in orbit. However the case has not been resolved and, in 2020, the partner who made the accusations was indicted on charges of making false allegations to law enforcement. <i>Primary/Contemporary References:</i> [415, 416]					
2019	Spoofing	Researcher	United States	Navigational	United States
A security researcher at Black Hat USA 2019 demonstrated a series of GPS spoofing attacks against an autonomous vehicle. The researcher was able to cause the vehicle to drive off the road by spoofing measurements of its current location. <i>Primary/Contemporary References:</i> [417]					
2019	Groundstation	Government	North Korea	Gov. Scientific	India
A North Korean attributed malware, dubbed “Dtrack” was reported to have been found on computer systems belonging to the Indian Space Research Organisation. Little information regarding the result of the compromise is publicly available, although some media sources surmise it may relate to the concurrent failure of the Chandrayaan 2 Lunar Lander. No evidence of this claim has been provided. <i>Primary/Contemporary References:</i> [418]					

Year	Attack Type	Attacker Type	Attacker Country	Victim Type	Victim Country
2020	Eavesdropping	Researcher	United Kingdom	Commercial	United States
A security researcher at Black Hat and DEFCON presented research demonstrating that satellite broadband signals could be intercepted by eavesdroppers using inexpensive home-television equipment. They further demonstrated that this impacted the security and privacy of terrestrial, maritime, and aviation customers. The research was based around some prior academic publications. <i>Primary/Contemporary References:</i> [20, 37, 114, 115]					
2020	Misc.	Individual	United States	Gov. Military	United States
The US Air Force and Defense Digital Service hosted “Hackasat.” A series of satellite hacking related events and competitions with the goal of increasing technical exposure to satellite cyber-security. The final challenge of the competition was to upload a mission plan to a live satellite (after exploiting a series of vulnerabilities in a ground-based system meant to replicate a satellite) and take a “cyber moon-shot” photograph of the moon using the satellite’s onboard camera. <i>Primary/Contemporary References:</i> [95]					

B

GSExtract Implementation Details

This appendix details the general approach used by the GSExtract tool to parse corrupted and incomplete raw DVB-S2 streams containing GSE data feeds. Several novel strategies are used to simplify the data extraction challenge to its core dimensions and reduce the complexity of an otherwise highly variable set of protocol standards. This appendix is intended to provide technical insight into the techniques employed which may be of academic interest. An open-source release of the GSExtract source code has also been made available following a year-long responsible disclosure effort.¹ A simplified overview of the entire GSExtract data extraction process can be found in Figure B.3 at the end of this section.

The first step in parsing raw transponder streams is to extract individual baseband frames (BBFrames). BBFrames are the lowest-level logical encapsulation layer inside a demodulated DVB-S2 stream. Each BBFrame begins with a 10-byte BBHEADER as defined by ETSI EN 302 307 and summarized in Figure B.1 [178]. The most important portion of this header for our purposes is the two-byte Data Field Length (DFL) value, which indicates the overall size of the data-field which follows the BBHEADER and the location in the stream where the next BBFrame begins. Additionally, the final byte of the BBHEADER contains a CRC-8 (using the polynomial represented by 0xD5) which protects the BBHEADER from corruption.

¹GSExtract source code can be found at <https://github.com/ssloxford/gsextract>.

In theory, the first two bytes of the BBHEADER, collectively referred to as the MATYPE may change arbitrarily in an Adaptive Coding and Modulation (ACM) feed. However, in practice, we observed that such changes occurred only rarely and that, in particular, the second byte of the MATYPE header in marine VSAT implementations was almost always 0x00. This observation acted as a ‘crib’ which significantly simplified GSExtract’s identification of corrupted and invalid BBFrames.

MATYPE 1	MATYPE 2	User Packet Length	Data Field Length	Sync	Syncd	CRC-8
1B	1B	2B	2B	1B	2B	1B

Figure B.1: The structure of a DVB-S2 BBFrame Header.

To extract BBFrames, GSExtract first attempts to identify a valid 2-byte MATYPE in the recorded raw DVB-S2 stream. This value can be identified through statistical analysis of a portion of the DVB-S2 recording where it will appear as one of the most frequently recurring 2-byte sequences. To validate this identification, the CRC-8 value in byte 10 of the BBHEADER can be used to confirm whether a 9-byte sequence beginning with the MATYPE is a plausible BBHEADER.

Once the correct MATYPE has been found, it is possible to parse the stream into complete BBFrames using the DFL BBHEADER value. Generally, the next BBFrame will begin immediately after the end of the previous BBFrame’s data field. However, in the case of signal processing errors, the data field may be truncated. In order to resolve issues with lower-end equipment, we thus validate each subsequent BBFrame header by checking its MATYPE against the known good value. While this decreases compatibility with some complex implementations that may use multiple MATYPEs in a single stream, for the maritime VSAT operators that we observed, these protocol features did not appear to be in use. This approach to error recovery ensures that no more than 2 BBFrames worth of data are lost as the result of a single signal processing failure. In most cases, only a fraction of a single corrupted frame is discarded before GSExtract recovers its synchronization with the DVB-S2 stream.




Figure B.2: An overview of the GSE header format. Only the first two bytes are required but the header can be of arbitrary length depending on the addition of optional extensions.

Next, we extract bytes up to the length indicated by the DFL BBHEADER value, less four bytes at the end of the BBFrame. While these four trailing bytes are not mentioned in the relevant DVB-S2 specifications, both service operators analyzed appeared to reserve these four bytes for a CRC-32 checksum calculated across the entire BBFrame.

Next, the contents of an extracted BBFrame are further parsed into GSE packets. GSE packets follow a format specified in ETSI TS 102 606 and outlined in Figure B.2 [170]. Each GSE packet has a variable-length header of at least 2-bytes which includes a 12-bit integer indicating the overall length of the GSE packet. Unlike indicated in the GSE standard, we found that for both maritime VSAT operators this value was the length of the entire GSE packet rather than the number of bytes which followed the mandatory 2-byte header. An arbitrary number of additional optional headers exist depending on the type of GSE packet encoded. Of particular importance are the headers related to GSE fragmentation. When a payload exceeds the maximum size of an individual GSE packet, it may be fragmented across several. This fragmentation process uses a 1-byte identifier to label related fragments and the entirety of any given fragmented payload must be completed within 255 BBFrames. GSEExtract attempts to deal with fragmentation by combining related fragments as they are identified. In a case where all fragments

are not successfully identified with a range of 255 BBFrames (e.g., due to signal corruption), GSEextract will attempt to recover partially completed GSE packets by padding the remaining bytes of the GSE payloads with null values (0x00). Individual GSE packets cannot traverse multiple BBFrames. Thus, if the final GSE packet inside a BBFrame appears truncated, this can be taken as an indication of signal processing error and the broken GSE packet will be discarded by GSEextract.

Finally, within the payloads of either complete GSE packets or re-assembled fragmented payloads, GSEextract will attempt to parse the payloads as if they contain raw IPv4 and IPv6 packets. At present, GSEextract is focused on IP based protocols. However, the process for parsing non-IP traffic should be largely the same as that currently employed by the utility. Once an IP packet is successfully parsed from the raw payloads, it is converted into a .pcap compatible format and stored for analysis.

Given the high frequency of signal processing errors caused by the use of low-end equipment, many of the IP payloads identified by GSEextract will appear abruptly truncated due to missing data. In these cases, the remainder of the IP packet length is optionally padded with null bytes (0x00) to ensure compatibility with packet analysis tools like Wireshark.





Figure B.3: A notional overview of the stream interpretation and recovery approach used by GSEExtract.

C

QPEP Testbed Configurations

For clarity, we have provided a number of network configuration diagrams to detail the testbed configuration used for each of the five scenarios commonly referenced in our benchmark comparisons from Chapter 6.



(a) Plain Satellite Connection Testbed Architecture (b) Integrated PEPsal Testbed Architecture




Figure C.1: QPEP Testbed Configurations.

References

- [1] Mark Doman, Alex Palmer, and Nathanael Scott. “The Scramble for Space at Earth’s Outer Limits”. In: *ABC (Australian Broadcasting Commission) News* (Aug. 10, 2020). URL: <https://www.abc.net.au/news/2020-08-07/spacex-amazon-satellites-scramble-for-space-around-earth/12512978> (visited on 09/09/2020).
- [2] Jason Fritz. “Satellite Hacking: A Guide for the Perplexed”. In: *Culture Mandala* 10.1 (2013), p. 5906. URL: <https://cm.scholasticahq.com/article/5906-satellite-hacking-a-guide-for-the-perplexed>.
- [3] M. Manulis et al. “Cyber Security in New Space”. In: *International Journal of Information Security* (May 12, 2020).
- [4] Union of Concerned Scientists. *UCS Satellite Database*. UCS Satellite Database. 2018. URL: <https://www.ucsusa.org/nuclear-weapons/space-weapons/satellite-database> (visited on 01/21/2019).
- [5] David Livingstone MBE Dsc and Dr Patricia Lewis. *Space, the Final Frontier for Cybersecurity?* Chatham House, Sept. 22, 2016. URL: <https://www.chathamhouse.org/publication/space-final-frontier-cybersecurity> (visited on 01/18/2019).
- [6] John R. Vacca. *Computer and Information Security Handbook*. 2nd ed. San Francisco, UNITED STATES: Elsevier Science & Technology, 2012. URL: <http://ebookcentral.proquest.com/lib/oxford/detail.action?docID=1195617> (visited on 01/18/2019).
- [7] Matthew Weinzierl. “Space, the Final Economic Frontier”. In: *Journal of Economic Perspectives* 32.2 (May 2018), pp. 173–192.
- [8] Andre Tartar and Yue Qiu. “Which Rockets Are Winning the Race to Make Space Affordable?” In: *Bloomberg Businessweek* (July 26, 2018). URL: <https://www.bloomberg.com/graphics/2018-rocket-cost/> (visited on 02/04/2019).
- [9] Harry Jones. “The Recent Large Reduction in Space Launch Cost”. In: *48th International Conference on Environmental Systems*. International Conference on Environmental Systems. Albuquerque, New Mexico, 2018. URL: <https://core.ac.uk/download/pdf/288485535.pdf>.
- [10] Interorbital. *Interorbital Storefront*. Interorbital Store. URL: <http://www.interorbital.com/Store> (visited on 02/04/2019).
- [11] Matthew E. Grant. *Space Dependence - A Critical Vulnerability of the Net-Centric Operational Commander*. Naval War College, May 17, 2005, pp. 1–23. URL: <https://apps.dtic.mil/docs/citations/ADA463682> (visited on 02/11/2019).

- [12] Joseph Lungerman. “What Happens If They Say No? Preserving Access to Critical Commercial Space Capabilities during Future Crises”. In: *Air and Space Power Journal* (Dec. 2014), pp. 103–116. URL: <https://apps.dtic.mil/docs/citations/ADA617824> (visited on 02/11/2019).
- [13] James Pavur and Ivan Martinovic. “The Cyber-ASAT: On the Impact of Cyber Weapons in Outer Space”. In: *2019 11th International Conference on Cyber Conflict (CyCon)*. 2019 11th International Conference on Cyber Conflict (CyCon). Vol. 900. Tallinn, Estonia: IEEE, 2019, pp. 1–18.
- [14] Gregory Falco. “The Vacuum of Space Cyber Security”. In: *2018 AIAA SPACE and Astronautics Forum and Exposition*. AIAA SPACE and Astronautics Forum and Exposition. Orlando, FL: American Institute of Aeronautics and Astronautics, 2018.
- [15] Todd Harrison, Kaitlyn Johnson, and Thomas Roberts. *Space Threat Assessment 2018*. Center for Strategic and International Studies (CSIS), Apr. 12, 2018. URL: <https://www.csis.org/analysis/space-threat-assessment-2018> (visited on 01/25/2019).
- [16] CCSDS. *Security Threats Against Space Missions*. Report concerning space data system standards. Dec. 2015. URL: <https://public.ccsds.org/Pubs/350x1g2.pdf>.
- [17] David Fidler. *Cybersecurity and the New Era of Space Activities*. Council on Foreign Relations, Apr. 3, 2018. URL: <https://www.cfr.org/report/cybersecurity-and-new-era-space-activities> (visited on 01/18/2019).
- [18] Daria Lane et al. “High-Assurance Cyber Space Systems for Small Satellite Mission Integrity”. AIAA/USU Small Satellite Conference. Aug. 8, 2017. URL: <https://digitalcommons.usu.edu/smallsat/2017/all2017/95>.
- [19] Luca del Monte. “Towards a Cybersecurity Policy for a Sustainable, Secure and Safe Space Environment”. In: *Proceedings of the 64th International Astronautical Congress (IAC)*. 2013. URL: <https://iafastro.directory/iac/archive/browse/IAC-13/E3/4/16989/>.
- [20] James Pavur et al. “Secrets in the Sky: On Privacy and Infrastructure Security in DVB-S Satellite Broadband”. In: *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks*. WiSec ’19: Conference on Security and Privacy in Wireless and Mobile Networks. Miami, Florida: ACM, May 15, 2019.
- [21] James D. Rendleman and Robert Ryals. “Cyber Operations to Defend Space Systems?” In: *AIAA SPACE 2013 Conference and Exposition*. AIAA SPACE Conference and Exposition. San Diego, CA: American Institute of Aeronautics and Astronautics, 2013.
- [22] Nils Ole Tippenhauer et al. “On the Requirements for Successful GPS Spoofing Attacks”. In: *Proceedings of the 18th ACM Conference on Computer and Communications Security*. ACM Conference on Computer and Communications Security (CCS) (Chicago, Illinois, USA). CCS ’11. 2011, pp. 75–86.

- [23] C. Knez et al. “Lessons Learned from Applying Cyber Risk Management and Survivability Concepts to a Space Mission”. In: *2016 IEEE Aerospace Conference*. 2016 IEEE Aerospace Conference. Big Sky, MT, USA: IEEE, Mar. 2016, pp. 1–8.
- [24] Brian Young. “Commercial Satellites, Critical Information Infrastructure Protection, and Preventing Today’s Threat Actors from Becoming Tomorrow’s Captain Midnight”. In: *Strategic Cyber Defense: A Multidisciplinary Perspective* 48 (2017), p. 86.
- [25] Ted Vera. “Cyber Security Awareness for SmallSat Ground Networks”. In: *AIAA/USU Conference on Small Satellites*. AIAA/USU Conference on Small Satellites. AIAA, Aug. 10, 2016. URL: <https://digitalcommons.usu.edu/smallsat/2016/TS9GroundSystems/2>.
- [26] Julio Vivero and Ricardo Marin. “Cyber Situational Awareness in Space Organizations Operations Centres”. In: *AIAA SpaceOps Conference 2018*. 2018 SpaceOps Conference. Marseille, France: AIAA, 2018, p. 2481.
- [27] Reliability and Interoperability Council Communications Security. *Cybersecurity Risk Management and Best Practices Working Group 4: Final Report*. FCC, Mar. 2015. URL: https://transition.fcc.gov/pshs/advisory/csrc4/CSRIC_IV_WG4_Final_Report_031815.pdf.
- [28] D. J. Byrne et al. “Cyber Defense of Space-Based Assets: Verifying and Validating Defensive Designs and Implementations”. In: *Procedia Computer Science: 2014 Conference on Systems Engineering Research*. Conference on Systems Engineering Research. Vol. 28. Jan. 1, 2014, pp. 522–530.
- [29] Nicholas Cohen et al. “Spacecraft Embedded Cyber Defense- Prototypes & Experimentation”. In: *AIAA SPACE Forum 2016*. Space Forum 2016. AIAA SPACE Forum. Long Beach, California: American Institute of Aeronautics and Astronautics, Sept. 9, 2016.
- [30] Jessica A Steinberger. “A Survey of Satellite Communications System Vulnerabilities”. Air Force Institute of Technology, June 2008. URL: <https://core.ac.uk/download/pdf/288295156.pdf>.
- [31] NASA Office of the Inspector General. *Cybersecurity Management and Oversight at the Jet Propulsion Laboratory*. IG-19-022. NASA Office of the Inspector General, June 18, 2019. URL: <https://www.oversight.gov/report/nasa/cybersecurity-management-and-oversight-jet-propulsion-laboratory>.
- [32] Ben Elgin. “Network Security Breaches Plague NASA”. In: *Bloomberg Businessweek* (Nov. 20, 2008). URL: <https://www.bloomberg.com/news/articles/2008-11-19/network-security-breaches-plague-nasa> (visited on 02/06/2019).
- [33] Stefano Zatti. “The Protection of Space Missions: Threats and Cyber Threats”. In: *International Conference on Information Systems Security*. International Conference on Information Systems Security. Ed. by Rudrapatna K. Shyamasundar, Virendra Singh, and Jaideep Vaidya. Lecture Notes in Computer Science. Springer International Publishing, 2017, pp. 3–8.

- [34] Wayne A. Wheeler et al. “Cyber Resilient Flight Software for Spacecraft”. In: *AIAA SPACE and Astronautics Forum and Exposition*. SPACE and Astronautics Forum and Exposition. Orlando, FL: American Institute of Aeronautics and Astronautics, 2017.
- [35] Courier Mail. “China Targets US Satellite”. In: *The Courier Mail* (Oct. 7, 2006).
- [36] AFP. “Moscow Admits Satellite Phone Jamming”. In: (Nov. 24, 1999).
- [37] James Pavur et al. “A Tale of Sea and Sky: On the Security of Maritime VSAT Communications”. In: *IEEE Symposium on Security and Privacy (S&P)*. To Appear in 2020 IEEE Symposium on Security and Privacy (S&P). Oakland, CA: IEEE, May 2020.
- [38] Leonard David. “How Amateur Satellite Trackers Are Keeping an ‘eye’ on Objects around the Earth”. In: *Space.com* (May 3, 2020). URL: <https://www.space.com/amateur-satellite-trackers-on-global-lookout.html> (visited on 10/21/2020).
- [39] BBC Monitoring World Media. “ Hamas “Hacks into” Satellite Transmission of Israeli Channel 10 TV”. In: *BBC Worldwide Monitoring* (July 15, 2014).
- [40] Damien Francis. “Computer Virus Infects Orbiting Space Station”. In: *The Guardian. Technology* (Aug. 27, 2008). URL: <https://www.theguardian.com/technology/2008/aug/28/spacetechnology.spaceexploration> (visited on 02/06/2019).
- [41] Jeffrey Bardin. “Satellite Cyber Attack Search and Destroy”. In: *Computer and Information Security Handbook*. 2nd ed. Elsevier Science & Technology, July 5, 2013, pp. 1093–1102. URL: <https://www.sciencedirect.com/science/article/pii/B9780124166813000148>.
- [42] Paul Martin. *NASA Cybersecurity: An Examination of the Agency’s Information Security*. Testimony before the Subcommittee on Investigations and Oversight, House Committee on Science, Space, and Technology. Natinoal Aeronautics and Space Administration, Feb. 29, 2012. URL: https://oig.nasa.gov/docs/FINAL_written_statement_for_%20IT_%20hearing_February_26_edit_v2.pdf.
- [43] Davey Winder. “Confirmed: NASA Has Been Hacked”. In: *Forbes* (June 2019). URL: <https://www.forbes.com/sites/daveywinder/2019/06/20/confirmed-nasa-has-been-hacked/> (visited on 08/17/2020).
- [44] R. Banu and T. Vladimirova. “On-Board Encryption in Earth Observation Small Satellites”. In: *Proceedings 40th Annual 2006 International Carnahan Conference on Security Technology*. Proceedings 40th Annual 2006 International Carnahan Conference on Security Technology. IEEE, Oct. 2006, pp. 203–208.
- [45] New York Times. “Radio Satellites Open to Jamming”. In: *New York Times* (Feb. 28, 1962), p. 15.
- [46] Washington Post. “Soviets Ask Controls on Satellite TV”. In: *The Washington Post, Times Herald (1959-1973)* (Aug. 10, 1972), p. 1.
- [47] Bob Davis. “Captain Midnight Unmasked by FCC, Enters Guilty Plea”. In: *Wall Street Journal* (July 23, 1986), p. 1.
- [48] Andrew Pollack. “Technology: Barring Jams of Satellite TV”. In: *New York Times* (Mar. 6, 1986), p. 1.

- [49] Rebecca B. Smith. *Appeal from the United States District Court for the Eastern District of Virginia, at Norfolk. Rebecca B. Smith, District Judge*. United States Courts of Appeals, Fourth Circuit, Aug. 14, 1991. URL: <https://law.resource.org/pub/us/case/reporter/F2/940/940.F2d.653.91-5000.html> (visited on 02/06/2019).
- [50] Times of India. “Space Piracy”. In: *The Times of India* (Jan. 23, 1986), p. 8.
- [51] Gareth Parry. “Youths Hacked into Secret Nasa Network”. In: *The Guardian* (Sept. 15, 1987), p. 1.
- [52] AFP. “Foreign Satellite Programs Jammed in Iran”. In: *Agence France Presse – English* (June 18, 1994).
- [53] Small Media. *Satellite Jamming in Iran: A War Over Airwaves*. Small Media, 2012. URL: <https://smallmedia.org.uk/sites/default/files/Satellite%20Jamming.pdf>.
- [54] Wilson Wong and James Gordon Fergusson. *Military Space Power: A Guide to the Issues*. ABC-CLIO, 2010. 166 pp. Google Books: GFG5CqCojqQC.
- [55] Intelligence Newsletter. “Anybody Need A GPS Jammer?” In: *Intelligence Newsletter* (Feb. 5, 1998).
- [56] Nicholas Hellen. “Hackers Target BskyB Channels”. In: *Evening Standard* (May 26, 1993).
- [57] Sharon Walsh. “Cable TV ‘Pirate’ Faces Sentencing; Prosecutors Hope to Send Signal Of Toughness on High-Tech Crime”. In: *The Washington Post* (Dec. 13, 1994).
- [58] Associated Press. “Falun Gong Hijacks Chinese TV”. In: *Wired* (Sept. 24, 2002). URL: <https://www.wired.com/2002/09/falun-gong-hijacks-chinese-tv/> (visited on 02/06/2019).
- [59] South China Morning Post. “Falun Gong Accused of Pirate Broadcast”. In: *South China Morning Post* (Oct. 17, 2003).
- [60] Xinhua. “AsiaSat Accuses Falungong of Intercepting Satellite Signal”. In: (Nov. 21, 2004).
- [61] John Daly. “LTTE: Technologically Innovative Rebels”. In: *Asian Tribune* (June 14, 2007). URL: <https://web.archive.org/web/20190118174630/http://www.asiantribune.com/node/6151> (visited on 02/06/2019).
- [62] Siobhan Gorman, Yochi J. Dreazen, and August Cole. “Insurgents Hack U.S. Drones”. In: *Wall Street Journal. US* (Dec. 18, 2009). URL: <https://www.wsj.com/articles/SB126102247889095011> (visited on 06/01/2018).
- [63] AFP. “Iranian Government Jamming Satellite TV: Opposition Group”. In: (July 30, 2000).
- [64] Sandra Marquez. “US Condemns Cuba for Jamming Signals to Iran”. In: *Associated Press* (July 16, 2003).
- [65] Charles Arthur. “Chinese Hackers Suspected of Interfering with US Satellites”. In: *The Guardian. Technology* (Oct. 27, 2011). URL: <https://www.theguardian.com/technology/2011/oct/27/chinese-hacking-us-satellites-suspected> (visited on 02/06/2019).

- [66] Samuel Gibbs. “International Space Station Attacked by ‘Virus Epidemics’”. In: *The Guardian. Technology* (Nov. 12, 2013). URL: <https://www.theguardian.com/technology/2013/nov/12/international-space-station-virus-epidemics-malware> (visited on 02/06/2019).
- [67] Carin Zissis. *China’s Anti-Satellite Test*. Council on Foreign Relations. Feb. 22, 2007. URL: <https://www.cfr.org/backgroundunder/chinas-anti-satellite-test> (visited on 02/11/2019).
- [68] BBC Monitoring World Media. “UK-Based Al-Hiwar Satellite TV off-Air after ‘Deliberate Jamming’”. In: *BBC Worldwide Monitoring* (May 8, 2009).
- [69] BBC. “EU Pressures Iran to End Jamming”. In: *BBC News* (Mar. 22, 2010). URL: http://news.bbc.co.uk/1/hi/world/middle_east/8579719.stm (visited on 02/06/2019).
- [70] BBC Monitoring Middle East. “Jamming of Al-Jazeera TV Broadcasts Traced to Jordan; Experts Comment”. In: *BBC Worldwide Monitoring* (Oct. 2, 2010).
- [71] Nancy Messieh. *Bahrain Satellite Channel Jammed, Launches on Livestation Instead*. The Next Web. Aug. 14, 2011. URL: <https://thenextweb.com/me/2011/08/14/bahrain-satellite-channel-jammed-launches-on-livestation-instead/> (visited on 02/06/2019).
- [72] ECADF. *China Accused of Jamming TV, Websites in Ethiopia – COMPUTERWORLD Reported*. ECADF Ethiopian News. June 30, 2011. URL: <https://ecadforum.com/2011/06/30/china-accused-of-jamming-tv-websites-in-ethiopia-computerworld-reported/> (visited on 02/06/2019).
- [73] Thuraya Press Office. *Thuraya Telecom Services Affected by Intentional Jamming in Libya*. Thuraya. Feb. 25, 2011. URL: <https://www.thuraya.com/content/thuraya-telecom-services-affected-intentional-jamming-libya> (visited on 02/06/2019).
- [74] BBC Monitoring World Media. “Iran’s Arabic TV Said “Jammed from Saudi Arabia””. In: *Supplied by BBC Worldwide Monitoring* (Mar. 17, 2011).
- [75] Leo. *Les Héros Ordinaires*. Les Erythréens. Jan. 23, 2013. URL: <https://erythreens.wordpress.com/2013/01/23/les-heros-ordinaires/> (visited on 02/06/2019).
- [76] Pual Richardson. “Eritrea Accuses Ethiopia of Blocking Satellite Transmissions”. In: *Bloomberg* (Jan. 11, 2012). URL: <https://www.bloomberg.com/news/articles/2012-01-11/eritrea-accuses-ethiopia-of-blocking-satellite-transmissions> (visited on 02/06/2019).
- [77] BBC Monitoring World Media. “World Broadcasters Condemn Satellite Jamming “Emanating from Syria””. In: *BBC Worldwide Monitoring* (Oct. 22, 2012).
- [78] BBC Monitoring Trans Caucasus Unit. “Azeri Editor Says Authorities to Lose despite Jamming Satellite Signals”. In: *BBC Worldwide Monitoring* (June 27, 2013).
- [79] BBC Monitoring Middle East. “Syrian TV Still Jammed, Minister Says Channel Back on Air”. In: *BBC Worldwide Monitoring* (Feb. 10, 2018).

- [80] BBC Monitoring Asia Pacific. “North Korea Increases Jamming Electronic Signals against South - Report”. In: *BBC Worldwide Monitoring* (Oct. 9, 2013).
- [81] Deborah Housen-Couriel. “When Hamas Comes into Your Living Room”. In: *The Times of Israel* (Mar. 16, 2016).
- [82] Leonardo Egea. “Playing in a Satellite Environment 1.2”. Blackhat 2010. 2010. URL: http://www.blackhat.com/presentations/bh-dc-10/Nve_Leonardo/BlackHat-DC-2010-Nve-Playing-with-SAT-1.2-wp.pdf.
- [83] Colby Moore. “Spread Spectrum Satcom Hacking: Attacking The Globstar Simplex Data Service”. Blackhat. 2015. URL: <https://www.blackhat.com/docs/us-15/materials/us-15-Moore-Spread-Spectrum-Satcom-Hacking-Attacking-The-GlobalStar-Simplex-Data-Service-wp.pdf>.
- [84] B. Driessen et al. “Don’t Trust Satellite Phones: A Security Analysis of Two Satphone Standards”. In: *2012 IEEE Symposium on Security and Privacy*. 2012 IEEE Symposium on Security and Privacy. San Francisco, CA: IEEE, May 2012, pp. 128–142.
- [85] Jiao Hu, Ruilin Li, and Chaojing Tang. “A Real-Time Inversion Attack on the GMR-2 Cipher Used in the Satellite Phones”. In: *Science China Information Sciences* 61 (2018). URL: <https://link.springer.com/article/10.1007/s11432-017-9230-8>.
- [86] Chris Forester. “Russia "Eavesdropping" on Satellite Operations”. In: *Inside Satellite TV* (Nov. 10, 2015). URL: <https://advanced-television.com/2015/11/10/russia-eavesdropping-on-satellite-operations/>.
- [87] Symantec. *Thrip: Espionage Group Hits Satellite, Telecoms, and Defense Companies*. Symantec Enterprise Blogs. 2017. URL: <https://www.symantec.com/blogs/threat-intelligence/thrip-hits-satellite-telecoms-defense-targets> (visited on 01/23/2019).
- [88] Mary Flaherty, Jason Samenow, and Lisa Rein. *Chinese Hack U.S. Weather Systems, Satellite Network*. Washington Post. Nov. 12, 2014. URL: https://www.washingtonpost.com/local/chinese-hack-us-weather-systems-satellite-network/2014/11/12/bef1206a-68e9-11e4-b053-65cea7903f2e_story.html (visited on 02/07/2019).
- [89] Lily Newman. “Report: Another Chinese Military Unit Has Been Hacking U.S. Systems-This Time Satellite Networks”. In: *Slate Magazine* (June 10, 2014).
- [90] BBC. “Hackers Controlled Nasa Computers”. In: *BBC News. Technology* (Mar. 8, 2012). URL: <https://www.bbc.com/news/technology-17231695> (visited on 02/06/2019).
- [91] Allen Crawley. *Expedited Efforts Needed to Remediate High-Risk Vulnerabilities in JPSS Ground System*. Washington, DC: US Department of Commerce Office of the Inspector General, Aug. 21, 2014. URL: <https://www.oig.doc.gov/OIGPublications/OIG-14-027-M.pdf>.
- [92] Ruben Santamarta. “Last Call for SATCOM Security”. Blackhat USA 2018. Aug. 2018. URL: <https://i.blackhat.com/us-18/Thu-August-9/us-18-Santamarta-Last-Call-For-Satcom-Security-wp.pdf>.

- [93] Ruben Santamarta. *SATCOM Terminals: Hacking by Air, Sea, and Land*. BlackHat Whitepaper. 2014, p. 26. URL: <https://www.blackhat.com/docs/us-14/materials/us-14-Santamarta-SATCOM-Terminals-Hacking-By-Air-Sea-And-Land-WP.pdf>.
- [94] Stefan Tanase. *Satellite Turla: APT Command and Control in the Sky*. Securelist - Kaspersky Lab's cyberthreat research and reports. Sept. 9, 2015. URL: <https://securelist.com/satellite-turla-apt-command-and-control-in-the-sky/72081/> (visited on 06/20/2018).
- [95] Defense Digital Service. *Hackasat*. hackasat. 2020. URL: <https://www.hackasat.com> (visited on 08/13/2020).
- [96] AeroSpace Village. *AeroSpace Village – Securing the Skies and Beyond*. 2020. URL: <https://aerospacevillage.org/> (visited on 08/13/2020).
- [97] L. Simone, N. Salerno, and M. Maffei. “Frequency-Hopping Techniques for Secure Satellite TT+C: System Analysis +Trade-Offs”. In: *2006 International Workshop on Satellite and Space Communications*. 2006 International Workshop on Satellite and Space Communications. Madrid, Spain: IEEE, Sept. 2006, pp. 13–17.
- [98] Lachlan Gunn et al. “Anomaly Detection in Satellite Communications Systems Using LSTM Networks”. In: *2018 Military Communications and Information Systems Conference (MilCIS)*. 2018 Military Communications and Information Systems Conference (MilCIS). Canberra, ACT, Australia: IEEE, Nov. 2018, pp. 1–6.
- [99] Cortney Weinbaum, Steven Berner, and Bruce McClintock. *Sigint for Anyone: The Growing Availability of Signals Intelligence in the Public Domain*. PE-273-OSD. RAND Corporation Washington United States, Jan. 1, 2017. URL: <https://apps.dtic.mil/docs/citations/AD1053269> (visited on 01/24/2019).
- [100] A. Roy-Chowdhury et al. “Security Issues in Hybrid Networks with a Satellite Component”. In: *IEEE Wireless Communications* 12.6 (Dec. 2005), pp. 50–61.
- [101] Kai Wirt. “Fault Attack on the DVB Common Scrambling Algorithm”. In: *Computational Science and Its Applications – ICCSA 2005*. International Conference on Computational Science and Its Applications. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, May 9, 2005, pp. 577–584.
- [102] W. Li and D. Gu. “Security Analysis of DVB Common Scrambling Algorithm”. In: *The First International Symposium on Data, Privacy, and E-Commerce (ISDPE 2007)*. The First International Symposium on Data, Privacy, and E-Commerce (ISDPE 2007). Chengdu, China: IEEE, Nov. 2007, pp. 271–273.
- [103] Erik Tews, Julian Wälde, and Michael Weiner. “Breaking DVB-CSA”. In: *Research in Cryptology*. Western European Workshop on Research in Cryptology. Lecture Notes in Computer Science. Berlin: Springer, July 20, 2011, pp. 45–61.
- [104] Colibri. *PowerVu Management Keys Hacked*. May 12, 2014. URL: http://colibri.bplaced.net/PowerVu_management_keys_hacked.pdf.
- [105] Lishoy Francis et al. “Countermeasures for Attacks on Satellite Tv Cards Using Open Receivers”. In: *Third Australasian Information Security Workshop (AISW2005)*: Proceedings of the 2005 Australasian Workshop on Grid Computing and E-Research-Volume 44. Australian Computer Society, Inc., 2005, pp. 153–158.

- [106] M. P. Howarth et al. “Dynamics of Key Management in Secure Satellite Multicast”. In: *IEEE Journal on Selected Areas in Communications* 22.2 (Feb. 2004), pp. 308–319.
- [107] Yingli Sheng et al. “Security Architecture for Satellite Services over Cryptographically Heterogeneous Networks”. In: *2011 6th International ICST Conference on Communications and Networking in China (CHINACOM)*. 2011 6th International ICST Conference on Communications and Networking in China (CHINACOM). Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Aug. 2011, pp. 1093–1098.
- [108] R.J. Hughes et al. “Quantum Cryptography for Secure Satellite Communications”. In: *2000 IEEE Aerospace Conference. Proceedings*. 2000 IEEE Aerospace Conference Proceedings. Vol. 1. Big Sky, MT, USA: IEEE, 2000, pp. 191–200.
- [109] SES. *SES Announces 10 Project Partners in QUARTZ Satellite Cybersecurity Consortium*. SES. June 7, 2018. URL: <https://www.ses.com/press-release/ses-announces-10-project-partners-quartz-satellite-cybersecurity-consortium> (visited on 01/23/2019).
- [110] A. J. H. Fidler et al. “Satellite — A New Opportunity for Broadband Applications”. In: *BT Technology Journal* 20.1 (Jan. 1, 2002), pp. 29–37.
- [111] S. Iyengar et al. “Security Requirements for IP over Satellite DVB Networks”. In: *2007 16th IST Mobile and Wireless Communications Summit*. 2007 16th IST Mobile and Wireless Communications Summit. IEEE, July 2007, pp. 1–6.
- [112] Laurence Duquerroy et al. “SatiPsec : An Optimized Solution for Securing Multicast and Unicast Satellite Transmissions”. In: *22nd AIAA International Communications Satellite Systems Conference & Exhibit 2004 (ICSSC)*. AIAA International Communications Satellite Systems Conference & Exhibit (ICSSC). Monterey, California: American Institute of Aeronautics and Astronautics, May 2004.
- [113] James Pavur et al. “QPEP: A QUIC-Based Approach to Encrypted Performance Enhancing Proxies for High-Latency Satellite Broadband”. In: *Network and Distributed System Security Symposium 2021*. Network and Distributed System Security Symposium. Internet Society, Feb. 12, 2020. arXiv: 2002.05091 [cs]. URL: <https://www.ndss-symposium.org/ndss-paper/qpep-an-actionable-approach-to-secure-and-performant-broadband-from-geostationary-orbit/> (visited on 02/13/2020).
- [114] James Pavur. “Whispers Among the Stars: A Practical Look at Perpetrating (and Preventing) Satellite Eavesdropping Attacks”. Conference briefing. Conference briefing. Black Hat USA. Las Vegas, NV, Aug. 5, 2020. URL: <https://www.blackhat.com/us-20/briefings/schedule/index.html#whispers-among-the-stars-a-practical-look-at-perpetrating-and-preventing-satellite-eavesdropping-attacks-19391> (visited on 08/13/2020).
- [115] James Pavur. “Whispers Among the Stars”. Conference briefing. DEFCON 28 - Safe Mode. Aug. 5, 2020. URL: https://www.youtube.com/watch?v=ku0Q_Wey4K0 (visited on 08/13/2020).

- [116] CCSDS. *Space Data Link Security Protocol - Summary of Concept and Rationale*. Green book. June 2018. URL: <https://public.ccsds.org/Pubs/350x5g1.pdf>.
- [117] rtl-sdr.com. *Cheating at Pokémon Go with a HackRF and GPS Spoofing*. rtl-sdr.com. July 20, 2016. URL: <https://www.rtl-sdr.com/cheating-at-pokemon-go-with-a-hackrf-and-gps-spoofing/> (visited on 08/13/2020).
- [118] Todd E. Humphreys. “Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer”. In: *In Proceedings of the Institute of Navigation GNSS (ION GNSS)*. Savanna, GA, 2008. URL: <http://hdl.handle.net/2152/63316>.
- [119] Kyle Wesson, Daniel Shepard, and Todd Humphreys. “Straight Talk on Anti-Spoofing”. In: *Gps World* 23.1 (2012), pp. 32–39. URL: https://radionavlab.ae.utexas.edu/images/stories/files/papers/antiSpoofStraightTalk_Wesson.pdf.
- [120] Jon S Warner and Roger G Johnston. “GPS Spoofing Countermeasures”. In: *Homeland Security Journal* 25.2 (2003), pp. 19–27.
- [121] Hengqing Wen et al. “Countermeasures for GPS Signal Spoofing”. In: *Proceedings of the 18th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS)*. International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS). Long Beach, CA, 2005, pp. 1285–1290. URL: <https://www.ion.org/publications/abstract.cfm?articleID=6325>.
- [122] Ali Jafarnia-Jahromi et al. “GPS Vulnerability to Spoofing Threats and a Review of Antispoofing Techniques”. In: *International Journal of Navigation and Observation* 2012 (2012).
- [123] Julio Vivero and Luca del Monte. “Space Missions Cybersecurity”. In: *AIAA SpaceOps 2014*. SpaceOps 2014 Conference. 2014, p. 1765.
- [124] eoPortal. *SwissCube*. eoPortalDirectory. URL: <https://directory.eoportal.org/web/eoportal/satellite-missions/s/swisscube> (visited on 08/13/2020).
- [125] Mathias Pietzka. “Development and Characterization of a Propulsion System for CubeSats Based on Vacuum Arc Thrusters”. Universität der Bundeswehr München, 2016. URL: <https://d-nb.info/1115728423/34>.
- [126] Thomas Llanos and Dallas Pearson. “Achieving Space Mission Resilience To Cyber Attack: Architectural Implications”. In: *AIAA SPACE 2016*. AIAA Space Forum. AIAA SPACE Forum. AIAA, Sept. 2016.
- [127] Robert Lemos. “Satellite Control Codes Stolen by Hackers”. In: *ZDnet* (Mar. 7, 2001). URL: <https://www.zdnet.com/article/satellite-control-codes-stolen-by-hackers/> (visited on 02/07/2019).
- [128] Catherine Wilson. “Teen Given Six Months for Hacking Into NASA”. In: *ABC News* (Sept. 22, 2000). URL: <https://abcnews.go.com/Technology/story?id=119422&page=1> (visited on 02/07/2019).
- [129] Amazon Web Services. *AWS Ground Station*. Amazon Web Services, Inc. URL: <https://aws.amazon.com/ground-station/> (visited on 08/13/2020).

- [130] Daniel Cunningham, Geancarlo Palavincini, and Jose Romero-Mariona. “Towards Effective Cybersecurity for Modular, Open Architecture Satellite Systems”. In: *AIAA/USU Conference on Small Satellites*. AIAA, Aug. 9, 2016. URL: <https://digitalcommons.usu.edu/smallsat/2016/TS4AdvTech1/6>.
- [131] Julio Vivero. “Space Missions Cybersecurity Modelling”. In: *31st AIAA International Communications Satellite Systems Conference*. AIAA International Communications Satellite Systems Conference\, Florence, Italy: American Institute of Aeronautics and Astronautics, 2013.
- [132] U.S. Government Accountability Office. *Critical Infrastructure Protection: Commercial Satellite Security Should Be More Fully Addressed*. GAO-02-78. Government Accountability Office, Oct. 3, 2002. URL: <https://www.gao.gov/products/GAO-02-781> (visited on 01/21/2019).
- [133] Paul Meyer. “International Cyber Norms: Legal, Policy & Industry Perspectives”. In: *Simons Working Paper Series in Security and Development*. Tallinn, Estonia: NATO CCD COE Publications, 2016, pp. 155–169. URL: https://ccdcoe.org/uploads/2018/10/InternationalCyberNorms_Ch8.pdf (visited on 01/18/2019).
- [134] Deborah Housen-Couriel. “Cybersecurity and Anti-Satellite Capabilities (ASAT): New Threats and New Legal Responses”. In: *Journal of Law & Cyber Warfare* 4.3 (2015), p. 116. JSTOR: 26441259.
- [135] P. J. Blount. “Satellites Are Just Things on the Internet of Things”. In: *Air and Space Law* 42.3 (2017), pp. 273–293. URL: <https://kluwerlawonline.com/journalarticle/Air+and+Space+Law/42.3/AILA2017019>.
- [136] André Adelsbach and Ulrich Greveler. “Satellite Communication without Privacy - Attacker’s Paradise”. In: *Sicherheit – Schutz Und Zuverlässigkeit*. 2005, pp. 257–268. URL: <https://dl.gi.de/handle/20.500.12116/28380>.
- [137] Adam Laurie. “\$atellite Hacking for Fun & Profit!” Blackhat 2009. 2009. URL: <http://www.blackhat.com/presentations/bh-dc-09/Laurie/BlackHat-DC-09-Laurie-Satellite-Hacking.pdf>.
- [138] H. Cruickshank et al. “Securing Multicast in DVB-RCS Satellite Systems”. In: *IEEE Wireless Communications* 12.5 (Oct. 2005), pp. 38–45.
- [139] Cjcr-Software. *EBSpro*. 2016. URL: <http://ebspro.net/> (visited on 06/20/2018).
- [140] crazycat69. *CrazyScan: Satellite/Terrestrial/Cable Scan Software*. SourceForge. 2018. URL: <https://sourceforge.net/projects/crazyscan/> (visited on 06/20/2018).
- [141] World-Satellite. *World-Satellite Forum*. World Satellite Forum. 2018. URL: <http://www.world-satellite.net/> (visited on 06/20/2018).
- [142] DVB Project. *Digital Video Broadcasting (DVB); Second Generation Framing Structure, Channel Coding and Modulation Systems for Broadcasting, Interactive Services, News Gathering and Other Broadband Satellite Applications*. 2014. URL: <https://www.dvb.org/standards/dvb-s2> (visited on 06/01/2018).
- [143] Koen Williams. *DVB-S2X Demystified*. White paper. Newtec, Mar. 2014. URL: https://www.newtec.eu/frontend/files/userfiles/files/Whitepaper%20DVB_S2X.pdf.

- [144] Comsys. *VSAT Network Types*. Comsys. URL: https://www.comsys.co.uk/wvr_nets.htm (visited on 03/07/2019).
- [145] Inmarsat. *BGAN Voice and Broadband Service*. Inmarsat. URL: <https://www.inmarsat.com/service/bgan/> (visited on 04/01/2019).
- [146] International Telecommunication Union. *H.222.0 Information Technology – Generic Coding of Moving Pictures and Associated Audio Information: Systems*. ITU, Oct. 2014. URL: <https://www.itu.int/rec/T-REC-H.222.0-201410-S/en>.
- [147] Vahe Balabanian et al. “An Introduction to DSM-CC”. In: *IEEE Communications Magazine* (Nov. 1996). URL: <https://web.archive.org/web/20190301195809/http://www.iuma.ulpgc.es/~nunez/procmultimedia98-00/csel/mpeg/documents/dsmcc/dsmcc.htm> (visited on 06/01/2018).
- [148] Teh Chee Hong, Wan Tat Chee, and R. Budiarto. “A Comparison of IP Datagrams Transmission Using MPE and ULE over Mpeg-2/DVB Networks”. In: *2005 5th International Conference on Information Communications Signal Processing*. 2005 5th International Conference on Information Communications Signal Processing. Bangkok, Thailand: IEEE, Dec. 2005, pp. 1173–1177.
- [149] BroadSat. *Opensky: One Way Satellite Internet*. Opensky. URL: <https://www.broadsat.com/en/opensky/> (visited on 06/01/2018).
- [150] Dan Lester and Harley Thronson. “Human Space Exploration and Human Spaceflight: Latency and the Cognitive Scale of the Universe”. In: *Space Policy* 27.2 (May 1, 2011), pp. 89–93.
- [151] Matt Cocchiaro. “Vessel Accumulation and Cargo Value Estimation”. Genova IUMI (Genoa Italy). Sept. 2016. URL: https://iumi.com/images/documents/genoa-2016-program/7_matthew_cocchiaro_1474364889.pdf.
- [152] Martin Cox. *CMA CGM BENJAMIN FRANKLIN Gets Hollywood Welcome*. Maritime Matters. Dec. 27, 2015. URL: <https://web.archive.org/web/20200918194827/http://maritimematters.com/2015/12/cma-cgm-benjamin-franklin-gets-hollywood-welcome/> (visited on 04/24/2019).
- [153] United Nations Conference on Trade and Development. *Review of Maritime Transport 2018*. New York and Geneva: UNITED NATIONS, 2018. URL: https://unctad.org/en/PublicationsLibrary/rmt2018_en.pdf.
- [154] Stefano Brizzolara and Robert A. Brizzolara. “Autonomous Sea Surface Vehicles”. In: *Springer Handbook of Ocean Engineering*. Ed. by Manhar R. Dhanak and Nikolaos I. Xiros. Springer Handbooks. Cham: Springer International Publishing, 2016, pp. 323–340.
- [155] J. DiRenzo, D. A. Goward, and F. S. Roberts. “The Little-Known Challenge of Maritime Cyber Security”. In: *2015 6th International Conference on Information, Intelligence, Systems and Applications (IISA)*. 2015 6th International Conference on Information, Intelligence, Systems and Applications (IISA). Corfu, Greece: IEEE, July 2015, pp. 1–5.

- [156] Etherr. *Maritime VSAT System*. Jan. 21, 2011. URL: https://commons.wikimedia.org/wiki/File:Marine_vsat.jpg (visited on 04/25/2019).
- [157] Jahshan A. Bhatti and Todd E. Humphreys. “Hostile Control of Ships via False GPS Signals: Demonstration and Detection”. In: *Navigation* 64.1 (Spr. 2017), pp. 51–66.
- [158] Center for Advanced Defense Studies. *Above Us Only Stars: Exposing GPS Spoofing in Russia and Syria*. Center for Advanced Defense Studies, 2019. URL: <https://www.arcgis.com/apps/Cascade/index.html?appid=b919c8d91b0a4f868f02acfdabc428d7&classicembedmode> (visited on 06/11/2019).
- [159] Desmond Schmidt et al. “A Survey and Analysis of the GNSS Spoofing Threat and Countermeasures”. In: *ACM Computing Surveys* 48.4 (May 2016), 64:1–64:31.
- [160] myshiptracking.com. *My Ship Tracking*. My Ship Tracking. URL: <http://www.myshiptracking.com> (visited on 06/11/2019).
- [161] Marine Traffic. *MarineTraffic: Global Ship Tracking Intelligence / AIS Marine Traffic*. Marine Traffic. URL: <https://www.marinetraffic.com/en/ais/home> (visited on 06/11/2019).
- [162] Marco Balduzzi, Alessandro Pasta, and Kyle Wilhoit. “A Security Evaluation of AIS Automated Identification System”. In: *Proceedings of the 30th Annual Computer Security Applications Conference*. Annual Computer Security Applications Conference. ACSAC ’14. New Orleans, Louisiana: ACM, Dec. 2014, pp. 436–445.
- [163] Kevin D Jones, Kimberly Tam, and Maria Papadaki. “Threats and Impacts in Maritime Cyber Security”. In: *Engineering & Technology Reference* 1.1 (2016).
- [164] O. Jacq et al. “Detecting and Hunting Cyberthreats in a Maritime Environment: Specification and Experimentation of a Maritime Cybersecurity Operations Centre”. In: *2018 2nd Cyber Security in Networking Conference (CSNet)*. 2018 2nd Cyber Security in Networking Conference (CSNet). Paris, France: IEEE, Oct. 2018, pp. 1–8.
- [165] Rory Hopcraft and Keith M. Martin. “Effective Maritime Cybersecurity Regulation – the Case for a Cyber Code”. In: *Journal of the Indian Ocean Region* 14.3 (Sept. 2, 2018), pp. 354–366.
- [166] Aniseh Tabrizi. “What We Know about Gulf of Oman Tanker Attacks”. In: *BBC News. Middle East* (June 18, 2019). URL: <https://www.bbc.com/news/world-middle-east-48627014> (visited on 07/30/2019).
- [167] David Kirkpatrick. “Tankers Are Attacked in Mideast, and U.S. Says Video Shows Iran Was Involved”. In: *New York Times* (June 13, 2019). URL: <https://www.nytimes.com/2019/06/13/world/middleeast/oil-tanker-attack-gulf-oman.html> (visited on 07/23/2019).
- [168] A. R. Lee and H. P. Wogan. “All at Sea: The Modern Seascape of Cybersecurity Threats of the Maritime Industry”. In: *OCEANS 2018 MTS/IEEE Charleston*. OCEANS 2018 MTS/IEEE Charleston. Charleston, SC: IEEE, Oct. 2018, pp. 1–8.

- [169] Matthew Smith et al. “Undermining Privacy in the Aircraft Communications Addressing and Reporting System (ACARS)”. In: *Proceedings on Privacy Enhancing Technologies* 2018.3 (June 1, 2018), pp. 105–122.
- [170] ETSI. *ETSI TS 102 606 V1.1.1 Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE); Part 1: Protocol*. DVB BlueBook. Dec. 2013. URL: https://www.dvb.org/resources/public/standards/a116-1_gse_specification.pdf.
- [171] iDirect. *The Maritime VSAT Advantage: A Cost Analysis of VSAT Broadband versus L-Band Pay-per-Use Service*. URL: http://www.groundcontrol.com/Maritime_VSAT/Marine_VSAT_Comparison.pdf (visited on 06/12/2019).
- [172] Martyn Wingrove. *Cruise Ship Orders for VSAT Providers*. Sept. 4, 2018. URL: https://www.passengership.info/news/view,cruise-ship-orders-for-vsats-providers_54065.htm (visited on 06/12/2019).
- [173] MarineMec. *Fishing Vessel Owners Turn to VSAT*. Maritime Digitalisation & Communications. Nov. 17, 2015. URL: https://web.archive.org/web/20160426023124/https://www.marinemec.com/news/view,fishing-vessel-owners-turn-to-vsats_40942.htm (visited on 06/12/2019).
- [174] Tero Marine. *Products - Tero Marine*. Tero Marine. URL: <https://www.teromarine.com/products/> (visited on 06/12/2019).
- [175] Intellian. *V240mt*. Intellian V-Series. URL: <https://www.intelliantech.com/Satcom/v-series/v240mt> (visited on 06/11/2019).
- [176] iDirect. *iDirect Evolution*. iDirect. 2018. URL: <https://www.idirect.net/wp-content/uploads/2019/01/iDirectEvolution-ProductBrochure2018.pdf>.
- [177] ETSI. *ETSI 300 421 V1.1.2 Digital Video Broadcasting (DVB); Framing Structure, Channel Coding and Modulation for 11/12 GHz Satellite Services*. DVB blue book. 1997. URL: https://www.etsi.org/deliver/etsi_en/300400_300499/300421/01.01.02_60/en_300421v010102p.pdf.
- [178] ETSI. *ETSI EN 302 307 V1.3.1 Digital Video Broadcasting (DVB); Second Generation Framing Structure, Channel Coding and Modulation Systems for Broadcasting, Interactive Services, News Gathering and Other Broadband Satellite Applications*. DVB BlueBook. 2014. URL: https://www.dvb.org/resources/public/standards/a83-1_dvb-s2_den302307v141.pdf.
- [179] Newtec. *Satcom for National Security & Intelligence Gathering*. Newtec. 2015. URL: <https://www.newtec.eu/article/application-note/intelligence-gathering> (visited on 06/01/2018).
- [180] Via Satellite and iDirect. *The Coming Wave of Maritime VSAT Growth*. Via Satellite. 2015. URL: <https://www.satellitetoday.com/long-form-stories/maritime-vsats/> (visited on 07/26/2019).
- [181] Fortune. *Global 500*. Fortune. 2019. URL: <https://fortune.com/global500/2019/> (visited on 11/04/2019).

- [182] International Hydrographic Organization. *IHO Data Protection Scheme*. Jan. 2015. URL: https://web.archive.org/web/20171025222715/https://iho.int/iho_pubs/standard/S-63/S-63_e1.2.0_EN_Jan2015.pdf.
- [183] International Hydrographic Organization. *IHO Transfer Standard for Digital Hydrographic Data*. Nov. 2000. URL: https://web.archive.org/web/20180920170703/http://www.iho.int/iho_pubs/standard/S-57Ed3.1/31Main.pdf.
- [184] Markus Binder, Jillian Quigley, and Herbert Tinsley. *Islamic State Chemical Weapons: A Case Contained by Its Context?* Combating Terrorism Center at West Point, Mar. 29, 2018. URL: <https://ctc.usma.edu/islamic-state-chemical-weapons-case-contained-context/> (visited on 06/07/2019).
- [185] Hans Christian Rudolph and Nils Grundmann. *CipherSuite*. Cipher Suite Info. 2019. URL: <https://ciphersuite.info/> (visited on 07/17/2019).
- [186] Microsoft. *Microsoft Security Bulletin MS08-020*. Security Bulletins. Apr. 8, 2008. URL: <https://docs.microsoft.com/en-us/security-updates/securitybulletins/2008/ms08-020> (visited on 11/04/2019).
- [187] Tellitec. “TL200 TelliShape V2.6”. In: (). URL: <https://www.newtec.eu/frontend/files/leaflet/nop1700-bandwidth-manager-and-shaper-software.pdf> (visited on 06/10/2019).
- [188] Viveris Technologies. *OpenSAND*. Apr. 2019. URL: <https://forge.net4sat.org/opensand/opensand>.
- [189] Hughes. *What Should I Do If I Have Slow or No Connectivity?* HughesNet Support. URL: <https://web.archive.org/web/20200714013838/https://support.hughesnet.com/en/faq/internet/slow-no-connectivity>.
- [190] Viasat. *Answers to Your Questions Relating to Your Viasat Internet Service during This Difficult Time*. Viasat Help Center. 2020. URL: https://help.viasat.com/internet/articles/Denver_FAQ/Answers-to-your-Questions-relating-to-your-Viasat-Internet-service-during-this-difficult-time.
- [191] Margaret M McMahon and Robert Rathburn. *Measuring Latency in Iridium Satellite Constellation Data Services*. US Naval Academy, 2005. URL: <https://apps.dtic.mil/dtic/tr/fulltext/u2/a464192.pdf>.
- [192] SpaceX. “Starlink Mission”. In: *SpaceX* (Jan. 2020). URL: <https://www.spacex.com/news/2020/01/07/starlink-mission>.
- [193] OneWeb. *OneWeb*. OneWeb. URL: <http://oneweb.world/> (visited on 02/04/2019).
- [194] Carlo Caini and Rosario Firrincieli. “TCP Hybla: A TCP Enhancement for Heterogeneous Networks”. In: *International Journal of Satellite Communications and Networking* 22.5 (2004), pp. 547–566.
- [195] S. Oueslati-Boulahia et al. “TCP over Satellite Links: Problems and Solutions”. In: *Telecommunication Systems* 13.2 (July 2000), pp. 199–212.
- [196] C. Caini et al. “TCP, PEP and DTN Performance on Disruptive Satellite Channels”. In: *2009 International Workshop on Satellite and Space Communications*. Sept. 2009, pp. 371–375.

- [197] Igor Bisio, Mario Marchese, and Maurizio Mongelli. “Performance Enhanced Proxy Solutions for Satellite Networks: State of the Art, Protocol Stack and Possible Interfaces”. In: *Personal Satellite Services (PSATS) 2009*. Ed. by Kandeepan Sithamparanathan and Mario Marchese. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Berlin, Heidelberg: Springer, 2009, pp. 61–67.
- [198] Thomas R. Henderson and Randy H. Katz. *TCP Performance over Satellite Channels*. UCB/CSD-99-1083. EECS Department, University of California, Berkeley, Dec. 1999. URL: <https://www2.eecs.berkeley.edu/Pubs/TechRpts/1999/CSD-99-1083.pdf>.
- [199] Per Hurtig, Wolfgang John, and Anna Brunstrom. “Recent Trends in TCP Packet-Level Characteristics”. In: *Proceedings of the 7th International Conference on Networking and Services (ICNS)*. International Conference on Networking and Services (ICNS). Venice, Italy: IARIA, 2011, pp. 179–195. URL: <https://www.semanticscholar.org/paper/Recent-Trends-in-TCP-Packet-Level-Characteristics-Hurtig-John/6fef36d3285997391a9d7d5259dd78c9d63a81aa>.
- [200] Kostas Pentikousis and Hussein Badr. “Quantifying the Deployment of TCP Options-a Comparative Study”. In: *IEEE Communications Letters* 8.10 (2004), pp. 647–649.
- [201] J. Border et al. *RFC 3135 - Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations*. IETF Tracker. 2001. URL: <https://tools.ietf.org/html/rfc3135>.
- [202] Carlo Caini and Rosario Firrincieli. “DTN and Satellite Communications”. In: *Delay Tolerant Networks: Protocols and Applications*. 1st ed. CRC Press, 2012. URL: <https://www.taylorfrancis.com/chapters/dtn-satellite-communications-carlo-caini-rosario-firrincieli/e/10.1201/b11309-10>.
- [203] Alain Pirovano and Fabien Garcia. “A New Survey on Improving TCP Performances over Geostationary Satellite Link”. In: *Network and Communication Technologies* 2.1 (Jan. 2013).
- [204] Lun Li et al. “Secure Spectrum-Efficient Frequency Hopping for Return Link of Protected Tactical Satellite Communications”. In: *2016 IEEE Military Communications Conference*. MILCOM 2016. Baltimore, MD, USA: IEEE, Nov. 2016, pp. 254–258.
- [205] Gan Zheng, Pantelis-Daniel Arapoglou, and Bjorn Ottersten. “Physical Layer Security in Multibeam Satellite Systems”. In: *IEEE Transactions on Wireless Communications* 11.2 (Feb. 2012), pp. 852–863.
- [206] Lekhemissi Djedjai and Rong Ke Liu. “IPSecOPEP: IPSec over PEPs Architecture, for Secure and Optimized Communications over Satellite Links”. In: *2016 7th IEEE International Conference on Software Engineering and Service Science (ICSESS)*. International Conference on Software Engineering and Service Science (ICSESS). Beijing, China: IEEE, Aug. 2016, pp. 264–268.
- [207] Encore Networks. *BANDIT 2*. Encore Networks. URL: <https://www.encorenetworks.com/products/bandit-ii/>.

- [208] DSD Telecom. *DSD Satellite VPN*. Satellite VPN. 2019. URL: <http://www.dsdsatellite.com/index.php/support/satellite-vpn>.
- [209] Newtec. *EL810 Mobile PEP-Box Terminal*. Oct. 2008. URL: <http://www.tellitec.be/tellinet/Newtec-Elevation-EL810%20Mobile%20PEP-Box%20Terminal%20-%20web%20print.pdf>.
- [210] RigNet. *CyphreLink - End-to-End Data Protection*. Cyphre. URL: <https://www.cyphre.com/cyphrelink/>.
- [211] Hughes. *HN7000S Modem Product Page*. Satellite Ground Systems: HN7000S. URL: <https://web.archive.org/web/20170718035749/https://www.hughes.com/technologies/broadband-satellite-systems/hn-systems/hn7000s>.
- [212] Tellitec GmbH. *TC-Spy Documentation*. Wikileaks. Mar. 2006. URL: https://wikileaks.org/wiki/Tellitec_Tellinet_Sat_%20Spy_manual,_6_Mar_2006.
- [213] André Adelsbach and Ulrich Greveler. “Insider Attacks Enabling Data Broadcasting on Crypto-Enforced Unicast Links”. In: *European Symposium On Research In Computer Security*. ESORICS. Ed. by Joachim Biskup and Javier López. Lecture Notes in Computer Science. Dresden, Germany: Springer, Sept. 2007, pp. 469–484.
- [214] Ground Control. “IG-VPN: Using Application Layer Technology to Overcome the Impact of Satellite Circuit Latency on VPN Performance”. In: (2003). URL: https://www.groundcontrol.com/IG-VPN_Intro.pdf.
- [215] Ludovic Thomas et al. “Google QUIC Performance over a Public SATCOM Access”. In: *International Journal of Satellite Communications and Networking* 37.6 (2019), pp. 601–611.
- [216] Daniele Lacamera and Sergio Ammirata. *PEPsal: A TCP Performance Enhancing Proxy for Satellite Links*. Sept. 2016. URL: <https://github.com/danielinux/pepsal>.
- [217] OpenVPN Inc. *OpenVPN*. OpenVPN. 2020. URL: <https://openvpn.net/>.
- [218] Jana Iyengar and Martin Thomson. *QUIC: A UDP-Based Multiplexed and Secure Transport*. Request for Comments RFC 9000. Internet Engineering Task Force, May 2021. 151 pp.
- [219] Shan Chen et al. “Secure Communication Channel Establishment: TLS 1.3 (over TCP Fast Open) vs. QUIC”. In: *Computer Security – ESORICS 2019*. Ed. by Kazuo Sako, Steve Schneider, and Peter Y. A. Ryan. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2019, pp. 404–426.
- [220] Antoine Delignat-Lavaud et al. “A Security Model and Fully Verified Implementation for the IETF QUIC Record Layer”. In: *2021 IEEE Symposium on Security and Privacy (SP)*. 2021 IEEE Symposium on Security and Privacy (SP). May 2021, pp. 1162–1178.
- [221] Jingjing Zhang et al. “Formal Analysis of QUIC Handshake Protocol Using Symbolic Model Checking”. In: *IEEE Access* 9 (2021), pp. 14836–14848.
- [222] Robert Lychev et al. “How Secure and Quick Is QUIC? Provable Security and Performance Analyses”. In: *2015 IEEE Symposium on Security and Privacy*. 2015 IEEE Symposium on Security and Privacy. May 2015, pp. 214–231.

- [223] Cloudflare. *What Is a QUIC Flood DDoS Attack?* Cloudflare. URL: <https://www.cloudflare.com/learning/ddos/what-is-a-quic-flood/>.
- [224] Martin Thomson and Sean Turner. *Using TLS to Secure QUIC*. Request for Comments RFC 9001. Internet Engineering Task Force, May 2021. 52 pp.
- [225] Ian Swett. *QUIC FEC v1 - Google Groups*. QUIC Prototype Protocol Discussion Group. Feb. 2016. URL: <https://groups.google.com/a/chromium.org/forum/#!topic/proto-quic/Z5qKkk2XZe0>.
- [226] John P Rula et al. “Mile High Wifi: A First Look at in-Flight Internet Connectivity”. In: *Proceedings of the 2018 World Wide Web Conference*. International World Wide Web Conferences Steering Committee, 2018, pp. 1449–1458.
- [227] Nicholas Kuhn, John Border, and Emile Stephan. *QUIC for SATCOM*. IETF Tracker. Nov. 2019. URL: <https://www.potaroo.net/ietf/idref/draft-kuhn-quic-4-sat/>.
- [228] Marc Fischlin and Felix Günther. “Replay Attacks on Zero Round-Trip Time: The Case of the TLS 1.3 Handshake Candidates”. In: *2017 IEEE European Symposium on Security and Privacy (EuroS P)*. IEEE European Symposium on Security and Privacy (EuroSP). Paris, France: IEEE, Apr. 2017, pp. 60–75.
- [229] Grégoire Delannoy. *TCPeP*. May 2013. URL: <https://github.com/GregoireDelannoy/TCPeP>.
- [230] Dimitris Velenis, Dimitris Kalogeras, and Basil Maglaris. “SaTPEP: A TCP Performance Enhancing Proxy for Satellite Links”. In: *NETWORKING 2002: Networking Technologies, Services, and Protocols*. Ed. by Enrico Gregori et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2002, pp. 1233–1238.
- [231] Lucas Clemente. *Quic-Go*. Aug. 2019. URL: <https://github.com/lucas-clemente/quic-go>.
- [232] Google. *QUIC, a Multiplexed Stream Transport over UDP - The Chromium Projects*. Chromium.org. 2020. URL: <https://www.chromium.org/quic>.
- [233] Antoine Auger, Emmanuel Lochin, and Nicolas Kuhn. “Making Trustable Satellite Experiments: An Application to a VoIP Scenario”. In: *IEEE 89th Vehicular Technology Conference*. IEEE VTC2019-Spring. IEEE, 2019, pp. 1–5.
- [234] SitespeedIO. *Browsertime*. sitespeed.io. Dec. 2019. URL: <https://www.sitespeed.io/>.
- [235] Mozilla.org. *PerformanceTiming.navigationStart*. MDN Web Docs. Sept. 219. URL: <https://developer.mozilla.org/en-US/docs/Web/API/PerformanceTiming/navigationStart>.
- [236] Alexa. *The Top 500 Sites on the Web*. Alexa Top Sites. Jan. 2020. URL: <https://www.alexa.com/topsites>.
- [237] Ulrich Speidel. *Simulating Satellite Internet Traffic to a Small Island Internet Provider*. Information Society Innovation Fund (ISIF Asia), Jan. 2017. URL: <https://isif.asia/simulating-satellite-internet-traffic-to-a-small-island-internet-provider/>.

- [238] ETSI. *ETSI TR 102 768: Digital Video Broadcasting (DVB); Interaction Channel for Satellite Distribution Systems; Guidelines for the Use of EN 301 790 in Mobile Scenarios*. Tech. rep. ETSI TR 102 768. 2009.
- [239] ETSI. *ETSI TR 102 376-1 Digital Video Broadcasting (DVB); Implementation Guidelines for the Second Generation System for Broadcasting, Interactive Services, News Gathering and Other Broadband Satellite Applications; Part 1: DVB-S2*. Technical Report ETSI TR 102 376-2. 2015.
- [240] ITU. *ITU-T Y.1541 Network Performance Objectives for IP-Based Services*. Y-Series Recommendation. Dec. 2011.
- [241] Amal Boubaker. *OpenSAND Example of Delay Variations*. Net4Sat OpenSAND Wiki. June 2019. URL: https://wiki.net4sat.org/doku.php?id=opensand:emulated_satcom_features:physical:delay:delay_example:index.
- [242] A. Custura, T. Jones, and G. Fairhurst. “Rethinking ACKs at the Transport Layer”. In: *2020 IFIP Networking Conference (Networking)*. Paris, France: IEEE, June 2020, pp. 731–736. URL: <https://ieeexplore.ieee.org/abstract/document/9142733>.
- [243] Tom Jones, Ana Custura, and Gorrry Fairhurst. *Changing the Default QUIC ACK Policy*. Sept. 2020. URL: <https://tools.ietf.org/id/draft-fairhurst-quic-ack-scaling-03.html>.
- [244] NASA. *Frequently Asked Questions*. ARES: Orbital Debris Program Office. 2018. URL: <https://orbitaldebris.jsc.nasa.gov/faq.html> (visited on 12/10/2018).
- [245] Union of Concerned Scientists. *Satellite Database*. UCS Satellite Database. Apr. 1, 2020. URL: <https://www.ucsusa.org/resources/satellite-database> (visited on 07/30/2020).
- [246] N.N. Smirnov, A.I. Nazarenko, and A.B. Kiselev. “Continuum Model for Space Debris Evolution with Account of Collisions and Orbital Breakups”. In: *Space Debris 2.4* (Jan. 1, 2000), pp. 249–271.
- [247] SpaceQuest. *GNSS-701 Satellite GNSS Receiver*. SpaceQuest Ltd. 2017. URL: <http://www.spacequest.com/attitude-determination-control/gps12-v1> (visited on 04/28/2020).
- [248] Megan Wallace. *Tracking and Data Relay Satellite (TDRS)*. NASA. URL: http://www.nasa.gov/directorates/heo/scan/services/networks/tdrs_main (visited on 04/28/2020).
- [249] Ball Aerospace. *Star Trackers*. Ball Aerospace. 2020. URL: <https://www.ball.com/aerospace/markets-capabilities/capabilities/technologies-components/star-trackers> (visited on 04/28/2020).
- [250] Mykola Kaliuzhnyi et al. “International Network of Passive Correlation Ranging for Orbit Determination of a Geostationary Satellite”. In: *Odessa Astronomical Publications* 29 (2016), pp. 203–206.
- [251] T. S. Kelso. *Orbit Determination*. Celestrak. 1995. URL: <https://www.celestrak.com/columns/v01n06/> (visited on 03/06/2020).

- [252] Mike Gruss. “U.S. Plans \$6 Billion Investment in Space Situational Awareness”. In: *Space News* (Oct. 19, 2015). URL: <https://spacenews.com/planned-u-s-investment-in-space-awareness-is-6-billion-gao-says/> (visited on 07/30/2020).
- [253] Tommaso Sgobba, Firooz A. Allahdadi, and Fernand Alby. “Orbital Operations Safety”. In: *Safety Design for Space Operations*. Butterworth-Heinemann, Mar. 2013, pp. 411–431. URL: <https://learning.oreilly.com/library/view/safety-design-for/9780080969213/>.
- [254] Bahavya Lal et al. *Global Trends in Space Situational Awareness and Space Traffic Management*. Institute for Defense Analyses, Apr. 2018, pp. 1–153. URL: <https://www.ida.org/idamedia/Corporate/Files/Publications/STPIPubs/2018/D-9074.pdf> (visited on 12/10/2018).
- [255] ISON. *International Scientific Optical Network (ISON) & Low Frequency VLBI Network (LFVN)*. ISON. 2011. URL: <http://lfvn.astronomer.ru/index.htm> (visited on 03/09/2020).
- [256] Brian Weeden. *Space Situational Awareness Fact Sheet*. Secure World Foundation, May 2017. URL: http://swfound.org/media/205874/swf_ssa_fact_sheet.pdf.
- [257] Gadfium. *English: Chinese Ship "Yuan Wang 2" in Waitemata Harbour, Auckland, New Zealand*. Oct. 27, 2005. URL: <https://commons.wikimedia.org/wiki/File:YuanWang2c.JPG> (visited on 10/05/2020).
- [258] David A Vallado and Jacob D Griesbach. “Simulating Space Surveillance Networks”. In: *AAS/AIAA Astrodynamics Specialist Conference*. AAS/AIAA Astrodynamics Specialist Conference. AIAA, 2011, pp. 2769–2788.
- [259] LEOLABS. *LEOLABS*. LEOLabs. 2020. URL: <https://www.leolabs.space/> (visited on 03/09/2020).
- [260] Numerica. *Space Domain Awareness*. 2020. URL: <https://www.numerica.us/space-defense/> (visited on 03/09/2020).
- [261] ExoAnalytic Solutions. *ExoAnalytic Solutions*. ExoAnalytics. 2020. URL: <https://exoanalytic.com/> (visited on 03/09/2020).
- [262] Felix R. Hoots and Ronald L. Roehrich. *Models for Propagation of NORAD Element Sets*. Aerospace Defense Command Peterson AFB CO Office of Astrodynamics, 1988. URL: <https://apps.dtic.mil/docs/citations/ADA093554> (visited on 04/28/2020).
- [263] Space Track. *SSA Sharing and Orbital Data Requests*. Space-Track.org. 2020. URL: <https://www.space-track.org/documentation#odr> (visited on 04/28/2020).
- [264] Saika Aida and Michael Kirschner. “Accuracy Assessment of SGP4 Orbit Information Conversion into Osculating Elements”. In: 723 (Aug. 1, 2013), p. 160. URL: <http://adsabs.harvard.edu/abs/2013ESASP.723E.160A> (visited on 04/28/2020).

- [265] Lasunncty. *English: Digram Illustrating and Explaining Various Terms in Relation to Orbits of Celestial Bodies*. Oct. 10, 2007. URL: <https://commons.wikimedia.org/wiki/File:Orbit1.svg> (visited on 04/28/2020).
- [266] C. Brandon Halstead. "The Ultimate High Ground - U.S. Intersector Cooperation in Outer Space". In: *Journal of Air Law and Commerce* 81.4 (2016), pp. 595–610. URL: <https://scholar.smu.edu/jalc/vol81/iss4/2/> (visited on 12/17/2018).
- [267] Kevin Pollpeter. "Space, the New Domain: Space Operations and Chinese Military Reforms". In: *Journal of Strategic Studies* 39.5-6 (Sept. 18, 2016), pp. 709–727.
- [268] David Koplow. "ASAT-Isfaction: Customary International Law and the Regulation of Anti-Satellite Weapons". In: *Michigan Journal of International Law* 30.4 (Jan. 1, 2009). URL: <https://heinonline.org/HOL/P?h=hein.journals/mjil30&i=1199>.
- [269] David E. Lupton. *On Space Warfare: A Space Power Doctrine*. 1st ed. Air University Press, 1988.
- [270] Dana J. St. James. "The Legality of Antisatellites Recent Development". In: *Boston College International and Comparative Law Review* 3.2 (1980), pp. 467–494. URL: <https://heinonline.org/HOL/P?h=hein.journals/bcic3&i=473> (visited on 12/17/2018).
- [271] Bhupendra Jasani and Christopher Lee. *Countdown to Space War*. 1st ed. Taylor & Francis, Jan. 1, 1984. 132 pp. Google Books: [bpYgAAAAMAAJ](#).
- [272] Steven J. Bruger. *Not Ready for the 'First Space War,' What About the Second*. Navar War College Department of Operations, May 17, 1993, pp. 1–35. URL: <https://apps.dtic.mil/docs/citations/ADA266557> (visited on 12/17/2018).
- [273] John E Hyten. "A Sea of Peace or a Theater of War? Dealing with the Inevitable Conflict in Space". In: *Air & Space Power Journal* 16.3 (2002), p. 78. URL: <https://go.gale.com/ps/anonymous?id=GALE%7CA94269862&issn=1555385X>.
- [274] Vishnu Anantatmula. "U.S. Initiative to Place Weapons in Space: The Catalyst for a Space-Based Arms Race with China and Russia". In: *Astropolitics* 11.3 (Sept. 1, 2013), pp. 132–155.
- [275] Mark A Gubrud. "Chinese and US Kinetic Energy Space Weapons and Arms Control". In: *Asian Perspective* 35.4 (2011), pp. 617–641. JSTOR: 42704774.
- [276] Ministry of External Affairs, Government of India. *Frequently Asked Questions on Mission Shakti, India's Anti-Satellite Missile Test Conducted on 27 March, 2019*. Mar. 27, 2019. URL: https://www.mea.gov.in/press-releases.htm?dtl/31179/Frequently_Asked_Questions_on_Mission_Shakti_Indias_AntiSatellite_Missile_test_conducted_on_27_March_2019.
- [277] Nina Tannenwald. "Law versus Power on the High Frontier: The Case for a Rule-Based Regime for Outer Space". In: *Yale Journal of International Law* 29.2 (2004), pp. 363–422. URL: <https://heinonline.org/HOL/P?h=hein.journals/yjil29&i=373> (visited on 12/17/2018).

- [278] Roger Handberg. “Is Space War Imminent? Exploring the Possibility”. In: *Comparative Strategy* 36.5 (Oct. 20, 2017), pp. 413–425.
- [279] Paul Stares. “Space and US National Security”. In: *Journal of Strategic Studies* 6.4 (Dec. 1, 1983), pp. 31–48.
- [280] Steven Freeland. “Peaceful Purposes - Governing the Military Uses of Outer Space”. In: *European Journal of Law Reform* 18 (2016), pp. 35–51.
- [281] Jennifer Ann Urban. “Soft Law: The Key to Security in a Globalized Outer Space”. In: *Transportation Law Journal* 43.1 (2016), pp. 33–50. URL: <https://heinonline.org/HOL/P?h=hein.journals/tportl43&i=39> (visited on 12/17/2018).
- [282] Paul Meyer. “Dark Forces Awaken: The Prospects for Cooperative Space Security”. In: *The Nonproliferation Review* 23.3-4 (July 3, 2016), pp. 495–503. URL: <https://www.tandfonline.com/doi/full/10.1080/10736700.2016.1268750> (visited on 12/17/2018).
- [283] Francis Grimal and Jae Sundaram. “The Incremental Militarization of Outer Space: A Threshold Analysis”. In: *Chinese Journal of International Law* 17.1 (Mar. 1, 2018), pp. 45–72.
- [284] Bates Gill and Martin Kleiber. “China’s Space Odyssey: What the Antisatellite Test Reveals about Decision-Making in Beijing”. In: *Foreign Affairs* 86.3 (2007), pp. 2–6. JSTOR: 20032344. URL: <https://www.foreignaffairs.com/articles/china/2007-05-01/chinas-space-odyssey-what-antisatellite-test-reveals-about-decision>.
- [285] James Moltz. *The Politics of Space Security: Strategic Restraint and the Pursuit of National Interests, Second Edition*. Redwood City, UNITED STATES: Stanford University Press, 2014. URL: <http://ebookcentral.proquest.com/lib/oxford/detail.action?docID=744004> (visited on 12/18/2018).
- [286] Zaeem Shabbir and Ali Sarosh. “Counterspace Operations and Nascent Space Powers”. In: *Astropolitics* 16.2 (Aug. 24, 2018), pp. 119–140.
- [287] Bonnie L. Triezenberg. *Deterring Space War: An Exploratory Analysis Incorporating Prospect Theory into a Game Theoretic Model of Space Warfare*. Product page. Santa Monica, CA: Rand Corporation, 2017. URL: https://www.rand.org/pubs/rgs_dissertations/RGSD400.html (visited on 12/18/2018).
- [288] Mark Levine. “Russia Tops List of Countries That Could Launch Cyberattacks on US”. In: *ABC News* (May 19, 2017). URL: <https://abcnews.go.com/US/russia-tops-list-100-countries-launch-cyberattacks-us/story?id=47487188> (visited on 12/18/2018).
- [289] Johan Sigholm. “Non-State Actors in Cyberspace Operations”. In: *Journal of Military Studies* 4.1 (2016), pp. 1–37.
- [290] Thomas Rid and Ben Buchanan. “Attributing Cyber Attacks”. In: *Journal of Strategic Studies* 38.1-2 (Jan. 2, 2015), pp. 4–37.

- [291] Jon R. Lindsay. “Tipping the Scales: The Attribution Problem and the Feasibility of Deterrence against Cyberattack”. In: *Journal of Cybersecurity* 1.1 (Sept. 1, 2015), pp. 53–67.
- [292] Nicholas Tsagourias. “Cyber Attacks, Self-Defence and the Problem of Attribution”. In: *Journal of Conflict and Security Law* 17.2 (July 1, 2012), pp. 229–244.
- [293] Brendan I. Koerner. “Inside the OPM Hack, the Cyberattack That Shocked the US Government”. In: *Wired* (Oct. 23, 2016). URL: <https://www.wired.com/2016/10/inside-cyberattack-shocked-us-government/> (visited on 12/18/2018).
- [294] Nirav Shah et al. “System of Systems Architecture: The Case of Space Situational Awareness”. In: *AIAA SPACE 2007 Conference & Exposition*. AIAA SPACE. AIAA SPACE Forum. American Institute of Aeronautics and Astronautics, Sept. 18, 2007.
- [295] W.J. Hennigan. “Russian Craft Shadowing U.S. Spy Satellite, Space Force Commander Says”. In: *Time* (Feb. 10, 2020). URL: <https://time.com/5779315/russian-spacecraft-spy-satellite-space-force/> (visited on 04/28/2020).
- [296] Bruce McClintock. *Space Safety Coordination: A Norm for All Nations*. RAND Corporation. Apr. 16, 2019. URL: <https://www.rand.org/blog/2019/04/space-safety-coordination-a-norm-for-all-nations.html> (visited on 03/09/2020).
- [297] Caleb Henry. “Space Situational Awareness Experts Urge Russia to Join Orbital Neighborhood Watch”. In: *Space News* (Mar. 16, 2018). URL: <https://spacenews.com/space-situational-awareness-experts-urge-russia-to-join-orbital-neighborhood-watch/> (visited on 03/09/2020).
- [298] ai-solutions. *FreeFlyer® Software*. Version 7.3 (Mission). 2018. URL: <https://ai-solutions.com/freeflyer/> (visited on 12/20/2018).
- [299] Frank Johnson. *1 Million Subscribers Connected: Iridium Helps Prevent Shark Attacks While Protecting Local Ecosystems*. Iridium Satellite Communications. June 18, 2018. URL: <https://www.iridium.com/blog/2018/06/18/1-million-subs-helping-prevent-shark-attacks/> (visited on 12/14/2018).
- [300] Peter Selding. “U.S. Defense Agency Encourages Allied Nations to Join Unlimited-Use Iridium Program”. In: *Space News* (Nov. 11, 2016). URL: <https://spacenews.com/u-s-defense-agency-encourages-allied-nations-to-join-unlimited-use-iridium-program/> (visited on 12/14/2018).
- [301] JFSCC. *SSA Sharing & Orbital Data Requests*. Space-Track.org. 2018. URL: <https://www.space-track.org/documentation#/odr> (visited on 12/10/2018).
- [302] James David. “Was It Really ‘Space Junk’? Us Intelligence Interest in Space Debris That Returned to Earth”. In: *Astropolitics* 3.1 (Apr. 1, 2005), pp. 43–65.
- [303] The Guardian. “What Is Object 2014-28E – a Russian Military Satellite or a Piece of Unidentified Debris?”. In: *The Guardian. Science* (Nov. 18, 2014). URL: <https://www.theguardian.com/science/shortcuts/2014/nov/18/object-2014-28e-space-russian-satellite-unidentified> (visited on 03/10/2020).

- [304] The Telegraph. “Russia ’Busts Foreign Satellite Spy Ring”. In: *The Telegraph* (Apr. 12, 2015). URL: <https://www.telegraph.co.uk/news/worldnews/europe/russia/11531013/Russia-busts-foreign-satellite-spy-ring.html> (visited on 10/02/2020).
- [305] Joe Pappalardo. *America’s Next Spy Satellites Will Disappear. Here’s How*. Popular Mechanics. Nov. 30, 2018. URL: <https://www.popularmechanics.com/military/research/a25349950/nro-satellite-space-junk/> (visited on 03/10/2020).
- [306] Theresa Hitchens. “New Satellite Imagery Rules Hover In Interagency Limbo”. In: *Breaking Defense* (2020). URL: <https://breakingdefense.com/2020/03/new-satellite-imagery-rules-hover-in-interagency-limbo/> (visited on 10/02/2020).
- [307] T. S. Kelso. *CelesTrak: “FAQs: Two-Line Element Set Format”*. CelesTrak. Dec. 28, 2019. URL: <http://celestrak.com/columns/v04n03/> (visited on 04/28/2020).
- [308] Thomas M Johnson. “SSA Sensor Calibration Best Practices”. In: Advanced Maui Optical and Space Surveillance Technologies Conference. Maui, Hawaii, 2015. URL: <https://amostech.com/TechnicalPapers/2015/Poster/JohnsonT.pdf> (visited on 10/06/2020).
- [309] Forrest Gasdia. “Optical Tracking and Spectral Characterization of Cubesats for Operational Missions”. PhD Dissertations and Master’s Theses. Embry-Riddle Aeronautical University, May 1, 2016. URL: <https://commons.erau.edu/edt/212>.
- [310] Peerapong Torteeka et al. “Enhancing the Capability of a Ground-Based Optical Telescope for Thai National Space Objects Observation”. In: *Proceedings of Innovation Aviation & Aerospace Industry - International Conference 2020* 39.1 (2020), p. 15.
- [311] J. Puig-Suari, C. Turner, and W. Ahlgren. “Development of the Standard CubeSat Deployer and a CubeSat Class PicoSatellite”. In: *2001 IEEE Aerospace Conference Proceedings*. 2001 IEEE Aerospace Conference Proceedings (Cat. No.01TH8542). Vol. 1. Big Sky, MT, USA: IEEE, Mar. 2001, 1/347–1/353 vol.1.
- [312] S. Jahirabadkar et al. “Space Objects Classification Techniques: A Survey”. In: *2020 International Conference on Computational Performance Evaluation (ComPE)*. 2020 International Conference on Computational Performance Evaluation (ComPE). Shillong, India: IEEE, July 2020, pp. 786–791.
- [313] M. Nayak, J. Beck, and B. Udrea. “Real-Time Attitude Commanding to Detect Coverage Gaps and Generate High Resolution Point Clouds for RSO Shape Characterization with a Laser Rangefinder”. In: *2013 IEEE Aerospace Conference*. 2013 IEEE Aerospace Conference. Big Sky, MT, USA: IEEE, Mar. 2013, pp. 1–14.
- [314] Bin Li et al. “A Machine Learning-Based Approach for Improved Orbit Predictions of LEO Space Debris With Sparse Tracking Data From a Single Station”. In: *IEEE Transactions on Aerospace and Electronic Systems* 56.6 (Apr. 2020), pp. 4253–4268.

- [315] Roya Afshar and Shuai Lu. “Classification and Recognition of Space Debris and Its Pose Estimation Based on Deep Learning of CNNs”. In: *HCI International 2020 - Posters*. International Conference on Human-Computer Interaction. Vol. I. Copenhagen, Denmark: Springer, July 2020, pp. 605–613.
- [316] Jiangbo Xi et al. “Space Debris Detection Using Feature Learning of Candidate Regions in Optical Image Sequences”. In: *IEEE Access* 8 (2020), pp. 150864–150877.
- [317] R. Furfaro et al. “Space Debris Identification and Characterization via Deep Meta-Learning”. In: *First Int’l. Orbital Debris Conf. (2019)*. First International Orbital Debris Conference. Dec. 1, 2019. URL: <https://www.hou.usra.edu/meetings/orbitaldebris2019/orbital2019paper/pdf/6123.pdf> (visited on 10/12/2020).
- [318] Bin Liu, Li Yao, and Dapeng Han. “Harnessing Ontology and Machine Learning for RSO Classification”. In: *SpringerPlus* 5.1 (Sept. 26, 2016). pmid: 27730017.
- [319] M. Krüger. “Detection of AIS Spoofing in Fishery Scenarios”. In: *2019 22th International Conference on Information Fusion (FUSION)*. Ottawa, ON, Canada: IEEE, July 2019, pp. 1–7. URL: <https://ieeexplore.ieee.org/document/9011328>.
- [320] Diogo Barradas, Nuno Santos, and Luís Rodrigues. “Effective Detection of Multimedia Protocol Tunneling Using Machine Learning”. In: *Proceedings of the 27th USENIX Security Symposium*. 27th {USENIX} Security Symposium ({USENIX} Security 18). 2018, pp. 169–185.
- [321] Bandar Alotaibi and Khaled Elleithy. “A New MAC Address Spoofing Detection Technique Based on Random Forests”. In: *Sensors* 16.3 (Feb. 24, 2016), p. 281.
- [322] scikit-learn. *DecisionTreeClassifier*. Scikit Learn Documentaiton. URL: <https://scikit-learn.org/stable/modules/generated/sklearn.tree.DecisionTreeClassifier.html>.
- [323] scikit-learn. *BaggingClassifier*. Scikit Learn Documentaiton. URL: <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.BaggingClassifier.html>.
- [324] scikit-learn. *RandomForest*. Scikit Learn Documentaiton. URL: <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html>.
- [325] scikit-learn. *HistGradientBoostingClassifier*. Scikit Learn Documentaiton. URL: <https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.HistGradientBoostingClassifier.html>.
- [326] André Altmann et al. “Permutation Importance: A Corrected Feature Importance Measure”. In: *Bioinformatics* 26.10 (May 15, 2010), pp. 1340–1347.
- [327] Space Track. *User Agreement*. Space-Track.org. Aug. 1, 2019. URL: https://www.space-track.org/documentation#user_agree (visited on 04/28/2020).
- [328] Erik Kulu. *Cubesat Tables*. Nanosats Database. URL: <https://www.nanosats.eu/tables.html> (visited on 09/24/2020).

- [329] EnduroSat. *1U CubeSat Platform Cubesat Platforms*. CubeSat by EnduroSat. 2020. URL: <https://www.endurosat.com/cubesat-store/all-cubesat-modules/1u-cubesat-platform/> (visited on 09/24/2020).
- [330] Nanoracks. *Nanoracks*. Nanoracks. URL: <https://nanoracks.com/products/iss-deployment/> (visited on 09/24/2020).
- [331] Innovative Solutions in Space. *ISIS ISIPOD 3-Unit CubeSat Deployer*. CubeSatShop.com. URL: <https://www.cubesatshop.com/product/3-unit-cubesat-deployer/> (visited on 09/24/2020).
- [332] Clyde C. Helms. “A Survey of Launch Services 2016-2020”. In: *AIAA Propulsion and Energy 2020 Forum*. AIAA Propulsion and Energy 2020 Forum. Virtual: American Institute of Aeronautics and Astronautics, Aug. 2020.
- [333] eoPortal. *Vega PoC Flight for SSMS*. eoPortalDirectory. 2020. URL: <https://directory.eoportal.org/web/eoportal/satellite-missions/v-w-x-y-z/vega-ssms> (visited on 12/02/2020).
- [334] Northrop Grumman. *Minotaur Rocket*. Northrop Grumman Space. URL: <https://www.northropgrumman.com/space/minotaur-rocket> (visited on 12/02/2020).
- [335] William Graham. *Russia’s Rokot Vehicle Successfully Launches Geo-IK-2 Satellite*. NASASpaceFlight.com. Aug. 30, 2019. URL: <https://www.nasaspaceflight.com/2019/08/russias-rokot-geo-ik-2-satellite/> (visited on 12/03/2020).
- [336] Philipp Olbrich and David Shim. “Symbolic Practices of Legitimation: Exploring Domestic Motives of North Korea’s Space Program”. In: *International Relations of the Asia-Pacific* 19.1 (Jan. 1, 2019), pp. 33–61.
- [337] Defense Intelligence Agency. *Challenges to Security in Space*. Jan. 2019. URL: https://www.dia.mil/Portals/27/Documents/News/Military%20Power%20Publications/Space_Threat_V14_020119_sm.pdf.
- [338] J.C. Mauduit. “Collaboration around the International Space Station: Science for Diplomacy and Its Implication for US-Russia and China Relations”. In: *7th Annual SAIS Asia Conference (SAIS 2018)*. 7th Annual SAIS Asia Conference (SAIS 2018). Washington, DC: Secure World Foundation. URL: <https://discovery.ucl.ac.uk/id/eprint/10083727/>.
- [339] Michael Byers. “Cold, Dark, and Dangerous: International Cooperation in the Arctic and Space”. In: *Polar Record* 55.1 (June 2019), pp. 32–47.
- [340] Sheng-Chih Wang. “The Making of New ‘Space’: Cases of Transatlantic Astropolitics”. In: *Geopolitics* 14.3 (Aug. 21, 2009), pp. 433–461.
- [341] Elizabeth L. Chalecki. “Knowledge in Sheep’s Clothing: How Science Informs American Diplomacy”. In: *Diplomacy & Statecraft* 19.1 (Mar. 13, 2008), pp. 1–19.
- [342] Christian Davenport. “Implication of Sabotage Adds Intrigue to SpaceX Investigation”. In: *The Washington Post* (Sept. 30, 2016). URL: https://www.washingtonpost.com/business/economy/implication-of-sabotage-adds-intrigue-to-spacex-investigation/2016/09/30/5bb60514-874c-11e6-a3ef-f35afb41797f_story.html (visited on 12/11/2020).

- [343] RT International. *Sabotage Considered in Proton Rocket Crash – Investigator*. RT International. May 29, 2014. URL: <https://www.rt.com/news/162228-proton-rocket-failure-sabotage/> (visited on 12/11/2020).
- [344] Fred Weir. “Russia Hints Foreign Sabotage May Be behind Space Program Troubles”. In: *Christian Science Monitor* (Jan. 10, 2012). URL: <https://www.csmonitor.com/World/Global-News/2012/0110/Russia-hints-foreign-sabotage-may-be-behind-space-program-troubles> (visited on 12/11/2020).
- [345] Debra Werner. “Small Satellite Sector Grapples with Cybersecurity Requirements, Cost”. In: *Space News* (Aug. 8, 2018). URL: <https://spacenews.com/small-satellite-sector-grapples-with-cybersecurity-requirements-cost/> (visited on 09/21/2020).
- [346] K. W. Ingols and R. W. Skowrya. *Guidelines for Secure Small Satellite Design and Implementation: FY18 Cyber Security Line-Supported Program*. MIT Lincoln Laboratory Lexington United States, Feb. 6, 2019. URL: <https://apps.dtic.mil/sti/citations/AD1099003> (visited on 09/24/2020).
- [347] M. Langer and J. Bouwmeester. “Reliability of CubeSats – Statistical Data, Developers’ Beliefs and the Way Forward”. In: *Proceedings of the 30th Annual AIAA/USU Conference on Small Satellites*. AIAA/USU Conference on Small Satellites. AIAA, 2016. URL: <http://resolver.tudelft.nl/uuid:4c6668ff-c994-467f-a6de-6518f209962e> (visited on 09/23/2020).
- [348] Michael Swartwout. “You Say “Picosat”, I Say “CubeSat”: Developing a Better Taxonomy for Secondary Spacecraft”. In: *2018 IEEE Aerospace Conference*. 2018 IEEE Aerospace Conference. Big Sky, MT, USA: IEEE, Mar. 2018, pp. 1–17.
- [349] NASA. *CubeSat 101: Basic Concepts and Processes for First-Time CubeSat Developers*. Oct. 2017. URL: https://www.nasa.gov/sites/default/files/atoms/files/nasa_csli_cubesat_101_508.pdf.
- [350] Cal Poly SLO. *CubeSat Design Specification (CDS) REV 13*. 2014. URL: https://blogs.esa.int/philab/files/2019/11/RD-02_CubeSat_Design_Specification_Rev._13_The.pdf.
- [351] HQ AFSPC/SEK. *Air Force Space Command Manual 91-710, Volume 3*. May 15, 2019. URL: <https://static.e-publishing.af.mil/production/1/afspc/publication/afspcman91-710v3/afspcman91-710v3.pdf>.
- [352] Christine Gebara and David Spencer. “Verification and Validation Methods for the Prox-1 Mission”. 30th Annual AIAA/USU Conference on Small Satellites: Frank J. Redd Student Competition. Aug. 10, 2016. URL: <https://digitalcommons.usu.edu/smallsat/2016/TS8StudentComp/3>.
- [353] *AMSAT Fox-1 DITL Test*. In collab. with Jerry Buxton. Jan. 17, 2015. URL: <https://www.youtube.com/watch?v=TjGAYvMyz4Q> (visited on 12/31/2020).
- [354] Department of Defense. *MIL-STD-882E*. May 11, 2012. URL: <https://www.dau.edu/cop/armyesoh/DAU%20Sponsored%20Documents/MIL-STD-882E.pdf>.

- [355] Gary L. Prater. “NPSAT1 Missile System Pre-Launch Safety Package (MSPSP)”. Thesis. Naval Postgraduate School, June 1, 2004. URL: <https://apps.dtic.mil/sti/citations/ADA424941> (visited on 01/01/2021).
- [356] SpaceX. *Falcon User’s Guide*. Apr. 2020. URL: https://www.spacex.com/media/falcon_users_guide_042020.pdf.
- [357] United Launch Alliance. *Delta IV Launch Services User’s Guide*. June 2013. URL: <https://www.ulalaunch.com/docs/default-source/rockets/delta-iv-user’s-guide.pdf>.
- [358] Lisa Valencia. “Autonomous Flight Termination System (AFTS)”. 2019. URL: <https://www.gps.gov/cgsic/meetings/2019/valencia.pdf>.
- [359] FCC.report. *Space Exploration Technologies Corp. (SpaceX) Experimental License FCC Filings*. FCC.report: Database Report/Search Tool for FCC Information. 2020. URL: <https://fcc.report/ELS/Space-Exploration-Technologies-Corp-SpaceX> (visited on 01/02/2021).
- [360] Space Micro. *μSDR-C Software Defined Radio*. Space Micro. 2019. URL: spacemicro.com/products/communication-systems/%CE%BCSDR-C%E2%84%A2%20SOFTWARE%20DEFINED%20RADIO.pdf.
- [361] Flexitech Aerospace. *Satellite Communication Systems Products*. Flexitech Aerospace. URL: <https://flexitechaerospace.com/products/> (visited on 01/02/2021).
- [362] CubeSatShop. *Helios Deployable Antenna*. CubeSatShop.com. URL: <https://www.cubesatshop.com/product/helios-deployable-antenna/> (visited on 01/02/2021).
- [363] Ian F. Akyildiz, Josep M. Jornet, and Shuai Nie. “A New CubeSat Design with Reconfigurable Multi-Band Radios for Dynamic Spectrum Satellite Communication Networks”. In: *Ad Hoc Networks* 86 (Apr. 1, 2019), pp. 166–178.
- [364] Endrit Shehaj et al. “GPS Based Navigation Performance Analysis within and beyond the Space Service Volume for Different Transmitters’ Antenna Patterns”. In: *Aerospace* 4.3 (Sept. 2017), p. 44.
- [365] FCC. *GPS L1 Link Budget*. URL: <https://apps.fcc.gov/els/GetAtt.html?id=110032&x=..>
- [366] US Department of Transportation. *Global Positioning System (GPS) Adjacent Band Compatibility Assessment*. Dec. 2017. URL: <https://www.transportation.gov/sites/dot.gov/files/docs/subdoc/186/dot-gps-adjacent-band-final-report.pdf>.
- [367] W Lan et al. *Poly Picosatellite Orbital Deployer Mk. III Rev. E User Guide*. 2014. URL: https://static1.squarespace.com/static/5418c831e4b0fa4ecac1bacd/t/5806854d6b8f5b8eb57b83bd/1476822350599/P-POD_MkIIIRevE_UserGuide_CP-PPODUG-1.0-1_Rev1.pdf.
- [368] MathWorks. *RF Propagation - MATLAB & Simulink*. Version R2020b. 2020. URL: https://uk.mathworks.com/help/antenna/rf-propagation.html?s_tid=CRUX_lftnav (visited on 01/04/2021).

- [369] Inside GNSS. *Measuring GNSS Signal Strength*. Inside GNSS - Global Navigation Satellite Systems Engineering, Policy, and Design. Dec. 2, 2010. URL: <https://insidegnss.com/measuring-gnss-signal-strength/>.
- [370] Andrew Brierley-Green. “Global Navigation Satellite System Fundamentals and Recent Advances in Receiver Design”. IEEE Long Island Section. Sept. 2017. URL: https://www.ieee.li/pdf/viewgraphs/gnss_fundamentals.pdf.
- [371] G. X. Gao et al. “Protecting GNSS Receivers From Jamming and Interference”. In: *Proceedings of the IEEE* 104.6 (June 2016), pp. 1327–1338.
- [372] Tom Shales. “Cable’s ‘Captain Midnight’ Apprehended”. In: *The Washington Post* (July 23, 1986), p. 2.
- [373] Washington Post. “Doors Fail When Reagan Is Home”. In: *The Washington Post* (Apr. 5, 1986), p. 1.
- [374] Mark Stein. “Elbowing for a Piece of Space : The Parking Lot for Satellites Is Getting Jammed.” In: *LA Times* (Sept. 20, 1993). URL: <https://www.latimes.com/archives/la-xpm-1993-09-20-mn-37280-story.html> (visited on 08/17/2020).
- [375] Nora Boustany. “Kurdish TV Gets Static from Turks”. In: *The Washington Post* (Nov. 25, 1998). URL: <https://www.washingtonpost.com/archive/politics/1998/11/25/kurdish-tv-gets-static-from-turks/f41ee658-8258-495a-bc78-b644c2bcf56d/> (visited on 02/06/2019).
- [376] Stephen Kinzer. “Kurds Are Determined to Restore TV Station Shut by the British”. In: *New York Times. INTERNATIONAL* (1999), A13. URL: <https://search.proquest.com/docview/110109073/abstract/FE7E2E2363E74DF2PQ/2> (visited on 08/17/2020).
- [377] Amir Hassanpour. “Satellite Footprints as National Border: MED-TV and the Extraterritoriality of State Sovereignty”. In: *Journal of Muslim Minority Affairs* 18.1 (Apr. 1998), pp. 53–72. URL: <https://www.tandfonline.com/doi/abs/10.1080/13602009808716393?journalCode=cjmm20>.
- [378] James Glave. “Have Crackers Found Military’s Achilles’ Heel?” In: *Wired* (Apr. 21, 1998). URL: <https://www.wired.com/1998/04/have-crackers-found-militarys-achilles-heel/> (visited on 02/06/2019).
- [379] BBC. *Satellite Hijack ‘Impossible’*. BBC News. Mar. 2, 1999. URL: <http://news.bbc.co.uk/1/hi/sci/tech/288965.stm> (visited on 02/06/2019).
- [380] Lester Grau. “GPS Signals Jammed During Tank Trials”. In: *Military Review* (Mar. 2001). URL: <https://www.hsdl.org/?view&did=3694>.
- [381] Sam Costello. *Suspect Arrested in NASA Hack*. Computerworld. Sept. 25, 2000. URL: <http://www2.computerworld.com.au/article/78798/> (visited on 02/07/2019).
- [382] Paul McNulty. *United States of America v. Gary McKinnon*. Indictment. Nov. 2002.

- [383] Gary McKinnon. “Theresa May Saved My Life – Now She’s the Only Hope for the Human Rights Act | Gary McKinnon”. In: *The Guardian. Opinion* (Nov. 15, 2016). URL: <https://www.theguardian.com/commentisfree/2016/nov/15/theresa-may-saved-my-life-human-rights-act> (visited on 08/17/2020).
- [384] John Schwartz. “Compressed Data; Hacker Obtains Shuttle Design Files, Baffling NASA”. In: *The New York Times. Business* (Aug. 12, 2002). URL: <https://www.nytimes.com/2002/08/12/business/compressed-data-hacker-obtains-shuttle-design-files-baffling-nasa.html> (visited on 08/17/2020).
- [385] Bento Miguel Ribeiro Martins. “GNSS Vulnerabilites and Robustness”. Universidade Do Porto, 2014. URL: <https://core.ac.uk/download/pdf/143409373.pdf>.
- [386] Martyn Thomas, ed. *Global Naviation Space Systems: Reliance and Vulnerabilities*. The Royal Academy of Engineering, Mar. 2011. URL: https://rntfnd.org/wp-content/uploads/2013/09/Royal-Acad-of-Eng_Global_Navigation_Space-Systems_Report.pdf.
- [387] Mark Urban. “Enthusiast Watches Nato Spy Pictures”. In: (June 13, 2002). URL: <http://news.bbc.co.uk/1/hi/programmes/newsnight/2041754.stm> (visited on 08/17/2020).
- [388] Robert Windrem. *U.S. Satellite Feeds to Iran Jammed*. NBC News. Oct. 24, 2003. URL: <http://www.nbcnews.com/id/3340692/t/us-satellite-feeds-iran-jammed/> (visited on 02/06/2019).
- [389] Marlyn Kemper Littman. “Satellite Network Security”. In: *Encyclopedia of Information Science and Technology, Second Edition* (2009), pp. 3350–3355.
- [390] David Hencke and Owen Gibson. “Protest to Libya after Satellites Jammed”. In: *The Guardian. UK news* (Dec. 3, 2005). URL: <https://www.theguardian.com/uk/2005/dec/03/politics.libya> (visited on 02/06/2019).
- [391] Charles Q. Choi. *Libya Pinpointed as Source of Months-Long Satellite Jamming in 2006*. Space.com. Apr. 9, 2007. URL: <https://www.space.com/3666-libya-pinpointed-source-months-long-satellite-jamming-2006.html> (visited on 02/06/2019).
- [392] Kyle Spector. *Hacking Hezbollah*. Foreign Policy. Aug. 3, 2006. URL: <https://foreignpolicy.com/2006/08/03/hacking-hezbollah/> (visited on 02/06/2019).
- [393] Wired. “The Great Brazilian Sat-Hack Crackdown”. In: *Wired* (Apr. 20, 2009). URL: <https://www.wired.com/2009/04/fleetcom/> (visited on 02/06/2019).
- [394] Associated Press. *Glitch Shows How Much US Military Relies on GPS*. Phys.Org. June 1, 2010. URL: <https://phys.org/news/2010-06-glitch-military-gps.html> (visited on 02/06/2019).
- [395] BBC Monitoring Asia Pacific. “South Korean Satellite Comes under North Jamming Attack”. In: *BBC Worldwide Monitoring* (Nov. 16, 2012).

- [396] UT Austin Press Office. *UT Austin Researchers Successfully Spoof an \$80 Million Yacht at Sea*. UT News. July 29, 2013. URL: <https://news.utexas.edu/2013/07/29/ut-austin-researchers-successfully-spoof-an-80-million-yacht-at-sea/> (visited on 02/06/2019).
- [397] Al Jazeera. “Egypt Jamming Al Jazeera’s Satellite Signals”. In: *Al Jazeera* (Sept. 4, 2013). URL: <https://www.aljazeera.com/features/2013/9/4/egypt-jamming-al-jazeeras-satellite-signals>.
- [398] CBS New York. “N.J. Man In A Jam, After Illegal GPS Device Interferes With Newark Airport”. In: (Aug. 9, 2013). URL: <https://newyork.cbslocal.com/2013/08/09/n-j-man-in-a-jam-after-illegal-gps-device-interferes-with-newark-liberty-operations/> (visited on 08/17/2020).
- [399] Chris Forrester. “Thailand Suffers Satellite Jamming”. In: *Inside Satellite TV* (Jan. 20, 2014). URL: <https://advanced-television.com/2014/01/20/thailand-suffers-satellite-jamming/>.
- [400] AP. “Network: Signal Jammed in Egypt during Comedy Show”. In: (Mar. 8, 2014).
- [401] ICT Monitor. “MBC Hit by Libya Satellite Jamming”. In: (Oct. 3, 2014).
- [402] Der Spiegel. *DLR Mit Trojanern von Geheimdienst Ausgespäht*. Spiegel. Apr. 13, 2014. URL: <https://www.spiegel.de/netzwelt/web/dlr-mit-trojanern-von-geheimdienst-ausgespaecht-a-964099.html> (visited on 08/17/2020).
- [403] Arabsat. *Arabsat Is Subject to Jamming and Its Engineers Succeed in Locating Its Source*. May 29, 2014. URL: <https://www.arabsat.com/NewsDetails.aspx?pageid=428&lang=2#:~:text=Arabsat%20is%20subject%20to%20Jamming%20and%20its%20Engineers%20succeed%20in%20locating%20its%20source,-Partager&text=Arab%20Satellite%20Communication%20Organization%20%E2%80%93%20ARABSAT,past%20week%20up%20to%20today>. (visited on 08/17/2020).
- [404] *Iridium Hacking*. In collab. with Sec and Schneider. Chaos Communication Camp 2015, Aug. 2015. URL: <https://www.youtube.com/watch?v=ahZOGhV8qnc> (visited on 08/17/2020).
- [405] James Dean. “Briton Held after Cyberattacker Tells Pentagon: We Control Your Satellites”. In: *The Times* (Mar. 7, 2015). URL: <https://www.thetimes.co.uk/article/briton-held-after-cyberattacker-tells-pentagon-we-control-your-satellites-fdfpqhxjx2z> (visited on 02/07/2019).
- [406] BBC. “N Korea ‘jamming GPS Signals’ in South”. In: *BBC News. Asia* (Apr. 1, 2016). URL: <https://www.bbc.com/news/world-asia-35940542> (visited on 02/08/2019).
- [407] BBC Monitoring Kiev Unit. “Ukrainian TV Channels Report Targeted Signal Jamming”. In: *BBC Worldwide Monitoring* (July 18, 2016).
- [408] Yasser Okbi. “Hackers Take over Israeli News Broadcast, Post ‘Allahu Akbar’”. In: *The Jerusalem Post* (Nov. 29, 2016).

- [409] Adam Ali Zare Hudaib. "Satellite Network Hacking & Security Analysis". In: *International Journal of Computer Science and Security (IJCSS)* 10.1 (2016), p. 8. URL: <https://www.cscjournals.org/manuscript/Journals/IJCSS/Volume10/Issue1/IJCSS-1200.pdf>.
- [410] Sam Chambers. *Ship's Satellite Communication System Hacked with Ease*. Splash 247. July 19, 2017. URL: <https://splash247.com/ships-satellite-communication-system-hacked-ease/> (visited on 02/08/2019).
- [411] The Times. "Russia 'Disrupted Nato Wargames by Jamming GPS'". In: *The Times* (Nov. 13, 2018). URL: <https://www.thetimes.co.uk/article/russia-put-flights-at-risk-by-jamming-alliance-wargames-zttlmmknh>.
- [412] *Hacking Yachts Remotely via Satcom or Maritime Internet Router*. In collab. with Stephen Gerling. Dec. 3, 2018. URL: https://www.youtube.com/watch?v=mT7dXJ_ob8k&t= (visited on 08/18/2020).
- [413] David E. Sanger and William J. Broad. "U.S. Revives Secret Program to Sabotage Iranian Missiles and Rockets". In: *The New York Times. U.S.* (Feb. 13, 2019). URL: <https://www.nytimes.com/2019/02/13/us/politics/iran-missile-launch-failures.html> (visited on 08/17/2020).
- [414] Zak Doffman. "Crashed UAE Military Spy Satellite Raises Possibility Of Enemy Cyberattack". In: *Forbes* (July 2019). URL: <https://www.forbes.com/sites/zakdoffman/2019/07/12/did-an-iranian-cyberattack-force-a-military-spy-satellite-to-drop-from-the-sky/> (visited on 08/17/2020).
- [415] Robin McKie. "Nasa Astronaut 'Accessed Ex-Partner's Bank Account from Space Station'". In: *The Guardian. US news* (Aug. 24, 2019). URL: <https://www.theguardian.com/us-news/2019/aug/24/nasa-astronaut-allegedly-accessed-ex-partners-bank-account-while-living-on-iss> (visited on 08/17/2020).
- [416] Elisha Fieldstadt. *Woman Who Accused NASA Astronaut Wife of Hacking Bank Account Charged with False Allegations*. NBC News. Apr. 2020. URL: <https://www.nbcnews.com/news/us-news/woman-who-accused-nasa-astronaut-wife-hacking-bank-account-charged-n1178611> (visited on 08/17/2020).
- [417] Victor Murray. *Legal GNSS Spoofing and Its Effects on Autonomous Vehicles*. Black Hat USA 2019. 2019. URL: <https://www.blackhat.com/us-19/briefings/schedule/#legal-gnss-spoofing-and-its-effects-on-autonomous-vehicles-15497> (visited on 08/17/2020).
- [418] Jay Mazoomdaar. "Not Only Kudankulam, ISRO, Too, Was Alerted of Cyber Security Breach". In: *Indian Express* (Nov. 6, 2019). URL: <https://indianexpress.com/article/india/not-only-kudankulam-isro-too-was-alerted-of-cyber-security-breach-6105184/> (visited on 08/17/2020).