

A Tale of Sea and Sky

On the Security of Maritime VSAT Communications

James Pavur*, Daniel Moser†, Martin Strohmeier†, Vincent Lenders† and Ivan Martinovic*

*Oxford University

Email: *first.last@cs.ox.ac.uk*

†armasuisse

Email: *first.last@armasuisse.ch*

Abstract—Very Small Aperture Terminals (VSAT) have revolutionized maritime operations. However, the security dimensions of maritime VSAT services are not well understood. Historically, high equipment costs have acted as a barrier to entry for both researchers and attackers. In this paper we demonstrate a substantial change in threat model, proving practical attacks against maritime VSAT networks with less than \$400 of widely-available television equipment. This is achieved through GSExtract, a purpose-built forensic tool which enables the extraction of IP traffic from highly corrupted VSAT data streams.

The implications of this threat are assessed experimentally through the analysis of more than 1.3 TB of real-world maritime VSAT recordings encompassing 26 million square kilometers of coverage area. The underlying network platform employed in these systems is representative of more than 60% of the global maritime VSAT services market. We find that sensitive data belonging to some of the world's largest maritime companies is regularly leaked over VSAT ship-to-shore communications. This threat is contextualized through illustrative case studies ranging from the interception and alteration of navigational charts to theft of passport and credit card details. Beyond this, we demonstrate the ability to arbitrarily intercept and modify TCP sessions under certain network configurations, enabling man-in-the-middle and denial of service attacks against ships at sea. The paper concludes with a brief discussion of the unique requirements and challenges for encryption in VSAT environments.

I. INTRODUCTION

The maritime transportation industry has trended towards ever-larger vessels operated by ever-smaller crews, a change driven 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 [9], [10]. 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 [45]. Moreover, the use of computing technology for marine operations is expected to grow for the foreseeable future; perhaps even progressing to fully autonomous vessels [5].

One of the critical drivers of this digitization revolution has been improvements in ship-to-shore communications. Through terrestrial and space-based radio transmissions, landside 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 paper makes an initial contribution towards understanding and securing these increasingly critical linkages.

Specifically, the paper 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 paper concludes with a brief discussion of both immediate and long-term technical improvements which may address these issues.



Fig. 1: A typical marine VSAT system [14].

II. RELATED WORK

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 cybersecurity 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 Navigational System (ECDIS) [12]. 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 [3]. 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 [6]. Beyond maritime, a much wider body of research surrounding the general topic of GPS spoofing and countermeasures exists [40].

With regards to 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 [33], [30]. 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 [2].

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 [27]. Furthermore, given the high value of typical cargo payloads (on the scale hundreds of millions of dollars), information systems aboard ever more automated freight vessels may become targets of pirates [26], [20]. 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 [42], [28]. Moreover, given that no state has claimed responsibility for the act, 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” [29]. 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” [29].

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 [38], [39]. 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. Peripherally relevant research into the general security of satellite broadband exists as well. However, this research focuses on the MPEG-TS encoding method widely used for terrestrial satellite broadband services and not on the newer standards which tend to be used in specialized marine systems [1], [35]. Given that significant security issues have been found in specialized multi-purpose data links in other transportation sectors — such as aviation — a closer look at VSAT radio signals is likely warranted [41].

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 traditional satellite broadband [16]. While many open source and freely available tools exist for interpreting MPEG-TS recordings, to our knowledge no comparable software exists for GSE [11], [8]. Additionally, the equipment sold to maritime VSAT customers to receive and interpret these signals (such as the system in Figure 1) can cost upwards of \$50,000 [21]. These high costs act as a significant barrier to entry for researchers.

III. BACKGROUND

A. Uses for Maritime VSAT

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 [47], [31].

There are, however, several common use cases for VSAT connectivity with broad applicability [21]. 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 [44], [21]. Finally, VSAT connectivity enables critical safety and navigational aids ranging from remote medical support to up-to-date navigational charts [21].

B. VSAT Network Architectures

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 [23]. 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 paper 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 [22].

In this paper, 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.

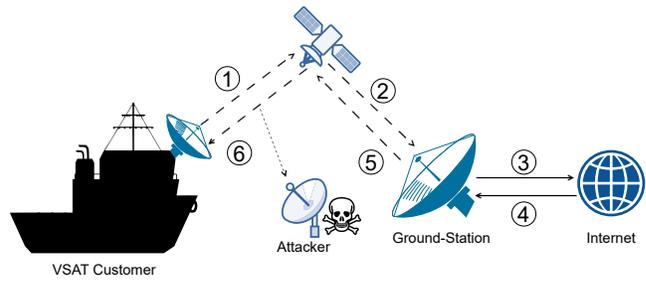


Fig. 2: The typical flow of data through a maritime VSAT network. The attacker in the diagram can eavesdrop on traffic from step 6 but has limited visibility into traffic at all other stages.

A maritime VSAT network is not significantly different from other satellite networking environments with respect to its basic architecture. As outlined in Figure 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 2 can easily observe responses from the satellite internet service provider (ISP) to the customer but would have 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.

IV. EXPERIMENTAL DESIGN

A. 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, Nordic and Mediterranean regions. An approximate map of the signal footprints involved in our research can be seen in Figure 3.

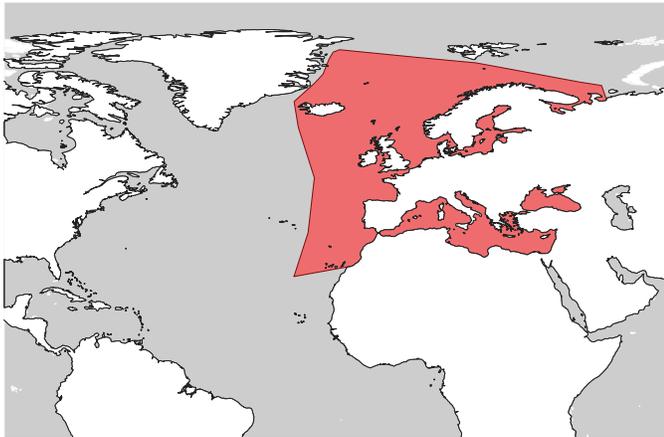


Fig. 3: Signal Coverage Footprint. Traffic from ships across the entire shaded area (more than 26 million square kilometers) was observable from our collection site in Europe.

As mentioned in section II, 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 a business-to-business or “VSAT as a service” model (generally 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 I).

In our specific experimental setup, our equipment was capable of receiving DVB-S2 signals in the Ku-band frequency range (10.7-12.75GHz). 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

TABLE I: Experimental Equipment

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

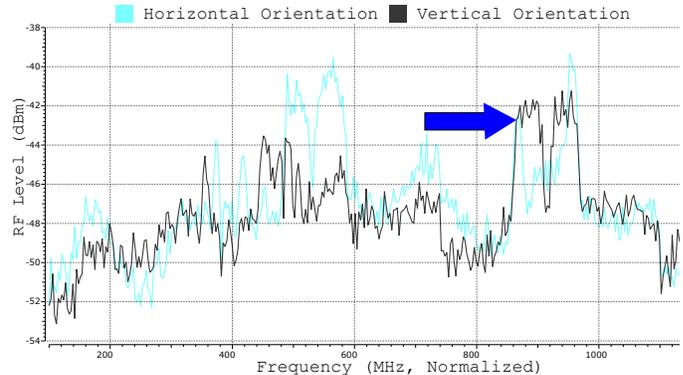


Fig. 4: Scanning for satellite streams across the Ku-band in two orientations. Distinct humps in the spectrum represent channels for potential analysis. NB: To maintain platform anonymity, the lower axis has been normalized.

noting that our research is 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 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 4). For this experiment, we identified a total of 15 VSAT streams on two geostationary platforms, mostly on the basis of signal modulation settings (e.g. 32-APSK) and strings detected in raw signal recordings.

B. 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 [15], [16], [17]. An overview of this encapsulation method can be found in Figure 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

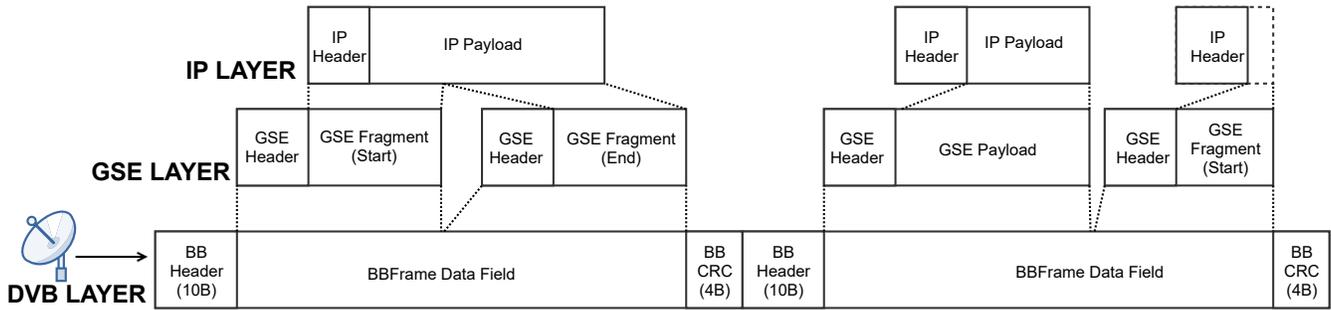


Fig. 5: A simplified overview of protocol layers which comprise maritime VSAT streams.

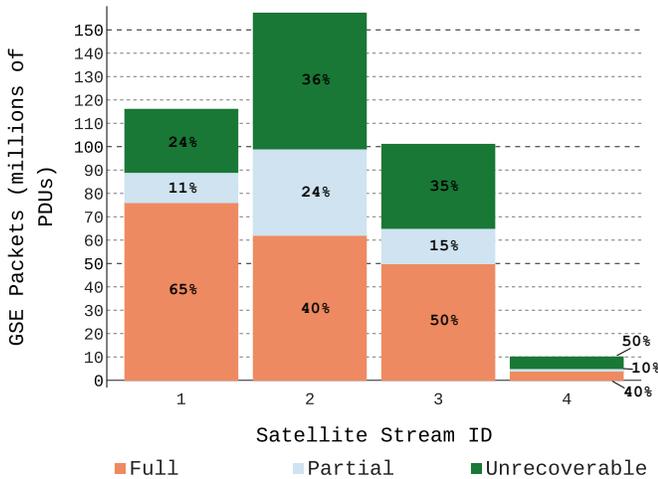


Fig. 6: The degree to which GSE packets within a given stream were recoverable. Stream 4 was of significantly lower throughput than the others and was included to assess GSEExtract’s performance in lower bandwidth contexts.

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

C. 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 IV-A, 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 PDUs contained within the targeted streams and partially recover a further 10-25% of corrupted PDUs (Figure 6).

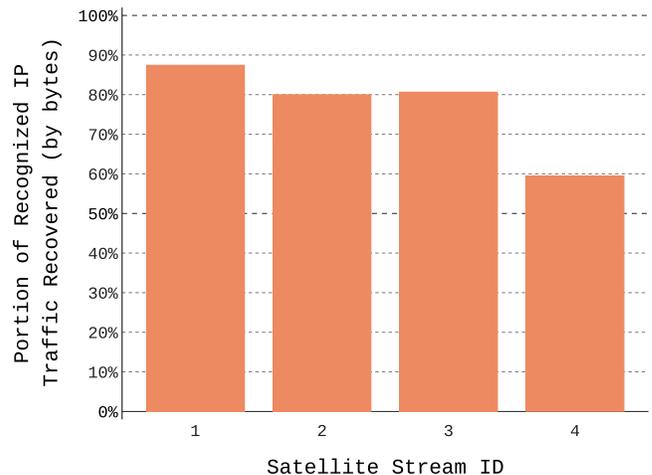


Fig. 7: The overall proportion of successfully reconstructed IP payload bytes using GSEExtract. These metrics were only calculable for successfully identified IP packet headers and do not apply to “unrecoverable” GSE packets lost in the signal processing stage (see Figure 6).

We lacked ground-truth regarding the quantity of internet traffic transmitted which makes it difficult to determine what

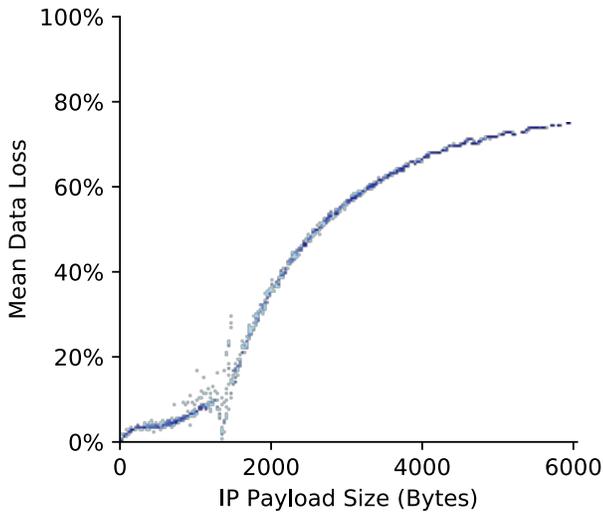


Fig. 8: The average percentage of a given IP packet which is unrecoverable by GSEExtract. As IP packets grow larger, especially above around 1.4kb, GSEExtract’s success rate diminishes due to dropped fragments in the signal processing stage.

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 GSEExtract into a recovered capture. In the case where a large number of IP packets were corrupted, it is expected that GSEExtract 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, GSEExtract will not add many additional bytes. This metric suggests that, at the IP packet level, GSEExtract recovered on average, approximately 92% of any given IP payload. However, in terms of overall data volume we estimate that GSEExtract was able to reconstruct between 60% to 85% of bytes transmitted on a given frequency (Figure 7). Performance was roughly correlated with signal quality, with the lowest-quality data signal also showing significantly higher rates of data corruption using GSEExtract. Additional variance in performance measurements may result from specific network properties and 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 GSEExtract were smaller packets which could be transmitted entirely within in a single 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

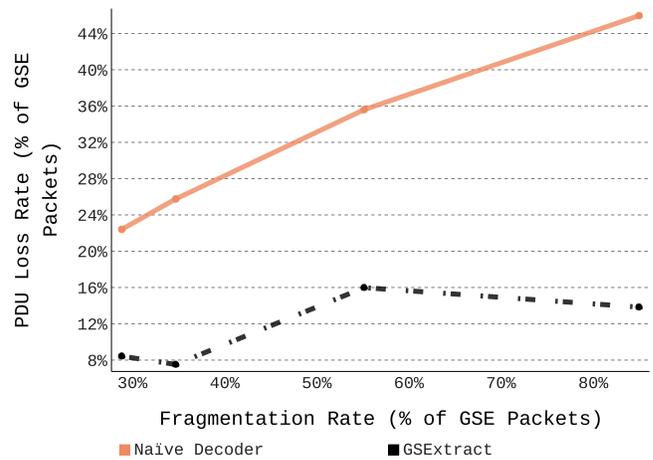


Fig. 9: A comparison of packet loss with and without GSEExtract’s PDU recovery. The solid line depicts a naive decoder which employs GSEExtract’s basic re-synchronization strategy but no other forensic techniques. The dashed line depicts GSEExtract’s performance, indicating only those packets for which partial recovery was impossible.

those fragments is not picked up by the signal hardware. The strength of this relationship can be observed in Figure 8.

Even in the case of fragmented packets, however, GSEExtract 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 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 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 6).

D. Additional Experimental 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), GSEExtract is capable of automatically reconstructing and resuming analysis of broken GSE data streams. While beyond the scope of this security analysis, GSEExtract 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 [34].

E. 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. At the time of submission, these conversations are ongoing but many contacted organizations have expressed surprise at the findings and an interest in taking steps to mitigate them. In some cases, 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 from this publication to give them time to address identified issues.

V. THREAT MODEL

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

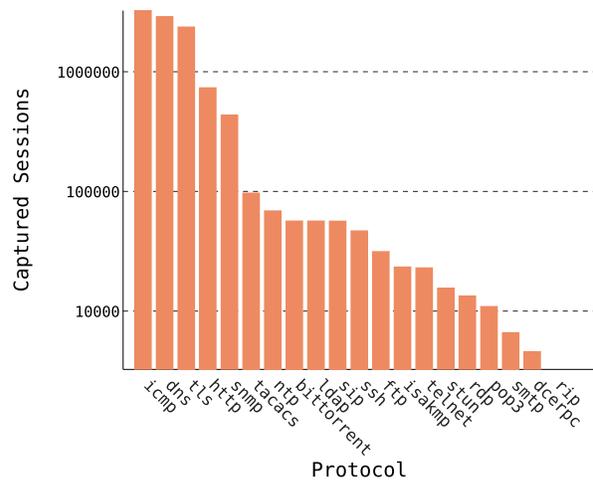


Fig. 10: The 20 most commonly observed protocols across all collected VSAT signals. Note that sessions are counted on a log scale.

observer. However, as mentioned in Section II, significant concerns have been raised regarding the threats posed by criminals, pirates, and terrorists to critical maritime systems. Throughout the paper, we note findings that appear intuitively relevant with regards to these specific threats.

VI. GENERAL FINDINGS

All four maritime VSAT networks included in our study 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 [46]. Moreover, one of the satellites included in our study was launched within the past 3 years, suggesting that these findings are not merely representative of security issues in legacy systems.

A. Applications and Protocols

The principal protocols identified in our recordings are outlined in Figure 10. 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.

TABLE II: 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

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 11). 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 II. More detailed case studies of specific services can be found in Sections VII and VIII.

B. 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

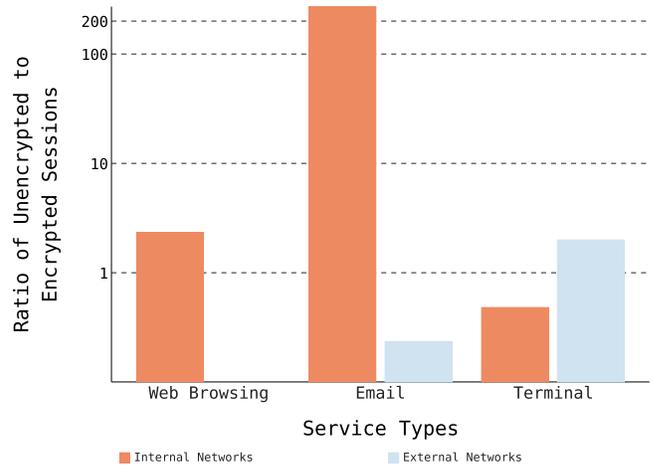


Fig. 11: A comparison of the ratio of sessions using unencrypted protocols vs encrypted alternatives on the basis of whether a session is contained with the local IP range or reaches out to globally addressable IPs. A higher preference for unencrypted protocols is observed in “internal” VSAT traffic. Note that this ratio is expressed on a log scale.

most significant maritime operators [29]. These included three members of the Fortune Global 500 and at least six publicly traded entities with combined annual revenues exceeding \$700 billion [19]. In the cargo sector alone, we observed sensitive traffic from organizations which, combined, account for more than one-third of all global maritime shipping.

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

TABLE III: 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.

- 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 III 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.

VII. FINDINGS: PHYSICAL SAFETY AND OPERATIONS

Section II 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 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.

A. 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.

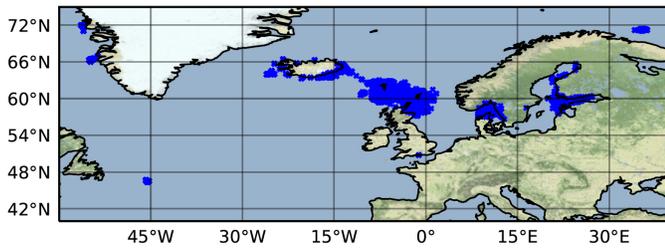


Fig. 12: A map of AIS positions reported in one VSAT stream with a heavy concentration of reported vessels near the Faroe Islands. A total of more than 4 million AIS messages were identified in the study.

As mentioned in Section II, 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 12. 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 IX) 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 [27]. 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 principle 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.
```

Fig. 13: Traffic from an FTP-based ECDIS update. This system is likely trivially vulnerable to the attacks in Section IX.

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 —

```
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.
```

Fig. 14: A captured NM which was transmitted via a clear-text HTTP API. This system is likely trivially vulnerable to the attacks detailed in section IX.

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 13). Were 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 [25]. 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) [24]. 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 14 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.

B. 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 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 15).

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

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 [37]. 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.

IX. ACTIVE ATTACKS

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.

A. 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 19). 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.

B. 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.

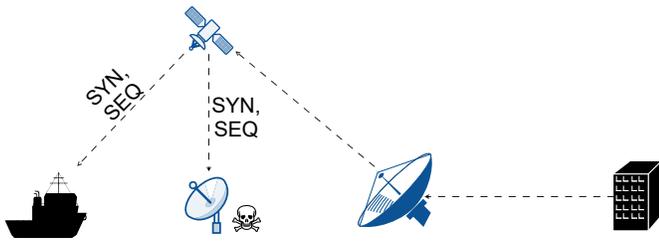
Unique to satellite ecosystems, there is also the risk that the observed TCP session over the air is not the same session as that observed by the receiving vessel and internet endpoint. This is due to the use of Performance Enhancing Proxies (PEPs). PEPs modify TCP connections and generate artificial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

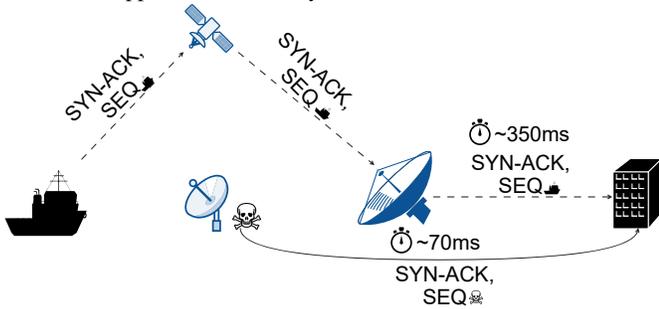
In our study, approximately 425 hosts, or around 5% of observed hosts, had publicly routable IP addresses. However, this is likely not a representative ratio as the provisioning of public IP addresses varies substantially between VSAT providers. Among 11 other VSAT service providers which were considered 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.

C. Hijacking Implementation

To hijack TCP sessions, GSEextract monitors live VSAT traffic for TCP SYN connections from a specified internet host



(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.

Fig. 19: Notional Overview of TCP-Hijacking in VSAT

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.

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.

D. 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 [32]. 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.

X. POSSIBLE SOLUTIONS AND FUTURE WORK

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 [36], [18]. Latency in TCP traffic is treated as an indication of network congestion and thus TCP conversations in satellite environments take much longer to maximize use of available bandwidth. As

a result, ISPs use PEPs to alter TCP headers and generate fake TCP ACK packets on the fly. VPNs prevent the deep-packet inspection necessary to perform these tasks. As such, further research into a link-layer security protocol suitable to the particularities of VSAT environments is likely warranted.

While some proprietary solutions exist, these implementations are not well studied and their security properties are unverified beyond marketing claims [43]. 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 [13]. Industry proposals for securing scientific space missions show promise but lack the key-management infrastructure and multiplexing capabilities for multi-user environments [7]. 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.

XI. CONCLUSION

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 paper 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.

REFERENCES

- [1] A. Adelsbach and U. Greveler, "Satellite Communication without Privacy - Attacker's Paradise," in *Sicherheit*, 2005.
- [2] M. Balduzzi, A. Pasta, and K. Wilhoit, "A Security Evaluation of AIS Automated Identification System," in *Proceedings of the 30th Annual Computer Security Applications Conference*, ser. ACSAC '14. New Orleans, Louisiana, USA: ACM, 2014, pp. 436–445.
- [3] J. A. Bhatti and T. E. Humphreys, "Covert Control of Surface Vessels via Counterfeit Civil GPS Signals," in *Unpublished*, 2015.
- [4] M. Binder, J. Quigley, and H. Tinsley, "Islamic State Chemical Weapons: A Case Contained by its Context?" Combating Terrorism Center at West Point, Tech. Rep., Mar. 2018.
- [5] S. Brizzolara and R. A. Brizzolara, "Autonomous Sea Surface Vehicles," in *Springer Handbook of Ocean Engineering*, ser. Springer Handbooks, M. R. Dhanak and N. I. Xiros, Eds. Cham: Springer International Publishing, 2016, pp. 323–340.
- [6] C4ADS, "Above Us Only Stars: Exposing GPS spoofing in Russia and Syria," Tech. Rep., 2019.
- [7] CCSDS, "Space Data Link Security Protocol - Summary of Concept and Rationale," Green Book, Jun. 2018.
- [8] Cjcr-Software, "EBSpro," <http://ebspro.net/>, 2016.
- [9] M. Cocchiario, "Vessel Accumulation and Cargo Value Estimation," Sep. 2016.
- [10] M. Cox, "CMA CGM BENJAMIN FRANKLIN Gets Hollywood Welcome," <http://maritimematters.com/2015/12/cma-cgm-benjamin-franklin-gets-hollywood-welcome/>, Dec. 2015.
- [11] crazycat69, "CrazyScan: Satellite/terrestrial/cable scan software," <https://sourceforge.net/projects/crazyscan/>, 2018.
- [12] 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)*, Jul. 2015, pp. 1–5.
- [13] L. Duquero, S. Josset, O. Alphand, P. Berthou, and T. Gayraud, "SatiPsec : An Optimized Solution for Securing Multicast and Unicast Satellite Transmissions," in *22nd AIAA International Communications Satellite Systems Conference & Exhibit 2004 (ICSSC)*. American Institute of Aeronautics and Astronautics, 2004.
- [14] Etherr: Wikimedia Commons, "Maritime VSAT system." Jan. 2011.
- [15] 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.
- [16] —, "ETSI TS 102 606 V1.1.1 Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE); Part 1: Protocol," DVB BlueBook, Dec. 2013.
- [17] —, "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.
- [18] A. J. H. Fidler, G. Hernandez, M. Lalovic, T. Pell, and I. G. Rose, "Satellite — A New Opportunity for Broadband Applications," *BT Technology Journal*, vol. 20, no. 1, pp. 29–37, Jan. 2002.
- [19] Fortune, "Global 500," <https://fortune.com/global500/2019/>, 2019.

- [20] R. Hopcraft and K. M. Martin, "Effective maritime cybersecurity regulation – the case for a cyber code," *Journal of the Indian Ocean Region*, vol. 14, no. 3, pp. 354–366, Sep. 2018.
- [21] iDirect, "The Maritime VSAT Advantage: A cost analysis of VSAT broadband versus L-band pay-per-use service," Marketing Materials.
- [22] —, "iDirect Evolution," Tech. Rep., 2018.
- [23] Intellian, "V240mt," <https://www.intelliantech.com/Satcom/v-series/v240mt>.
- [24] International Hydrographic Organization, "IHO Transfer Standard for Digital Hydrographic Data," Tech. Rep., Nov. 2000.
- [25] —, "IHO Data Protection Scheme," Tech. Rep., Jan. 2015.
- [26] O. Jacq, X. Boudvin, D. Brosset, Y. Kermarrec, and J. Simonin, "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)*, Oct. 2018, pp. 1–8.
- [27] K. D. Jones, K. Tam, and M. Papadaki, "Threats and impacts in maritime cyber security," 2016.
- [28] D. Kirkpatrick, "Tankers Are Attacked in Mideast, and U.S. Says Video Shows Iran Was Involved," <https://www.nytimes.com/2019/06/13/world/middleeast/oil-tanker-attack-gulf-oman.html>, Jun. 2019.
- [29] 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*, Oct. 2018, pp. 1–8.
- [30] Marine Traffic, "MarineTraffic: Global Ship Tracking Intelligence — AIS Marine Traffic," <https://www.marinetraffic.com/en/ais/home/centerx:2.5/centery:33.6/zoom:6>.
- [31] MarineMec, "Fishing vessel owners turn to VSAT," https://www.marinemec.com/news/view,fishing-vessel-owners-turn-to-vsats_40942.htm, Nov. 2015.
- [32] Microsoft, "Microsoft Security Bulletin MS08-020," <https://docs.microsoft.com/en-us/security-updates/securitybulletins/2008/ms08-020>, Apr. 2008.
- [33] myshiptracking.com, "My Ship Tracking," <http://www.myshiptracking.com>.
- [34] Newtec, "Satcom for National Security & Intelligence Gathering," https://www.newtec.eu/frontend/files/application_note/intelligence-gathering.pdf, 2015.
- [35] J. Pavur, D. Moser, V. Lenders, and I. Martinovic, "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*, ser. WiSec '19. Miami, Florida: ACM, 2019, pp. 277–284.
- [36] A. Roy-Chowdhury, J. S. Baras, M. Hadjitheodosiou, and S. Papademetriou, "Security issues in hybrid networks with a satellite component," *IEEE Wireless Communications*, vol. 12, no. 6, pp. 50–61, Dec. 2005.
- [37] H. C. Rudolph and N. Grundmann, "CIPHERSUITE," <https://ciphersuite.info/>, 2019.
- [38] R. Santamarta, "SATCOM Terminals: Hacking by Air, Sea, and Land," DEFCON White Paper, 2014.
- [39] —, "Last Call for SATCOM Security," Blackhat Whitepaper 2018, Tech. Rep., Aug. 2018.
- [40] D. Schmidt, K. Radke, S. Camtepe, E. Foo, and M. Ren, "A Survey and Analysis of the GNSS Spoofing Threat and Countermeasures," *ACM Comput. Surv.*, vol. 48, no. 4, pp. 64:1–64:31, May 2016.
- [41] M. Smith, D. Moser, M. Strohmeier, V. Lenders, and I. Martinovic, "Undermining Privacy in the Aircraft Communications Addressing and Reporting System (ACARS)," *Proceedings on Privacy Enhancing Technologies*, vol. 2018, no. 3, pp. 105–122, Jun. 2018.
- [42] A. Tabrizi, "What we know about Gulf of Oman tanker attacks," *BBC News*, Jun. 2019.
- [43] Tellitec, "TL200 TelliShape V2.6"
- [44] Tero Marine, "Products - Tero Marine," <https://www.teromarine.com/products/>.
- [45] United Nations Conference on Trade and Development, *Review of Maritime Transport 2018*. New York and Geneva: UNITED NATIONS, 2018, oCLC: 1083029692.
- [46] Via Satellite and iDirect, "The Coming Wave of Maritime VSAT Growth," <https://www.satellitetoday.com/long-form-stories/maritime-vsats/>, 2015.
- [47] M. Wingrove, "Cruise ship orders for VSAT providers," https://www.passengership.info/news/view,cruise-ship-orders-for-vsats-providers_54065.htm, Sep. 2018.

APPENDIX A GSEXTRACT IMPLEMENTATION

This appendix details the general approach used by the GSEextract 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. While these techniques may have some academic value, given the potential for abuse we have elected not to release GSEextract until the issues identified in this study are addressed by a significant proportion of the maritime VSAT industry. This appendix is thus intended to provide technical insight into the techniques employed which may be of academic interest without releasing a fully featured attack tool to the general public. A simplified overview of the entire GSEextract data extraction process can be found in Figure 22 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 20 [17]. 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. 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 GSEextract'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

Fig. 20: The structure of a DVB-S2 BBFrame Header

To extract BBFrames, GSEextract 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-

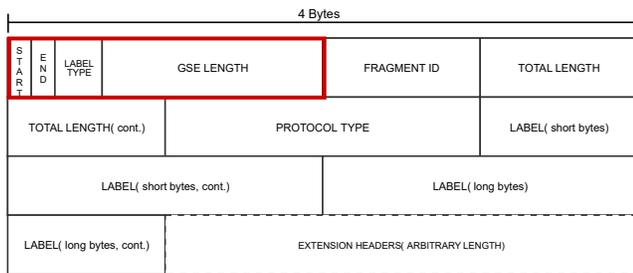


Fig. 21: An overview of the GSE header format. Only the first two bytes are required but the header can be of arbitrary length depending on the addition of optional extensions.

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 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 GSEextract recovers its synchronization with the DVB-S2 stream.

Next, we extract bytes up the 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 21 [16]. Each GSE packet has a variable-length header of at least 2-bytes which includes a 12bit 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.

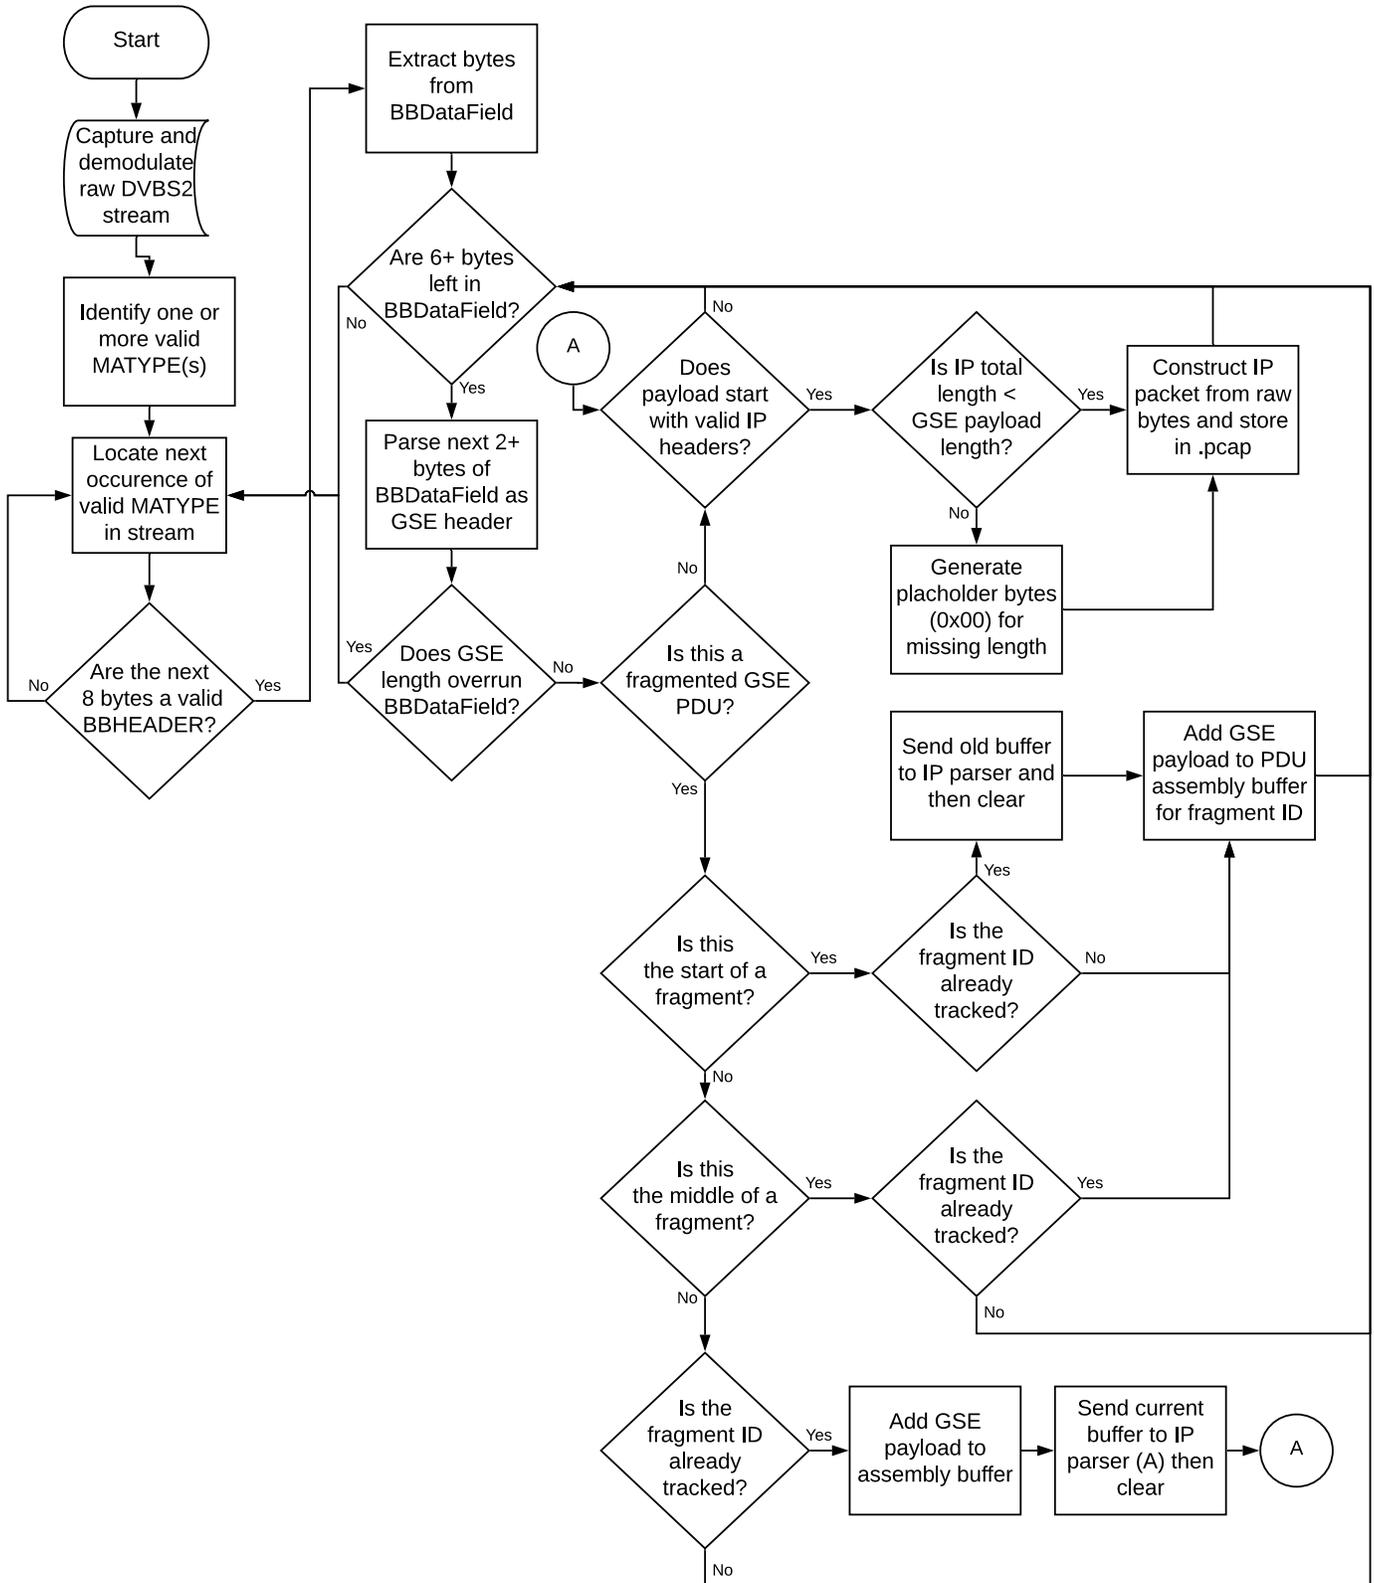


Fig. 22: A notional overview of the stream interpretation and recovery approach used by GSEextract