Jammer on the Horizon: A Robust Method for GPS Jammer Localization Using ADS-B Data

Matthias Schäfer - *SeRo Systems GmbH, Germany*
Steve Sõeruer, Taavi Kippak - *Estonian Air Navigation Services (EANS)*
Erkki Sadrak - *Estonian Consumer Protection and Technical Regulatory Authority (TTJA)*

BIOGRAPHY

Matthias Schäfer Dr. Schaefer has over 10 years of experience researching and developing novel technologies within the wireless and air traffic surveillance domains. He co-founded the OpenSky Network, a global-scale air traffic surveillance network that collects and analyzes data from more than 90,000 flights daily for non-profit research. In 2014, Matthias founded SeRo Systems, which offers sensor and big data solutions for spectrum monitoring and security of air traffic surveillance systems. He received his PhD in security of mobile systems from the University of Kaiserslautern, Germany, in 2018.

Steve Sõeruer Steve Sõeruer has been a Lead Engineer with Estonian Air Navigation Services since 2012. He is responsible for keeping the surveillance service up-to-date and ensuring systems are fully functional. Additionally, he is a co-founder of StaySharp OÜ (atseponline.com), a company dedicated to developing intuitive and engaging web-based training for CNS engineers worldwide.

Taavi Kippak Taavi Kippak has been a Surveillance Engineer with Estonian Air Navigation Services (EANS) since 2017. He has over 15 years of experience in software development and training air traffic safety electronics personnel (ATSEP). He is a co-founder of atseponline.com, an innovative provider of online ATSEP training.

Erkki Sadrak Erkki Sadrak earned his MSc in radio engineering from Tallinn Technical University in 1993. He has over 20 years of expertise in mobile networks and telecommunications, working at the Estonian Mobile Telephone Company (later Telia) and the Estonian Air Navigation Services (EANS). He is now with the Estonian Consumer Protection and Technical Regulatory Authority (TTJA) where he is responsible for radio spectrum monitoring. He has held an amateur radio license since 1987.

ABSTRACT

GPS jamming has emerged as a significant threat, with ramifications extending across various sectors, especially in the realm of air traffic surveillance. This paper delves into the synergistic application of ADS-B data as a tool for the detection of large-scale GPS jamming attacks. By tapping into this resource, we demonstrate a system that continuously and promptly detects these threats over expansive regions. Analyzing data from our European network segment over a one-year period, we offer an in-depth characterization of the detected interference. In addition, we propose a novel method for locating the source of the interference under circumstances where existing methods fail. More specifically, existing methods are inadequate when the jammed area is not fully covered or air traffic is sparse. Addressing these gaps, we introduce a novel localization method that uses the ADS-B-based observations to reconstruct the GPS jammer's radio horizon and then uses this information to determine its location.

I. INTRODUCTION

Since the Global Positioning System (GPS) became fully available to civil users in 2000, the technology has been widely adopted across many domains, from consumer electronics to critical infrastructure such as transportation systems and power grids. Its reliability and accuracy have made it a vital component of modern infrastructure, with millions of users worldwide relying on its continuous availability for positioning, navigation, and timing (PNT) services. Moreover, with an increasing number of such Global Navigation Satellite Systems (GNSSs) being deployed, the adoption of GNSSs is expected to continue growing in the coming years. The European Union predicts that the number of GNSS-enabled devices will grow from about 7 billion to 10.6 billion by 2031 (EUSPA, 2022).

With modern society's growing dependence on GPS[1], it has increasingly become a target of nefarious actors who aim to disrupt both civil and military operations. Enabled by the lack of security and resilience in GPS, criminals, terrorists and state actors use

[1]Throughout this paper, we will use the terms GPS and GNSS interchangeably. Although most systems, including ADS-B, primarily rely on GPS at present, it is important to note that GPS and other GNSSs often operate on the same or nearby frequencies. Consequently, when interference occurs, it usually impacts GPS and other GNSSs simultaneously. Thus, our results are also applicable to other GNSS systems.

https://doi.org/10.33012/2023.19393

intentional radio-frequency interference (RFI) to avoid being tracked, to cause economic damage, or to disrupt an opponent's operations. As a consequence, GNSS RFI, also known as GNSS jamming, has become a growing concern all over the world. In the United States, the president signed an executive order calling for a more responsible use of PNT services to strengthen national resilience (U.S. Government, 2020), while the European Union Aviation Safety Agency (EASA) has issued a Safety Information Bulletin (SIB) that warns authorities and airspace users about intensified GNSS jamming activities in the context of the Russian invasion of Ukraine (EASA, 2022).

A crucial step in addressing the growing GPS jamming issue is the deployment and operation of monitoring infrastructure. In fact, EASA has recommended that aviation authorities and air navigation service providers "establish a process to collect information on GNSS degradations" (EASA, 2022). Such a process requires the collection of data that provide situational awareness. Moreover, besides detecting interference, the location of the source must be identified to be able to take appropriate mitigation measures.

Typically, monitoring of GNSS interference is conducted by authorities through the use of ground-based installations or vehicles equipped with direction-finding devices. However, these methods have limitations. Ground-based sensors, for instance, are limited in their detection range to tens of kilometers due to the curvature of the Earth. Additionally, in the case of short-lived jamming events, authorities may not be able to react quickly enough to move the necessary equipment to the correct location to locate the source of the interference. As a result, there is a need for approaches that provide continuous monitoring and data collection over large areas, which traditional methods are unable to accomplish.

One promising solution to these challenges is to leverage the navigational integrity provided by aircraft through the Automatic Dependent Surveillance–Broadcast (ADS-B), an air traffic surveillance technology used by more than 94% of commercial air traffic in Europe and the US (EUROCONTROL, 2023). In ADS-B, aircraft continuously report their GNSS position along with a navigation integrity code (NIC) indicating the reliability of the transmitted information. When the aircraft is exposed to jamming, the NIC decreases, which indicates a GNSS degradation or even a GNSS denial. Therefore, by means of the NIC value, ground-based ADS-B receivers can infer information about the GNSS signal quality at the aircraft remotely and, by doing so, use ADS-B-equipped aircraft as flying GPS interference sensors.

The idea of using ADS-B data to detect GNSS RFI and locate their source is not new. Researchers have proposed several methods in recent years for both GNSS RFI detection and jammer localization based on ADS-B data. Table 2 provides an overview of the methods proposed in the literature. However, these methods have built-in assumptions that limit their applicability in practice. For instance, approaches that rely on signal strength models are susceptible to noise introduced by unknown factors such as signal attenuation by the aircraft's fuselage, the unknown gain pattern of the GNSS antenna on top of the aircraft, and other factors that affect signal strength or the NIC value itself. Similarly, methods that rely on centroid calculations have problems in scenarios with uneven traffic distributions and sparse data in regions with low traffic volumes.

In response to the limitations of existing methods, we propose a novel approach that combines new ADS-B data processing techniques with radio horizon calculations to locate the source of interference. Unlike previous methods that make assumptions about signal strength models or perform calculations on the shape of the affected region, our approach first identifies the locations where the aircraft intersect the radio horizon of the jammer. By calculating the intersections of these horizons, we can determine the location of the jammer. We evaluate the effectiveness of our approach by applying it to datasets from multiple interference events across Europe.

In contrast to existing methods, our method works well even in scenarios with sparse data, when the shape of the affected area is not circular or unknown, and when the interference event is short-lived. Since we do not make any assumptions about the jammer's unknown transmit power or the implementation-specific relationship between the NIC values and the jamming signal strength, our method is widely applicable and robust.

The remainder of this paper is organized as follows: In section II, we delve into the technological foundations, providing insight into the relevant aspects of ADS-B and GPS, laying the groundwork for the subsequent sections. Then, section III shifts the focus to GPS interference detection using ADS-B, discussing the methodologies employed and sharing results from our extensive 365-day analysis. Following this, section IV offers an overview of existing methods for GPS jammer localization, highlighting their limitations, and then introduces our novel approach, supplemented with example results that demonstrate its efficacy.

II. TECHNOLOGICAL FOUNDATIONS

1. Automatic Dependent Surveillance-Broadcast

For many decades, radar held its ground as the dominant technology in air traffic surveillance, with the Global Positioning System (GPS) finding little to no application in this sector. However, over the last two decades, a dedicated push of authorities around the world towards modernizing the air traffic management system began to reshape the landscape. Central to this modernization initiative is the GPS-based surveillance technology known as Automatic Dependent Surveillance-Broadcast

Table 1: Comparison of GPS Interference Threat Models

Model	Typical Power	Cost of Equipment	Range	Impact on Aviation
Unintentional Interference	Uncertain	Unknown	Varies	Varies
Script kiddies/hobbyists	Up to 100 mW	<$1,000	10s of meters	Little to none
Criminals	Up to 100 Watts	<$10,000	100-1000s of meters	Low altitude traffic / airports
Nation state or military actors	1000s of Watts	>$10,000	100s of kilometers	Significant fractions of airspace, including en route traffic

(ADS-B). Essentially, aircraft equipped with ADS-B utilize GPS for their positioning and then relay their positional data and other status information to any nearby receiving system.

To ensure that users of these positional data understand its reliability, each ADS-B position report is tagged with indicators of integrity and accuracy. Among these indicators is the Navigational Integrity Category (NIC). The NIC is a value that ranges between 0 and 11, where a value of 0 signifies uncertain position integrity, and 11 indicates the reported position comes with the highest levels of integrity and accuracy (RTCA Inc., 2011).

Under standard operating conditions, aircraft regularly report their positional data at an average rate of 2 updates every second. The accompanying NIC values typically hover between 7 and 10, depending on a variety of factors such as the aircraft's avionics and GPS equipment. Moreover, transient and regional factors such as the current geometry of the GPS satellites and the increased noise due to atmospheric effects can affect the NIC values as well.

Given the considerable transmission power used by ADS-B (up to 500 Watts) the range of ADS-B signal reception is primarily constrained by the radio horizon. For ground-based receivers positioned in relatively flat terrains, this radio horizon extends to roughly 500 km for aircraft flying in the en-route airspace. The consequence and one of the primary benefits of ADS-B is that expansive coverage of higher altitudes can be achieved with a fairly low receiver density.

2. GPS Jamming Threat Models

Systems that rely on radio signals, such as GPS, are inherently susceptible to interference. However, in order to design effective detection and mitigation strategies, it is essential to understand the multifaceted nature of such interference. Different interference sources have varying impacts and thus necessitate distinct countermeasures. In this study, we have grouped interference into four main categories, but we will focus our attention on just one of them.

1. *Unintentional Interference:* This covers everything from counter drone systems, over devices with electromagnetic compatibility issues to 5G infrastructure. This type of interference is largely uncharacterized, with uncertain transmit power and waveform, resulting in erratic and unpredictable impacts. This kind of interference is not targeted and there is no malicious intent.
2. *Script kiddies/hobbyists:* These are individuals who tinker with comparably cheap radio equipment which they either purchased online or built themselves. These devices can typically generate interference of up to 100 mW of power. Their motives are usually curiosity without any harmful intent.
3. *Criminals:* People with ciminal intent often use slightly more advanced equipment to either shield themselves from tracking or manipulate systems that rely on GPS. They, too, typically purchase their equipment online and are able to produce interference of up to 100 Watts.
4. *Nation state or military actors:* They typically employ high-tech and expensive equipment to generate very powerful interference (1000s of Watts) and their motives range from demonstrating power to disrupting opponents to protecting their assets from threats like GPS-guided missiles.

A comparison of the four threat models is shown in Table 1.

Once detected or reported, the typical response to interference is "interference hunting". This method uses highly specialized equipment that is deployed on cars or aircraft to track down and characterize the interference. The primary advantage of this method is the high data resolution it provides by directly accessing the interference using high quality state-of-the-art equipment. Moreover, this method provides a ground truth of the situation and typically leaves no questions unanswered. However, the big disadvantage of this method is that it is reactive, regionally limited and often lags behind the complex and expansive dynamics

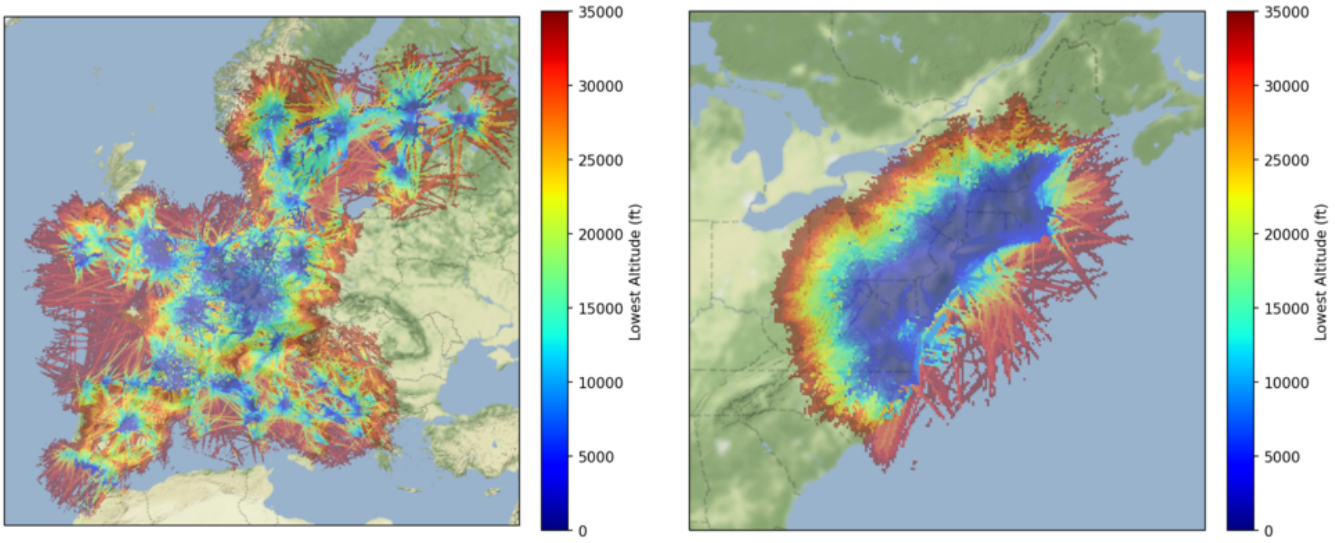

(a) The EU network segment, operated by SeRo Systems, covers most of Europe, though some regions in Eastern Europe and the far North are absent. The covered area totals about 5.8 million km^2.

(b) The network segment in the United States, managed by Nexteon Technologies, predominantly covers the Eastern Seaboard. This segment spans approximately 1 million km^2.

Figure 1: Coverage of our SecureTrack network in Europe and the U.S.

of interference. Since the approach is essentially ad-hoc and interference often temporary, teams might not be able to reach the site in time. Moreover, transporting personnel and specialized equipment is costly and time-consuming which limits the scale of the area that can be surveyed. Finally, while interference does not respect borders, interference hunting is usually conducted by government agencies and as such almost always bound by jurisdiction. Especially large-scale interference presents a significant challenge since the interference source often lies beyond national borders.

Our study particularly focuses on nation-state and military actors as their impact on ADS-B is the most significant, resulting in expansive GPS-denied zones covering several tens of thousands of square kilometers. In contrast, the other three threat models often have little to no effect on air traffic – and thus ADS-B – as their range is limited to a few tens or hundreds of meters. It is worth noting, however, that incidents of smaller scales affecting ADS-B do occur. For instance, in 2022, significant disruption occurred around Dallas airport due to GPS interference, compelling the FAA to redirect traffic. This event, widely covered in the media and documented by open-source tools, undeniably impacted air traffic and ADS-B (Goodin, 2022). However, to provide a tailored and effective solution, we will focus on large-scale GPS interference as caused by the fourth threat model for the remainder of this paper.

3. SecureTrack Data

The data used for this study was collected by the SecureTrack system, developed and operated by SeRo Systems[2] and Nexteon Technologies[3]. SecureTrack comprises a network of ground-based GPS and 1090 MHz signal receivers that continuously stream the received radar and ADS-B data, as well as data on the health of the considered frequencies, to a central data processing system. This central system conducts a range of real-time data analyses and anomaly detection calculations on the data stream, offering real-time alerting, situational awareness, and online analysis capabilities to its users.

As of this writing, the network comprises two segments: a European segment and a U.S. segment. The coverage of each segment is illustrated in Figure 1. In this work, our primary focus is on the European segment, as large-scale interference is more prevalent in this region. Interference observed in the U.S. tends to align more with the characteristics of the other three threat models, which are outside of the scope of this study.

Contrasting with crowdsourced systems such as Flightradar24[4], FlightAware[5], or the OpenSky Network[6], SecureTrack is a homogeneous network relying solely on vertically integrated high-performance receivers for data collection. This approach

[2]https://sero-systems.de
[3]https://www.nexteon.aero/
[4]https://www.flightradar24.com
[5]https://flightaware.com
[6]https://opensky-network.org

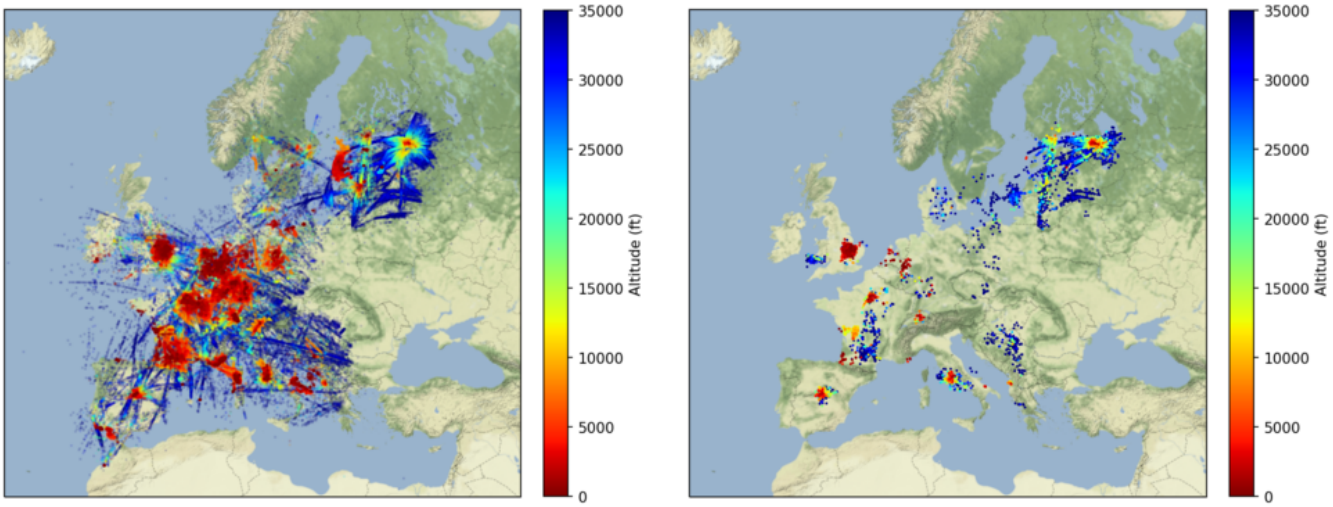

(a) Unfiltered NIC anomalies.

(b) Filtered NIC anomalies.

Figure 2: NIC anomalies observed in SecureTrack's EU segment over the 365 days period from August 22, 2022 and August 22, 2023.

ensures that coverage is more reliable and data quality and integrity are very high due to the absence of the data collection and receiver noise. Nonetheless, for the sake of reproducibility and broader applicability of our methods, we do not utilize any proprietary SecureTrack data items in this study. It is worth noting, however, that achieving similar results with data from other networks may necessitate more rigorous data cleansing and filtering, as crowdsourcing ADS-B data often comes with lower resolution, higher noise levels and data integrity challenges (Schäfer et al., 2018).

III. DETECTING GPS INTERFERENCE VIA ADS-B

To detect GPS interferences in ADS-B data, we rely on the NIC values reported by the avionics. The idea of using the NIC value for interference detection is not new and has been employed by other researchers for the same purpose (Liu et al., 2020b, 2021, 2022a,b, 2023; Dacus et al., 2022). The novelty of our approach lies in how we process and filter the NIC values to detect GPS interference. It is also worth noting that ADS-B offers other features that can be used for interference detection, such as gaps in the trajectory as used by (Jonáš and Vitan, 2019), and the navigational accuracy category as used by Darabseh et al. (2019); Lukeš et al. (2020); Figuet et al. (2022); Fol and Felux (2022). However, we found the NIC value to be a superior feature, as it can be tracked easily and provides high sensitivity with a comparably low level of noise.

The overall concept of our detection approach is as follows: SecureTrack monitors the NIC values of all ADS-B-equipped aircraft within its range. When the NIC value drops below a threshold of 6, SecureTrack flags the event as an anomaly and creates an entry in its NIC anomalies database. This entry includes the timestamp of the event and other aircraft state information, such as its position. Once the NIC value of the aircraft recovers, the database entry is updated by adding the same information, marking the end of the anomaly. These anomalies are then post-processed to eliminate noise in the data and to identify potential interference events. Further details on each of these steps are provided in the following subsections.

Note that we chose 6 as a threshold for NIC anomalies as a trade-off between a false alarm rate and the sensitivity of the approach. As NIC values typically hover between 7 and 10 under normal conditions, considering values of 5 or lower as anomalous turned out to provide the right balance.

1. NIC Anomalies

Throughout the span of one year, from August 22, 2022, to August 22, 2023, SecureTrack registered a total of 311,925 NIC anomalies. The spatial distribution of these irregularities is shown in Figure 2a. However, given the vast number and broad geographic spread of these anomalies throughout the entire coverage area, it becomes evident that factors other than interference might be causing these NIC anomalies. If interference was solely responsible for all these anomalies, then GPS functionality in Europe would be virtually compromised – a scenario we know not to be true. Delving deeper into the spatial density of these anomalies, we found that a significant proportion of them were isolated incidents. By additionally considering the temporal distribution of these anomalies, our notion was reinforced by the fact that most of them were not only isolated in space but also in time. Our hypothesis now is that if interference were prevalent in a region, it would affect a multitude of aircraft rather than a

single one. Yet, our density analysis largely rejects this hypothesis as the majority of the anomalies appear to be brief disruptions experienced by individual aircraft, with no surrounding aircraft confirming the disturbance. This observation underscores the need for refined filtering in order to leverage these NIC anomalies for interference detection.

2. Filtering

Further analysis of the data revealed several causes for anomalies unrelated to interference.

A significant proportion of the anomalies, about 80%, were brief isolated drop-outs lasting less than 60 seconds. Moreover, these could not be associated with any concurrent nearby anomalies. We classify these as transient avionics disturbances. This abundance of glitches emphasizes a principal drawback of ADS-B: its reliance on numerous avionic components. Contrary to radar, an ADS-B-based air traffic surveillance system bears a heightened risk of malfunction due to the increased number of points of failures in the system. The substantial number of anomalies we observed further highlights this issue of ADS-B as an air surveillance system.

Beyond isolated drop-outs, several aircraft consistently exhibited NIC values fluctuating above and below 6. This pattern created a trail of anomalies throughout our monitoring area. Although the precise cause remains unidentified and may vary among the aircraft, we are confident that external interference is not a contributing factor. This conclusion is drawn from observing this behavior of these aircraft regardless of their location. Potential causes of such issues might include onboard installation issues like damaged cables, antenna defects, or incompatibilities between avionic components.

A less prevalent, yet still misleading source of noise was tied to consecutive NIC anomalies from specific aircraft within limited airspace segments. Such anomalies are often rooted in aerobatics, military training, or, more generally, aircraft undertaking severe aerial maneuvers. The extreme pitch and roll angles of these movements misalign the GPS antennas, typically positioned on the aircraft's top, away from the satellites. Further blocked by the aircraft's fuselage, these antennas lose their direct satellite line of sight. This results in a degradation of the NIC value which in turn results in a sequence of NIC anomalies for that aircraft.

To harness NIC anomalies for GPS interference detection, it's crucial to filter out the aforementioned noise. We therefore applied the following filtering steps:

1. **Persistent Avionics Issues:** To identify anomalies from continuous avionics malfunctions, we gauged the total duration during which each aircraft registered low NIC values. By normalizing this against the aircraft's overall flight time and setting a threshold of 0.25 (indicating 25% problematic reporting time), aircraft surpassing this limit were labeled as problematic, and their associated anomalies were discarded.
2. **Extreme Flight Maneuvers:** Anomalies arising from aircraft demonstrating intensive flight behaviors were next. By analyzing the altitude profiles and identifying aircraft that exhibited multiple sharp ascents and descents of several hundred feet, we eliminated this noise source as well.
3. **Agglomerative Clustering:** Lastly, an agglomerative clustering method was applied to the residual NIC anomalies. This algorithm spotlighted NIC anomalies across aircraft with close temporal and spatial proximity. By setting a dual threshold of 50 km spatially and 30 minutes temporally, isolated anomalies were further eliminated.

3. Results

The remaining anomalies, classified as potential GPS interference instances, are depicted in Figure 2b. Our filtering removed approximately 97% of the primary anomalies, yielding 9330 anomalies that were potentially associated with interference. Our clustering algorithm further separated these anomalies into 1901 distinct jamming events. It is worth noting that the count of these events is sensitive to the clustering algorithm's thresholds. In regions with dense air traffic, smaller thresholds enhance precision but may also separate anomalies belonging to the same event into multiple events. Conversely, in sparser regions, wider thresholds prevent overlooking events due to infrequent air traffic.

We further analyzed these potential interference events deeper to gauge their validity. By further inspecting the NIC anomalies and air traffic within the detected interference regions we noticed a trend: multiple aircraft typically displayed a NIC value drop concurrently, recovering either upon exiting the affected area or collectively once the interference ceased. This pattern strongly indicates temporary interference in these regions, with the degree of certainty increasing as more aircraft were affected. In certain cases, supplemental information was also available to provide further context. For instance, in August 2022, in the United Kingdom, a NOTAM was issued on two days coinciding with detected large-scale interference. Although the NOTAM did not explicitly mention GPS interference, it did allude to military exercises in the affected area. Furthermore, pilot reports confirming GPS loss to air traffic control confirmed several events we identified in Northeast Europe. While there remains a degree of uncertainty regarding a few smaller events impacting a limited number of aircraft, the overarching characteristics of the majority of the detected events suggest interference.

In summary, potential interference was identified on 263 out of the 365 days. Figure 3a shows the temporal distribution of these

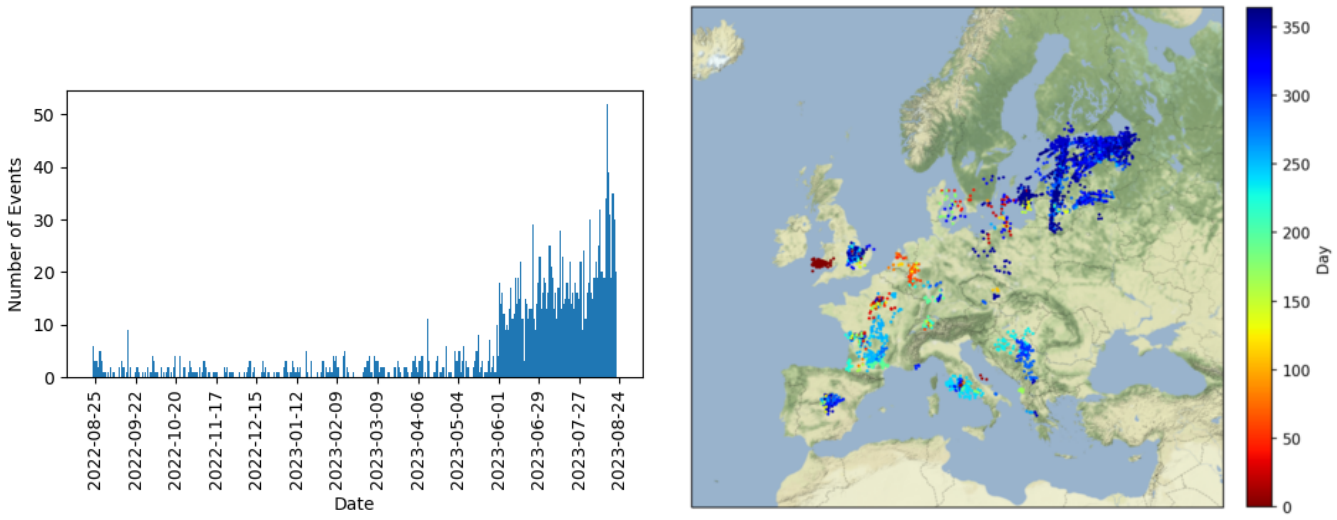

(a) Daily potential interference event count. A significant increase after June 1, 2023 is predominantly observed in Northeast Europe.

(b) Geo-spatial distribution of potential interference events, color-coded by the number of days after the start of our dataset (August 22, 2022).

Figure 3: Timeline of the detected interference events.

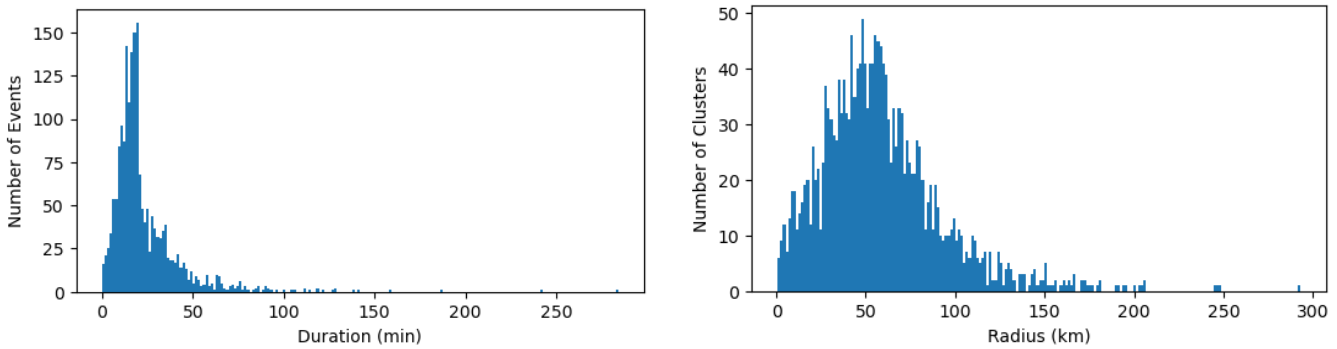

(a) Distribution of the duration of detected events. Most events were under 60 minutes, with some extending up to 5 hours.

(b) Distribution of event radii, defined as the distance from the center of all anomalies associated with a specific event to its farthest anomaly.

Figure 4: Characterization of the detected interference events.

events. The sharp rise in the number of events since June 1, 2023 can be almost fully attributed to the Northeast of Europe. This region has experienced a surge in jamming activities recently, culminating in near-constant GPS service disruption in certain regions, as further visualized in Figure 3b.

The main conclusion of this analysis is that the problem of large-scale interference is more prevalent in Europe than previously assumed. The reason for this is that, while reporting procedures do exist, manual reporting is often neglected and therefore unreliable. In addition, the tools that are typically available for detection, such as open-source tools[7] are not designed for detailed event detection and characterization. While these tools can highlight issues and raise (media) awareness in regions facing extreme long-term interference, they lack the temporal resolution required to detect the majority of events evident in our data.

The underlying reason becomes clearer upon examining Figure 4. It shows both the duration (Figure 4a) and geographic scope (Figure 4b) of the events. It is important to mention that both distributions might vary based on the parameters used during clustering. Nevertheless, regardless of clustering parameter choices, most events are short-lived (under 60 minutes for the 95th percentile) and span regions with a radius of roughly 120 km (95th percentile). The brevity of these events primarily accounts for their absence in publicly available tools. GPSJam, for example, processes data in 24-hour batches. With such a coarse

[7]e.g., GPSJam, accessible via `https://gpsjam.org`

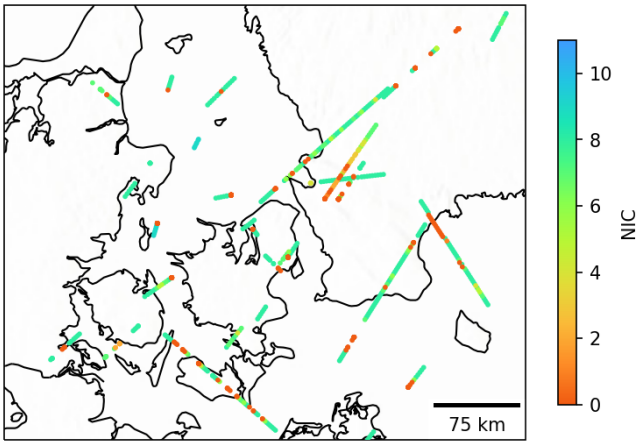

(a) Air traffic in the affected Kattegat region and the observed NIC values.

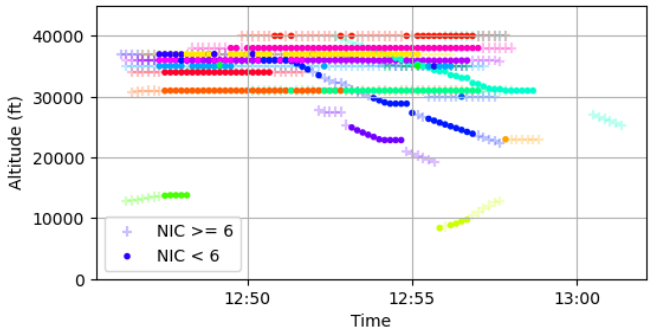

(b) Reported altitudes of aircraft present during the event. Although aircraft ceased reporting their position, they continuously provided their barometric altitude throughout the interference event.

Figure 5: An interference event around Kattegat, a sea area between Denmark and Sweden connecting the Baltic Sea to the North Sea. The event occurred on October 3, 2022, from 12:47 to 12:57 UTC. Both figures show data from 1 minute before the event began to 1 minute after it ended.

time resolution, short events go undetected since the affected traffic percentage is minuscule compared to the overall air traffic observed in a single day.

a) Case Study: Kattegat

To provide a clearer perspective on the detected events, we present a more detailed analysis of a representative event from the data set. This event occurred on October 3, 2022, near Kattegat, a sea area situated between Denmark and Sweden. It began at 12:47 UTC and ended at 12:57 UTC, lasting 10 minutes. As can be seen in Figure 5a, the interference was not strong enough to fully deny the GPS service across its horizon, which is typically at around 400-500 km. However, it still affected all air traffic up to an altitude of 40,000 ft (compare Figure 5b) over an area with an approximate radius of 200 km.

Other notable observations can be drawn from Figure 5. Figure 5a displays all the aircraft trajectories in the area with the color of each trajectory point representing the reported NIC value. Almost all trajectories exhibit a reduction in NIC value, with most plummeting to 0 (red). Notably, many trajectories show gaps, a result of aircraft typically stopping to provide location reports when they lose their GPS fix. In contrast, Figure 5b provides continued insight into the event, as the aircraft still transmits barometric altitude. This figure uses different colors for each aircraft and different markers for NIC values above or below our threshold of 6. This display nicely demonstrates that all aircraft experienced a simultaneous NIC value drop and only recovered upon descending outside the jammer's interference range or after the 10-minute interference duration.

This event is representative for many others identified using our detection method. Another recurrent pattern we observed comprises intermittent interference episodes, each spanning 15-30 minutes, separated by 30-60 minute breaks. Extended uninterrupted jamming periods, on the other hand, are rare and predominantly emerge in politically tense regions.

IV. LOCALIZING GPS JAMMERS

Once interference has been detected and characterized, locating its source becomes the logical next step towards impact mitigation and a quick resolution of the issue. Pinpointing the origin of the interference not only aids authorities in taking timely corrective measures but also provides valuable insights into the nature and intent of the disruptions. Although there are several methods to achieve this, their applicability and effectiveness can vary, especially in the context of large-scale interference events. In this section, we will provide an overview of existing techniques, introduce our novel approach tailored for the specific challenges posed by large-scale events, and offer two comprehensive case studies to demonstrate the practical applications and benefits of our method in scenarios where other methods fail.

1. Existing Detection and Localization Methods

The challenge of detecting and localizing GNSS interference using ADS-B data has been addressed by several researchers in the recent years. Jonáš and Vitan (2019) first presented a technique that utilized ADS-B tracking gaps for detection to calculate the distance ratios between the aircraft and the interference source. However, the authors acknowledge that many of the assumptions

that were made cannot be made in practice. Nevertheless, they applied their method to data from a real-world event and were able to find a narrow peak.

In the same year, Darabseh et al. (2019) examined data collected by the OpenSky Network[8], focusing on anomalies in the reported navigational accuracy category (NAC) values. Despite compiling known incidents and analyzing data from aircraft near these incidents, they were unable to confirm any incidents in the data.

A year later, Liu et al. (2020b) used OpenSky ADS-B data for interference event characterization and source localization. Their method used the center of the affected region as an estimate for the jammer location. The authors also analyzed the behavior of the NIC value during an interference exercise at the Edwards Air Force Base located in California in another study published in the same year (Liu et al., 2020a) .

Also in 2020, Lukeš et al. (2020) analyzed the NAC values of ADS-B reports that were collected at a location in Prague. In addition, they conducted a controlled experiment with a jamming device on an ADS-B transponder to determine the mapping between interference power and the NAC value reported by the transponder. They used the experimental data as training data for their pattern-matching algorithm, which they then applied to the larger ADS-B data set.

In the next year, Liu et al. (2021) used machine learning to detect GPS interference in ADS-B data with the intent to capture the various unknown factors that influence the shape of jammed regions. They used a mix of real-world data from the OpenSky Network and simulated data around a known event in the Cypriot airspace. Similarly to Jonáš and Vitan (2019), the same authors switched in a follow-up work (Liu et al., 2022a) to signal strength modeling and simulations and applied the resulting model to the same data they used in (Liu et al., 2021). However, they faced challenges due to the uneven distribution of the data.

In Liu et al. (2022b), the authors then proposed two methods to detect interference and localize the source. They applied the methods to the ADS-B data around an incident at Denver International Airport. They used the Cypriot data from Liu et al. (2021) to train their Bayesian detection approach. The second approach is a least-squares minimization approach that tries to find the jammer location and transmit power by fitting the estimated receiver jamming power with power measurements. The power measurements are not direct measurements but were inferred from the NIC values.

Also in 2022, Dacus et al. (2022) picks up the notion of gaps similar to Jonáš and Vitan (2019) and focuses on interpolating these gaps to be able to calculate the centroid location of the interpolated positions as an estimate for the interference source location. They evaluated their approach using OpenSky data from an event around Denver International Airport.

Meanwhile, studies by Figuet et al. (2022) and Felux et al. (2023) analyzed batches of OpenSky data that covered several months in 2022. They employed a simple detection method that assumed interference whenever the NAC value reported by an aircraft dropped to 60 for at least 60 seconds. Using this method, the authors were able to find a lot of GPS interference around the Ukraine conflict.

Fol and Felux (2022) looked at OpenSky data from flights where cockpit crews reported GPS loss and tried to assess the impact of GPS interference on the operations onboard the aircraft. After the case study, the authors applied their method to a larger data set and detected several hot spots around Europe.

Earlier this year, Chen et al. (2023) built a testbed for analyzing the impact of different GNSS jamming strategies on the NIC values provided by aircraft. They tested this by implementing an algorithm to calculate the NIC based on the horizontal protection level that was provided by three different commercial off-the-shelf GPS receivers.

Finally, Liu et al. (2023) analyzed ADS-B data from an event around Dallas-Fort Worth International Airport (KDFW) in October 2022. The authors also extended their work from Liu et al. (2022b) and combined it with some of the findings of Chen et al. (2023) in order to get an estimate for the received jamming signal power at the aircraft. Using these information, they then used a least-squares minimization to find the most likely jammer location.

An overview of these works is provided in Table 2. Existing methods can be broadly separated into two categories: Methods that map NIC values to jammer signal power estimates and then use this estimate to infer information on the jammer's location and methods that analyze the shape of the affected region, using geometric features like the centroid, to estimate the jammer's location.

Methods in the first category aim to translate NIC values into estimated jamming signal levels at the aircraft antenna, subsequently using this estimation to pinpoint the jammer's location. Nonetheless, accurately modeling the strength of a signal that originates from the ground and is received by a directed GPS antenna on top of an airborne aircraft is extremely challenging. Several unknown factors, including the transmitter and receiver antenna gain patterns, jamming signal attenuation by the aircraft's fuselage, among others, play a significant role in influencing the received signal level. Predicting these myriad factors accurately is close to impossible without knowing the exact installations on each aircraft. The situation is further complicated when the

[8]https://opensky-network.org

Table 2: Overview of the literature on GNSS RFI detection and source localization through ADS-B data.

Reference	Parameter	Localization	Localization Method
Jonáš and Vitan (2019)	Gaps	yes	Signal strength modeling
Darabseh et al. (2019)	NACp	no	-
Liu et al. (2020b)	NIC	yes	Centroid of affected region
Lukeš et al. (2020)	NACp	no	-
Liu et al. (2021)	NIC	yes	Machine learning
Liu et al. (2022a)	NIC	yes	Signal strength modeling
Liu et al. (2022b)	NIC	yes	Bayesian + Signal strength modeling
Dacus et al. (2022)	NIC	yes	Centroid of affected region
Figuet et al. (2022)	NACp	no	-
Fol and Felux (2022)	NACp + SIL	no	-
Liu et al. (2023)	NIC	yes	Signal strength modeling

parameter used for inferring signal level only has 12 distinct levels and frequently simply drops to zero during interference events. Such complexities suggest that the precision and practicality of these methods in real-world scenarios remain notably limited.

The second category of methods focuses on analyzing the shape of the affected region. By computing features such as the centroid, these methods try to estimate the jammer's location. This approach proves effective in scenarios where the interference area is circular and many aircraft enter the interference area from various directions, thus providing a clear picture of the situation. However, its efficacy diminishes significantly in situations characterized by low traffic volume, uneven distribution of the traffic, and incomplete coverage of the affected area. Unfortunately, these limitations render these methods unusable for many of the events that we detected as we often face one or more of these limiting conditions.

2. Jammer on the Horizon Principle

To address the limitations of existing methods in situations characterized by sparse and unevenly distributed traffic, we propose a novel approach. Although our approach is not a panacea, it nicely complements existing techniques. More specifically, it is tailored for extensive interference events where the jamming signal is strong enough to interrupt GPS services up to the jammer's radio horizon. The approach works as follows.

We begin by assuming the presence of a powerful jammer that creates a GPS-denied area extending to its full radio horizon. In other words, this area is only restricted by the jammer's line of sight which in turn is limited by the Earth's curvature[9]. Under conditions of relatively flat terrain, such a radio horizon generally spans around 400-500km considering aircraft at cruising altitudes of up to 45,000ft.

Now, whenever an aircraft flies into or out of this interference range and thus exhibits a drop or recovery in its NIC value, it crosses this radio horizon. Hence, each start or end of a NIC anomaly provides us with a data point that traces the contour of the jammer's radio horizon. To infer information about the jammer's whereabouts from this information, we make use of the fact that the radio horizon is reciprocal, meaning that the moment when the aircraft crosses the jammer's radio horizon is also the moment when the jammer first appears on the aircraft's horizon. Analogously, the moment the aircraft disappears behind the jammer's radio horizon is the same moment when the jammer disappears from the aircraft's line of sight. Consequently, with each radio horizon observation at an aircraft's respective moment of interference entry or exit, we learn something about the jammer's location as it lies somewhere on each of these circles. Hence, obtaining multiple of these circles and calculating their intersections provides an estimate of the jammer's position.

This concept is depicted in Figure 6. In the next section, we put our methodology to the test, applying it to two different events observed in northeast Europe. By further correlating our results with infrastructure, we validate the effectiveness of our approach, provided our assumption of a strong jammer is met.

[9]While the Earth isn't a perfect ellipsoid and there might be obstacles like buildings, large-scale jammers are usually placed in locations offering an unobstructed line of sight in all directions to ensure optimal performance. Thus, for the scenarios discussed in this paper, this assumption holds reasonably well.

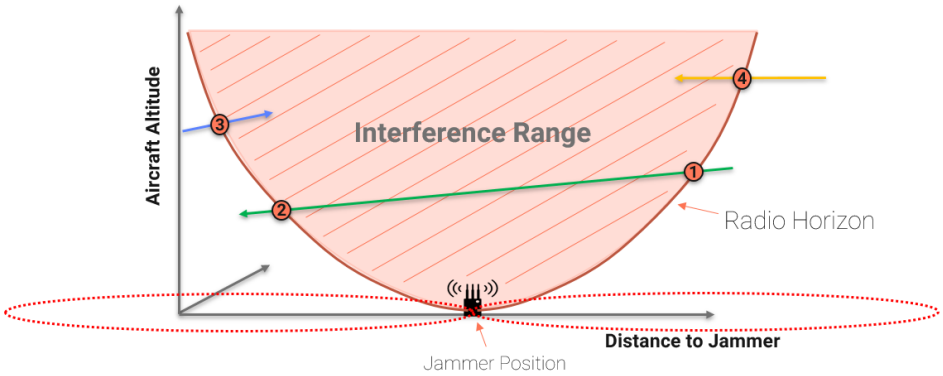

Figure 6: The concept of our jammer localization method. Each time an aircraft enters or exits the interference range, its NIC value drops or recovers, providing us with samples of the jammer's radio horizon. The jammer's position can be found by intersecting the respective radio horizons of the aircraft at their entry or exit of the interference range.

3. Case Studies

a) Baltic States, January 9, 2023

On January 9, 2023, our SecureTrack monitoring system detected an interference affecting all three Baltic States from 7:40am to 8:46am UTC. During this period, little air traffic was observed, with just 17 ADS-B-equipped aircraft in the airspace. Beginning at 7:40 UTC, NIC values reported by several aircraft began to decrease abruptly. Numerous aircraft on a North-South route across the Baltic states showed fluctuating NIC values between 0 and 5. Conversely, two East-bound aircraft ceased their position reports entirely. At the same time, West-bound traffic displayed no NIC value degradation, suggesting the interference originated from the East. The related data from this event is shown in Figure 7a.

In order to pinpoint the jammer's location, we evaluated the dataset in Figure 7a, focusing on the initial NIC value drops of aircraft that later entered the affected zone. We identified three such flights, presented in Figure 7b. For each of these flights, we determined the radio horizon at the moment they first displayed interference signs. These horizons are visualized as dotted circles. As Figure 7b reveals, all three horizons intersect at an eastern location, indicating our estimated jammer position.

To refine this location, we checked the vicinity for potential jammer-supporting infrastructure and located an airbase roughly 40 km from our estimated intersection by relying on open-source databases[10] and satellite imagery from Google Earth. Although lacking concrete evidence, we are confident in this finding, especially since it aligns with results from subsequent similar jamming incidents in later months.

Attentive readers might spot a secondary, less sharp intersection to the West. This is merely an artifact due to the vertical alignment of the selected positions. We can dismiss this western intersection based on the absence of interference detection in the West, contrasted with the full denial of service observed in the East.

b) Baltic States, July 27, 2023

As mentioned in subsection III.3, the Baltic region has witnessed a noticeable surge in jamming activity from June 1, 2023. Since then, a substantial portion of its airspace has been consistently denied GPS service. While at this point, this situation seems more enduring, it is also surprisingly dynamic. Almost on a daily basis, different jammers are activated and deactivated and new jammers. This case study focuses on a jammer that appeared on the evening of July 26, 2023, and remained active until the morning of July 28, impacting large parts of Estonian and Finnish airspace.

Utilizing the analysis approach described previously, we selected aircraft positions that displayed changes in NIC values. We then examined the radio horizons at these moments. The extended activity of this particular jammer allowed SecureTrack to gather an extensive set of data. With this larger dataset, we produced a density map of horizon intersections to pinpoint the jammer's location.

Figure 8 shows the outcome of our analysis. Figure 8a displays the radio horizons as aircraft began reporting declines in NIC values. The map in Figure 8b shows the density of these intersections. A location near Lake Lagoda stands out as a large fraction of the radio horizons intersect in this region. This map suggests that numerous aircraft encountered GPS problem the moment this red zone came into their view. Given the consistent behavior of multiple aircraft over an extended period, coupled with the presence of an airbase at the heart of this red zone (about 20km from the highest density spot), our confidence in this

[10] https://www.gfsis.org.ge/maps/russian-military-forces

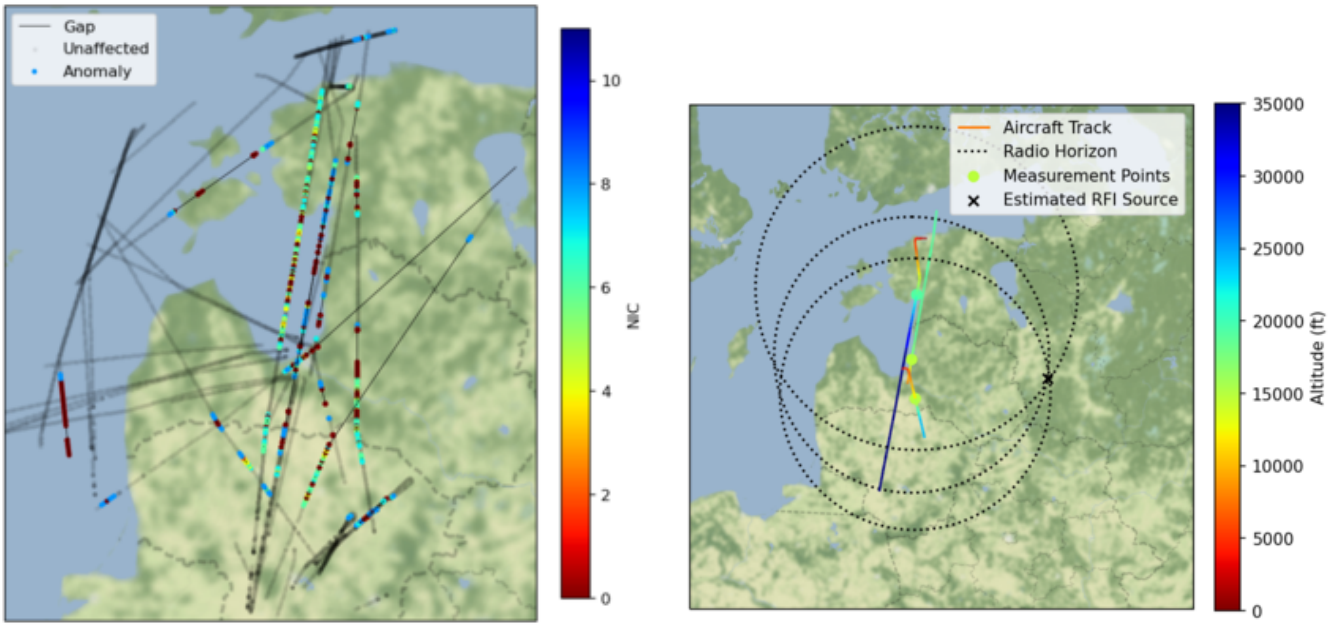

(a) Air traffic in the affected region and the observed NIC values. The two East-bound flights (solid thin black lines) were still tracked by our multilateration function, explaining our knowledge of their routes.

(b) The 3 selected flights and their positions for our jammer localization. Dotted circles show the radio horizons of the aircraft upon initial interference encounter. The intersection to the right is the suspected jammer location.

Figure 7: Interference event detected on January 9, 2023 over the Baltic States and estimation of the location of the interference source.

result is high.

V. CONCLUSION

In this work, we studied utilizing ADS-B data for the detection and characterization of GPS interference. This approach enables continuous, real-time surveillance across extensive regions. The main strengths of our detection methodology lie in its heightened sensitivity and cost-efficiency, offering significant advantages over traditional methods such as interference hunting.

Surprisingly, our investigation revealed that the prevalence of GPS interference is far more extensive than initial expectations. Our monitoring system detected potential interference in almost all countries that are covered by the network. One primary reason behind this unanticipated discovery is the lack of tools designed to offer a comprehensive situational awareness of such activities at a sensitivity that enables both a reliable detection and a characterization of GPS interference events. SecureTrack, the system we presented and used in this study, fills this void, already serving its users as a powerful tool for this very purpose.

We further reviewed the existing methods. The key takeaway is that there is no one-size-fits-all solution. The applicability of a method largely depends on the scenario at hand and we believe with our new method we filled another significant gap. A significant challenge when designing methods for locating the source of interference via ADS-B is the absence of a definitive ground truth. Without such a ground truth, a level of uncertainty about the accuracy and effectiveness of the methods proposed in this and the other works remains. Yet, there are silver linings: the capability to associate findings with other public data sources combined with the insights gained from radio propagation simulations, provides plenty of evidence and grants us a substantial degree of confidence in our results. Public data sources that can provide additional information to pinpoint the jammer or resolve ambiguities include satellite imagery, repositories of military infrastructure, media reports, and social media.

It is crucial to acknowledge that our method is not without limitations. For instance, its performance decreases in scenarios with multiple jammers. In such cases, the intersections of radio horizons may accumulate in multiple regions leading to ambiguities and false results. Moreover, our basic assumption – that a jammer would be sufficiently powerful to induce a denial of GPS service across its entire radio horizon – was not given across all detected events. In these circumstances, other methods need to be considered, such as those highlighted in subsection IV.1, or even to alternative manual methodologies, such as interference hunting or a triangulation of the interference from satellites, aircraft or high altitude balloons.

For those aiming to fine-tune the identification of jammer locations, there's a trove of resources at their disposal. Satellite imagery, repositories of military infrastructure data, public reports, and social media can offer valuable insights. Additionally,

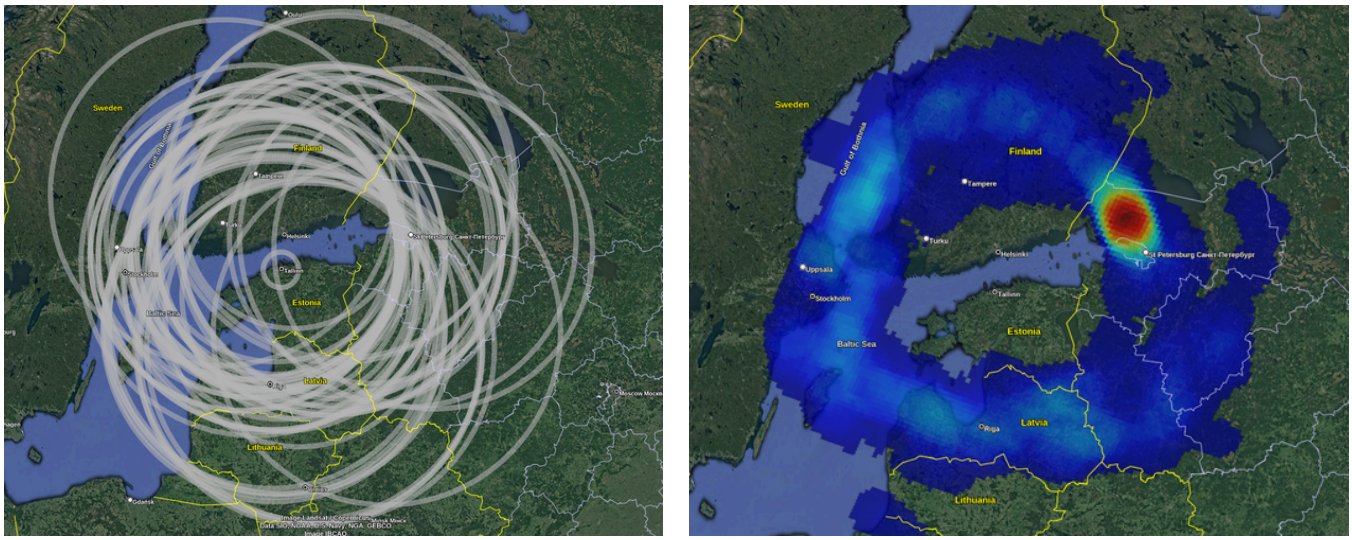

(a) Illustrating the radio horizons during detected interference transitions.

(b) Density map showing intersections of the radio horizons.

Figure 8: Pinpointing a jammer active in the Baltic region from July 26 to July 28, 2023.

technological assets like satellites, aircraft, and high-altitude balloons can be leveraged for precise triangulation. Once interference has been detected and characterized, a myriad of options become available, paving the way for more targeted, effective responses to GPS interference challenges.

REFERENCES

Chen, Y.-H., Liu, Z., Blanch, J., Lo, S., and Walter, T. (2023). A RFI Testbed for Examining GNSS Integrity in the Various Environments. In *Proceedings of the 2023 International Technical Meeting of The Institute of Navigation*.

Dacus, M., Liu, Z., Lo, S., and Walter, T. (2022). Improved RFI Localization Through Aircraft Position Estimation During Losses in ADS-B Reception. In *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*.

Darabseh, A., Bitsikas, E., and Tedongmo, B. (2019). Detecting GPS Jamming Incidents in OpenSky Data. In *Proceedings of the 7th OpenSky Workshop 2019*.

EASA (2022). Global Navigation Satellite System Outage Leading to Navigation / Surveillance Degradation. SIB No.: 2022-02. European Union Aviation Safety Agency.

EUROCONTROL (2023). Automatic Dependent Surveillance-Broadcast Airborne Equipage Monitoring. Accessed online on March 2nd, 2023 via `https://www.eurocontrol.int/service/adsb-equipage`.

EUSPA (2022). EO & GNSS Market Report. European Union Agency for the Space Programme.

Felux, M., Figuet, B., Waltert, M., Fol, P., Strohmeier, M., and Olive, X. (2023). Analysis of GNSS disruptions in European Airspace. In *Proceedings of the 2023 International Technical Meeting of The Institute of Navigation*.

Figuet, B., Waltert, M., Felux, M., and Olive, X. (2022). GNSS Jamming and Its Effect on Air Traffic in Eastern Europe. *Engineering Proceedings*, 28(1).

Fol, P. and Felux, M. (2022). Identification and Operational Impact Analysis of GNSS RFI Based on Flight Crew Reports and ADS-B Data. In *Proceedings of the 7th International Workshop on ATM/CNS (IWAC)*.

Goodin, D. (2022). Gps interference caused the faa to reroute texas air traffic. experts stumped. Ars Technica.

Jonáš, P. and Vitan, V. (2019). Detection and Localization of GNSS Radio Interference using ADS-B Data. In *2019 International Conference on Military Technologies (ICMT)*.

Liu, Z., Blanch, J., Lo, S., and Walter, T. (2023). Investigation of GPS Interference Events with Refinement on the Localization Algorithm. In *Proceedings of the 2023 International Technical Meeting of The Institute of Navigation*.

Liu, Z., Lo, S., and Walter, T. (2020a). Characterization of ADS-B Performance under GNSS Interference. In *Proceedings of the 33rd International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*.

Liu, Z., Lo, S., and Walter, T. (2020b). GNSS Interference Characterization and Localization Using OpenSky ADS-B Data. In *Proceedings of the 8th OpenSky Symposium*.

Liu, Z., Lo, S., and Walter, T. (2021). GNSS Interference Detection Using Machine Learning Algorithms on ADS-B Data. In *Proceedings of the 34th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2021)*.

Liu, Z., Lo, S., and Walter, T. (2022a). GNSS Interference Source Localization Using ADS-B Data. In *Proceedings of the 2022 International Technical Meeting of The Institute of Navigation*.

Liu, Z., Lo, S., Walter, T., and Blanch, J. (2022b). Real-time Detection and Localization of GNSS Interference Source. In *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*.

Lukeš, P., Topková, T., Vlček, T., and Pleninger, S. (2020). Recognition of GNSS Jamming Patterns in ADS-B Data. In *2020 New Trends in Civil Aviation (NTCA)*.

RTCA Inc. (2011). Minimum Operational Performance Standards for 1090 MHz Extended Squitter Automatic Dependent Surveillance – Broadcast (ADS-B) and Traffic Information Services – Broadcast (TIS-B). DO-260B with Corrigendum 1.

Schäfer, M., Strohmeicr, M., Smith, M., Fuchs, M., Lenders, V., and Martinovic, I. (2018). Opensky report 2018: Assessing the integrity of crowdsourced mode s and ads-b data. In *IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*.

U.S. Government (2020). Executive Order No. 13,905. 85 Fed. Reg. 9359.