Journal of Open Aviation Science (2023), Vol.**1**
doi:10.59490/joas.2023.7504

JOAS

PROCEEDINGS | *The 11th OpenSky Symposium*

Securing the Sky: Detecting Aircraft Location Drifting through Cross-Checking Receiver-Based Estimated and Received ADS-B Trajectories

Ala Darabseh,* Syed Khandker, and Christina Pöpper

NYU Abu Dhabi, UAE

*Corresponding author: ala.darabseh@nyu.edu

(Received: 5 April 2024; Revised: 3 June 2024; Accepted: 6 July 2024; Published: 11 July 2024)
(Editor: Tatiana Polishchuk; Reviewers: Vincent Lenders, Matthias Schäfer, Martin Strohmeier)

Abstract

ADS-B serves as a widely adopted protocol in aviation systems, more specifically, among commercial aircraft, transmitting vital information to other aircraft and ground-based receivers. However, the absence of robust security measures and the transmission of unencrypted data render this information susceptible to unauthorized tampering. Exploiting this vulnerability, malicious actors can manipulate the location details within ADS-B transmissions over time, potentially resulting in deviating aircraft from their intended trajectory and even leading to dangerous collisions.

In this paper, we introduce a new approach to identifying instances of aircraft location drifting. Our proposed method capitalizes on the established geographical coordinates of ground-based receivers, randomly positioned across the landscape by a group of dedicated volunteers, who share their data for research purposes. By leveraging these receiver coordinates, we propose a methodology to estimate the anticipated flight path of an aircraft. This projected trajectory is subsequently cross-checked with the trajectory derived from the incoming ADS-B messages, which contain real-time flight data. By conducting this comparison, our method can detect deviations or inconsistencies between the computed expected trajectory and actual flight paths that are derived from the location information in ADS-B messages. This differential analysis enables the prompt identification of potential aircraft trajectory deviating attempts. By integrating receiver-based trajectory estimation and real-time ADS-B data analysis, our approach contributes to enhancing aviation security and safeguarding against potential threats to the integrity of flight information. Our results indicate that our prediction method is valuable, but further improvements in its accuracy are necessary for it to be effectively used as a security validation approach against attacks.

Keywords: ADS-B; Aviation Security; Receiver Placement; Location Spoofing; Cross Checks

Abbreviations: ADS-B: Automatic Dependent Surveillance–Broadcast; ATM: Air Traffic Management

1. Introduction

Automatic Dependent Surveillance-Broadcast (ADS-B) is an air surveillance technology that uses GPS to determine an aircraft's position and then broadcasts this information to air traffic control (ATC) and other aircraft [1]. In contrast to conventional ground-based radar, ADS-B offers enhanced accuracy and reliability in surveillance. Considering the advantages, e.g., safety, efficiency, and cost reduction, this surveillance technology has been widely adopted around the world [2, 3]. Although ADS-B provides many benefits, its security is susceptible to various attacks because of its open de-

sign, where messages are transmitted without encryption. Researchers have unveiled several attack vectors on the ADS-B system, revealing vulnerabilities such as spoofing, flooding, jamming, virtual trajectory modification, and man-in-the-middle attacks [4, 5, 6, 7]. The lack of robust security measures and privacy safeguards, (e.g., encryption, authentication) in the ADS-B protocol creates these vulnerabilities, which then allow attackers to intercept broadcasted messages and manipulate them with the intention of diverting an aircraft's navigation information (e.g., latitude, longitude, altitude, etc.) from its intended path to an alternate trajectory. Such an attack holds the potential to deviate other aircraft from their intended trajectory and could even lead to dangerous collisions between aircraft. To prevent such attacks, various solutions have been proposed, encompassing cryptographic measures and location-verification techniques [8, 9, 10]. However, implementing cryptographic solutions may entail a comprehensive overhaul of the existing ADS-B protocol, posing significant implementation and realization challenges. Conversely, many current location verification methods, such as MLAT [11], necessitate the reception of messages by four or more sensors, a condition that proves challenging to meet. While alternative verification methods [12, 13, 14, 15, 16, 17] have been suggested, they often fall short of providing a comprehensive predictive solution that addresses challenges like the number of receptions, altitude considerations, and other pertinent factors. Consequently, there persists a demand for novel prediction methods that can tackle existing challenges without requiring modifications to the ADS-B protocol.

In this paper, we introduce a novel location-prediction approach based on the intersecting coverage areas of receiving sensors. When a specific set of sensors receives a message, our methodology involves identifying the intersection area of this sensor set, and the predicted location is then situated within that determined area. The challenge lies in the selection of this location and determining the parameters crucial for enhancing prediction accuracy.

The primary objectives of this research are as follows:

1. **Accurate Trajectory Derivation:** We seek to establish trajectory estimation based on the coordinates provided by ADS-B receivers and information from aircraft signals, such as its time of arrival (ToA) and received signal strength (RSS). Furthermore, we aim to quantify the difference between the trajectory estimation and the trajectory as received.
2. **Accuracy of Attack Detection:** We explore how accurately we can detect location-drift attacks using the trajectory estimation method. We discuss and analyze how we can improve the detection process of the attacks. We provide first results, but acknowledge that further improvement is needed by including more data to enhance the prediction result.

2. Related Work

ADS-B is an evolving air-surveillance technology with a broad range of potential applications. It has significantly enhanced air traffic management and safety by providing real-time aircraft tracking data. Nevertheless, the emergence of security issues [6, 5, 7] has raised concerns, leading to various proposed solutions. These solutions primarily fall into two categories: broadcast authentication methods and location verification methods. In the broadcast authentication method, the authenticity and integrity of ADS-B messages are ensured through digital signatures or message authentication codes (MAC) [8, 18]. Additionally, lightweight encryption techniques, such as Format-preserving Feistel-based encryption (FFX), have been suggested as part of these security solutions [19]. However, all these broadcast authentication methods require either adding extra information or changing the current message format, which may introduce complexities and compatibility challenges. On the other hand, location verification methods do not require any modification to the existing ADS-B message format. The latter method focuses on independently verifying the reported location of aircraft using various techniques such as multilateration, time difference of arrival (TDOA), triangulation,

secure track verification, secure motion verification using doppler effect, and cross-referencing with different sensor data.

The authors of [16] argue that current state-of-the-art methods of aircraft localization such as multilateration are insufficient, in particular for modern crowdsourced air traffic networks with random, unplanned deployment geometry. They utilized a combination of the k-nearest neighbor (kNN) algorithm and the expected TDOA of a received signal between multiple sensors to estimate signal's origin. During the training phase, the expected TDOAs for each position for the given sensor deployment are measured. At the verification time for each ADS-B signal, the k-NN of the messages' TDOAs looked up. These neighboring points are then averaged to produce the final estimate of the sender's location.Their experimental result showed that grid-based k-NN approach can increase the effective air traffic surveillance coverage compared to multilateration by a factor of up to 2.5. In [20], the authors proposed a similar scheme called AEALV. This approach leverages a grid to train a kNN regression model by constructing a rectangular grid plane and then dividing it into a large number of squares. Each square is assigned a TDOA vector called a fingerprint, which serves as a unique identifier. When an aircraft claims its location, AEALV first checks if the claimed location is within the airspace. If it is, AEALV then finds the k nearest grid squares to the claimed location. Once the k nearest grid squares have been identified, AEALV calculates the average of the fingerprints of the k nearest grid squares to estimate the aircraft's actual location.

In [13], a TDOA-based lightweight location verification method was proposed. In this method, TDOAs between at least two sensors that received the message are collected and used to verify the claimed position of the signal. In [21], a similar method was used, incorporating geospatial indexing, where the sensor's receiving range is measured first. Then, for a given ADS-B message, it is determined whether the received signal's coordinates are within the range. If so, the ADS-B signal is classified as legitimate; otherwise, it is not. Using timestamps in the ADS-B message for location verification was first proposed in [12]. According to them, for a spoofed signal claiming a false location, propagation delays to different verifiers would not satisfy the expected geometric relationships of the legitimate signal sources, thereby revealing the spoofing attempt. Later, in [22], a time-based location verification method called ADS-BT has been proposed. The authors suggest including an 8-bit timestamp in the ADS-B packet. This enables the determination of the distance between the sender and receiver through two methods: firstly, by the time difference, and secondly, by coordinate distance. In the event of spoofing, these two values will mismatch.

3. Methodology

3.1 Threat Model

We consider active attacks where the attacker manipulates ADS-B messages, impersonates legitimate senders, and broadcasts ADS-B messages with spoofed locations to divert the aircraft from its main trajectory. We assume that the attacker cannot change which receivers should receive a message; they can only spoof the location information as part of a spoofed message.

The attacker may be stationary or mobile. We consider two cases: (1) **Stationary Attacker:** The attacker remains in a fixed location and attempts to spoof messages within their range; in this case, the attacker's goal is to spoof the location of aircraft passing through the area covered by the attacker. (2) **Mobile Attacker:** The attacker is moving, using a ground vehicle (e.g., a car) or drone that moves at a speed different from the aircraft, or is on the same aircraft, moving at the same speed as it. Given an attacker with the described capabilities, we consider two categories of *location-drifting attacks*:

1. **Random Drifter Attack:** The adversary introduces random and unpredictable deviations from the original flight path.

2. **Targeted Precision Drifter Attack:** The attack starts from the actual location of the aircraft and makes the calculated location gradually drift away from the actual aircraft flight path. This type of attack involves precise planning by the attacker to execute subtle alterations that minimize the likelihood of detection.

3.2 System Approach

Our approach for identifying location-drifting attacks relies on the geographical positioning of receivers on the ground. Messages transmitted from aircraft are picked up by a group of sensors positioned within the aircraft's Line of Sight (LoS) at the time of message transmission. Each of these sensors is capable of receiving messages originating from various locations.

The methodology we propose for detecting location-drifting attacks is structured into three key phases:

Phase 1: Deriving Sensor Coverage and Sensor Locations: In the initial step, we determine the coverage area of each individual sensor based on the data received by these sensors.

We process and assess approximately 3 million messages extracted from OpenSky data [23]. Each message contains essential data, including the aircraft's geographic coordinates (latitude and longitude), the list of sensors (receivers) that received the message, the arrival time, and the received signal strength at each receiver.

From this data, we extract the locations in received ADS-B messages per sensor and ultimately create the coverage area for each sensor. The convex hull H of a set of locations L is the smallest convex polygon that contains all the locations in L. Let L_i be the set of locations detected by sensor i. We refer to the convex hull H_i of L_i by

$$H_i = \text{ConvexHull}(L_i).$$

We presume that the sensors exhibit a spherical coverage pattern, with their locations centralized within this area.

Phase 2: Predicting Aircraft Location: Next, we use the sensor coverage information obtained in the first step to predict the aircraft's location.

For each received message, we compile a list of sensors that received that message and determine the intersection area of these sensor coverage boundaries. From this area, we derive the predicted aircraft location. To predict the aircraft location, we explore two methods and evaluate their accuracy:

1. *Central Localization:* This method involves identifying the central point of the intersection area, without considering other factors. Assume the message was received by a set of N sensors and each sensor i is located at coordinates (x_i, y_i). Each sensor has a boundary or coverage area, and the intersection area I of all boundaries is defined as:

$$I = \bigcap_{i=1}^{N} P_i \tag{1}$$

where P_i is a polyshape (polygon) with vertices for each sensor in H_i:

$$P_i = \{(x_{i1}, y_{i1}), (x_{i2}, y_{i2}), \ldots, (x_{im}, y_{im})\}. \tag{2}$$

Let V be the set of vertices defining the boundary of the intersection area I. The central point (x_c, y_c) of the intersection area I is the centroid of the polygon defined by V. If $V = \{(x_1, y_1), (x_2, y_2), \ldots, (x_k, y_k)\}$

the centroid coordinates (x_c, y_c) are calculated as:

$$x_c = \frac{1}{k}\sum_{i=1}^{k} x_i \ , \ y_c = \frac{1}{k}\sum_{i=1}^{k} y_i. \tag{3}$$

2. *Weighted Localization*: In this method, we establish a mesh network G that represents a grid of potential locations:

$$G = \{(x_1, y_1), (x_2, y_2), \ldots, (x_n, y_n)\}. \tag{4}$$

To enhance the accuracy of our predictions, we leverage available data, including ToA and RSS for each message received by every receiver in our sensor network. We take these parameters into account to assign weights to individual points within the grid. Let T_i be the Time of Arrival and S_i be the Received Signal Strength for the message received by sensor i. We order the receivers based on their ToA or RSS values. Let $\mathbf{r} = (x_r, y_r)$ be the coordinates of the first receiver in the ordered list, we call this receiver the *reference sensor*. Then we assign a weight W_j to each potential location $\mathbf{g}_j$ using the weight function $f(d) = \frac{1}{d+0.1}$:

$$W_j = f(d_j) = \frac{1}{d_j + 0.1}, \tag{5}$$

where d_j is the Euclidean distance from each potential location $\mathbf{g}_j = (x_j, y_j)$ in the grid to the reference sensor $\mathbf{r}$:

$$d_j = \sqrt{(x_j - x_r)^2 + (y_j - y_r)^2}. \tag{6}$$

We factor in the distance from a reference sensor. The reference sensor's proximity to the message source enhances the accuracy of our calculations. The location with the highest weight W_j is selected as the final predicted location (x_p, y_p) of the aircraft.

$$(x_p, y_p) = \arg \max_{(x_j, y_j) \in G} W_j \tag{7}$$

This method for location prediction focuses on leveraging multiple data sources and refining the results based on specific criteria, making it a valuable component of our location-drifting attack detection system. We emphasize that while this method holds significant promise, it may require additional testing and refinement to reach its full potential.

Comparing the results of these two methods is crucial to determining the most accurate approach for predicting aircraft locations, which, in turn, is vital for robust location-drifting attack detection. As we will demonstrate in Section 4.4, the Central Localization method does not consistently provide accurate predictions because the same set of sensors receiving the message at different times would yield the same location, irrespective of other criteria (like the arrival time, distance from receivers, and density of the receivers). So the Weighted Localization method is a better choice.

Phase 3: Cross-Verifying Location Data: Finally, we cross-check the calculated aircraft location with the received location data by using two distinct evaluation methods: Trajectory-based Evaluation and Point-based Evaluation,

1. **Trajectory-based Evaluation:**
We begin by observing trajectories for all aircraft in the designated area, relying on received locations from ADS-B receivers. We analyze approximately 3 million messages from ADS-B, we observe 1661 distinct trajectory that are received by 45 sensors. Concurrently, we construct the expected trajectory for each aircraft based on the set of receivers capturing broadcast messages. The Fréchet distance serves as a metric to quantify the similarity between these two trajectories. A Fréchet distance of zero implies identical trajectories, while higher values indicate increasing dissimilarity, tailored to the specific application [24].

2. **Point-based Evaluation:**
   Our second evaluation focuses on assessing the standard deviation of the Haversine distance between actual and predicted locations. Pairwise comparisons of individual locations quantify their spatial separation, and the average deviation across the entire trajectory is determined. Given two trajectories, trajectory_1 (Received from ADS-B message) and trajectory_2 (Computed or predicted), each with coordinates (ϕ_i, λ_i) for the i-th point. We calculate the Haversine distance between each two points (locations), using the following formula:

$$d_i = 2R \cdot \arcsin\left(\sqrt{\sin^2\left(\frac{\phi_{2i} - \phi_{1i}}{2}\right) + \cos(\phi_{1i})\cos(\phi_{2i})\sin^2\left(\frac{\lambda_{2i} - \lambda_{1i}}{2}\right)}\right)$$

   where:
   - d is the Haversine distance.
   - R is the Earth's radius (mean radius = 6371 km).
   - ϕ_1 and ϕ_2 are the latitudes of the two points in radians.
   - λ_1 and λ_2 are the longitudes of the two points in radians.

utilizing this comparison to identify and flag any suspicious or anomalous activities that could indicate a potential attack.

4. Findings: Revealing Discoveries

4.1 Sensor Coverage

Each ADS-B receiver captures messages from aircraft within its range. The range of ADS-B reception can vary widely depending on the specific circumstances and equipment involved. Generally, ADS-B signals are designed for line-of-sight communication, and their effective range can extend up to several hundred nautical miles. However, several factors can influence the actual range, like the altitude of the receiver, the sensitivity of the receiver, and potential obstructions on the way from the aircraft to that receiver.

To assess the actual coverage of these receivers, we gathered one day's worth of ADS-B data for the geographic region between the range of 44 to 56 latitudes decimal degrees and 1 and 20 longitudes decimal degrees. Analyzing 3 million ADS-B messages received by 45 sensors in this area, we organized the data into clusters based on the receptors. The intersection area of all observed sensors in this region was then determined.

In Figure 1, the coverage area of a set of sensors is illustrated. Notably, receivers may sporadically receive data from distant areas, prompting us to normalize the data and eliminate outlier points for greater precision. Figures 1b, 1d, and 1f depict the coverage after this normalization process and the removal of outliers from these sensors.

4.2 Sensor Location

Given our primary focus on sensor location, position accuracy becomes a pivotal factor in achieving precise location estimates. The challenge arises when attempting to acquire the location data for ADS-B receivers. The OpenSky network provided a sample of receiver names and their locations, shared by users themselves, lacking a ground truth for location accuracy.

To address this, we explored estimating receiver locations based on the obtained coverage areas of the convex hull in Section 3.2 and gauged the disparity between estimated and actual locations.

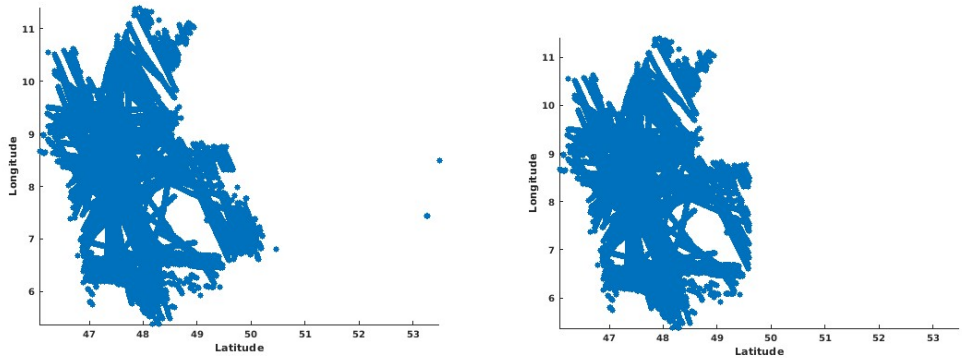

(a) Senor ID:129 before normalization

(b) Senor ID:129 after normalization

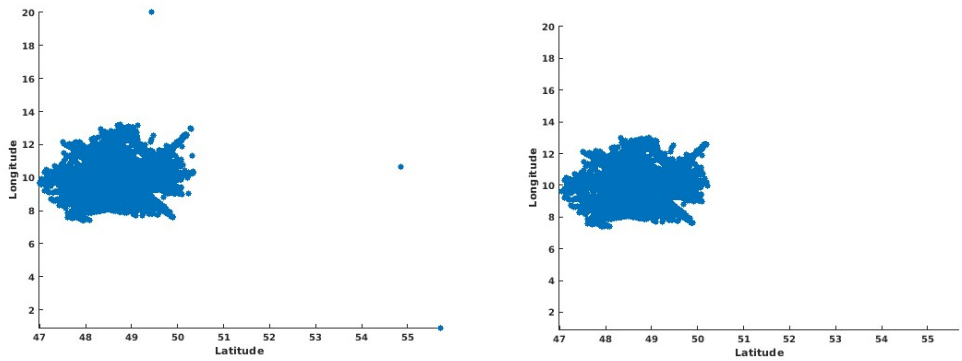

(c) Senor ID:254 before normalization

(d) Senor ID:254 after normalization

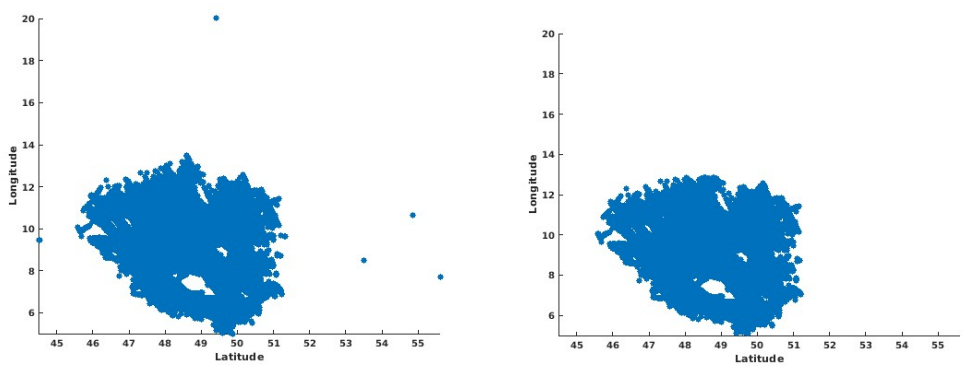

(e) Senor ID:320 before normalization

(f) Senor ID:320 after normalization

Figure 1. Sensors coverage before and after normalization

Assuming the coverage forms a circular area with the sensor's location at its center, our initial observation yielded estimates within a standard deviation of 50 km from the actual location. Furthermore, our approach of getting the center of the convex hull of all messages that have been received by that sensor demonstrates commendable accuracy when considering all the available sensors. Figure 2 illustrates the heatmap of the haversine distance between the actual locations and the computed

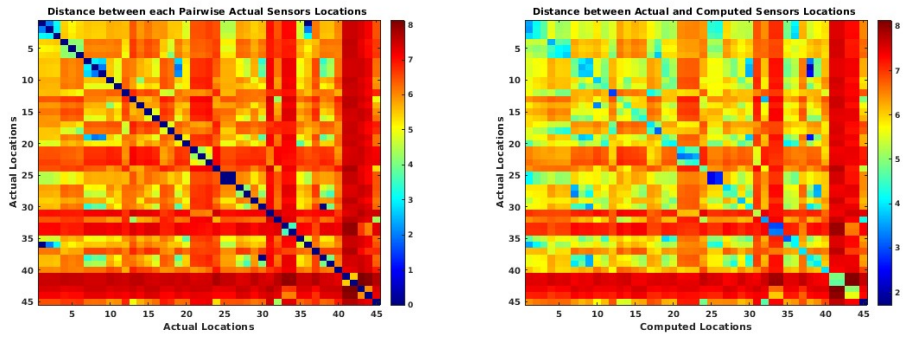

(a) Heatmap of pairwise sensor locations **(b)** Heatmap between actual location and computed location

Figure 2. Logarithmic heatmaps of Haversine distances between (a) pairwise sensor locations and (b) actual and computed sensor locations across all observed sensors. We applied a logarithmic transformation to narrow the range of distances for visualization purposes (distance [km] = $e^x - 1$). E.g., number 1 in the heatmap is $e^1 - 1 \approx 1.7$ km and number 8 corresponds to $e^8 - 1 \approx 2980$ km.

locations using our method. In Figure 2a, the distances between each pair of sensors based on the locations received from OpenSky are depicted. The diagonal holds particular significance in this context, representing pairwise distances; the blue color indicates shorter distances, while red signifies greater distances. The diagonal, being the distance of a sensor location from itself, is zero, and we include it in the plot to showcase the disparity between the ideal scenario and the computed distances. This disparity is visualized in Figure 2b, where the diagonal exhibits a slight difference from the optimal one, depicted with blue shades in distance.

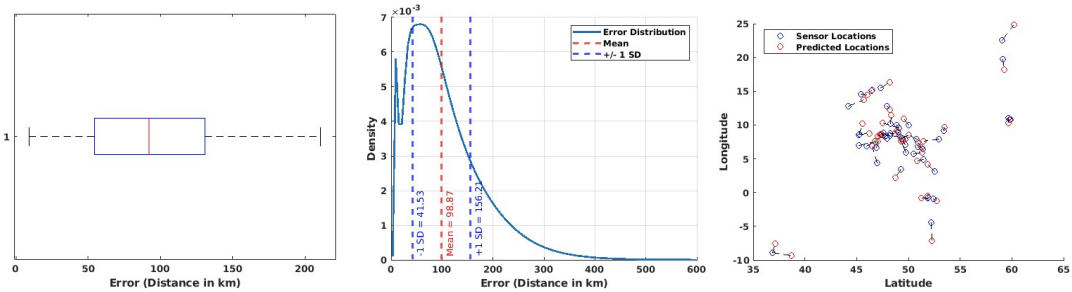

(a) Box Plot **(b)** Kernel Density Distribution (KDE) **(c)** Actual vs. Predicted Locations

Figure 3. Error Analysis between the Actual and Computed Locations for 45 sample receivers

As sensor location is crucial in our prediction, especially in our weighted localization method, we observed the error distribution of our sensor location estimation method (Figure 3). However, the results are not optimized, and other techniques like SkyPos [25] can achieve much better accuracy. We are still willing to explore other methods that can overcome challenges such as time drifting and offsets in the Time of Arrival (ToA) of the received messages by the receivers. Figure 4 visually represents the locations of a set of sensors alongside the estimated positions. Notably, the estimated positions closely approximate the actual ones. This suggests the potential for developing a method to derive locations based on coverage data rather than relying on users to provide location information. It is worth noting that some users regard their receiver locations as private, emphasizing the need for alternative methods. While our intent is not to compromise privacy, this underscores the feasibility of deriving locations in scenarios where such information is unavailable or lacks ground truth for reliable foundation.

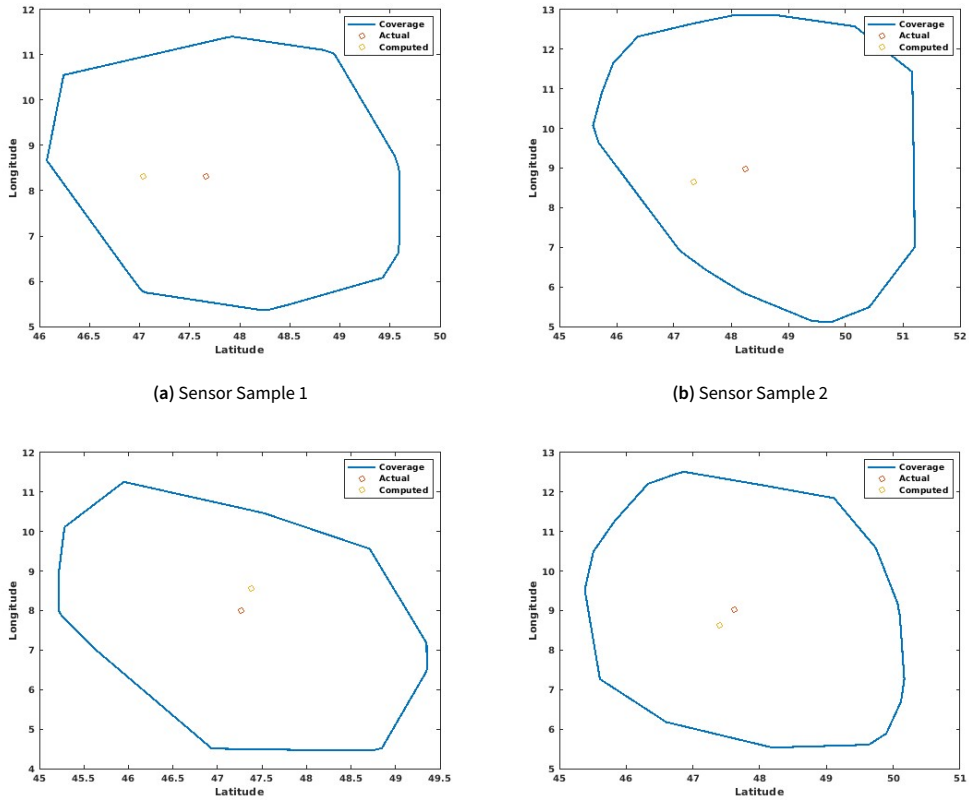

(a) Sensor Sample 1

(b) Sensor Sample 2

(c) Sensor Sample 3

(d) Sensor Sample 4

Figure 4. Reported sensors locations vs. computed locations

4.3 Sensor Pattern

As the aircraft moves from its origin to its destination, it traverses a series of sensors along its route. When a sensor falls within the line of sight of the aircraft, it begins receiving location data, losing contact as the aircraft moves beyond the transmission range. In light of this, we capture the appearance pattern of sensor interactions along the aircraft's path. By analyzing this pattern, we aim to see observable indications that can subsequently enhance the accuracy of our aircraft location predictions.

As illustrated in Figure 5, the sensor patterns are depicted for four distinct trajectories. The figure delineates the aircraft's entry into the coverage area of a particular sensor and its departure from that region over time. Analyzing these patterns provides insights that can be leveraged to enhance prediction accuracy. For instance, monitoring the appearance and disappearance of sensors along the aircraft's trajectory over time can reveal patterns that provide information about the estimated trajectory. Initially, this enhances the accuracy of the expected trajectory and subsequently improves the precision of attack prediction.

4.4 Accuracy Evaluation

We construct the anticipated trajectory after receiving a set of messages and subsequently determine its distances from the received trajectory. This process allows us to measure the accuracy of the

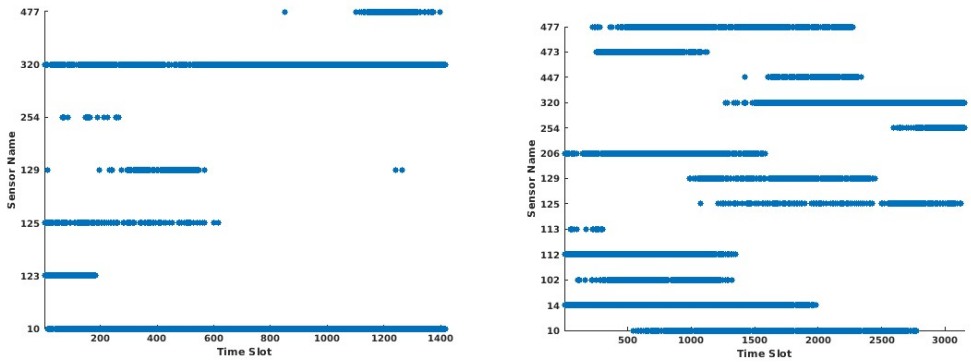

(a) Trajectory Sample 1

(b) Trajectory Sample 2

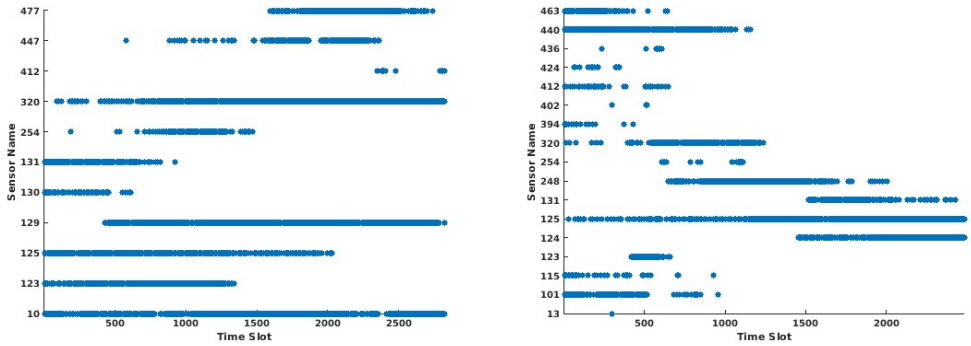

(c) Trajectory Sample 3

(d) Trajectory Sample 4

Figure 5. Sensors patterns of four different trajectories

derived trajectory and establish threshold values and error margins.

Ensuring prediction accuracy is crucial for later attacker detection, particularly in the context of smart drifter attackers. Our approach involves cross-checking received and predicted locations, and evaluating their proximity through two distinct measures:

1. **Trajectory-based Evaluation:**
   In our ADS-B application, we observed a maximum coupling measure of 17 for entirely different trajectories. A coupling measure close to zero suggests trajectory similarity, while a value near 17 indicates dissimilarity. Spotted results reveal a coupling measure of 0.8 for the central localization method, with the weighted localization method yielding improved results by reducing the coupling measure to 0.4.
2. **Point-based Evaluation:**
   The Central Localization method yields a standard deviation of 20 km, while the weighted localization method demonstrates a slight improvement with an standard deviation of 19 km. Although the difference is modest, the weighted localization method holds potential for enhancement, given its reliance on a grid of points rather than a singular central point.

These evaluations collectively contribute to the refinement of our prediction methods, with a focus on enhancing accuracy for effective attacker detection. Figure 6 illustrates the error distribution of our localization approach using point-based evaluation. Figures 6a and 6b depict the variance in

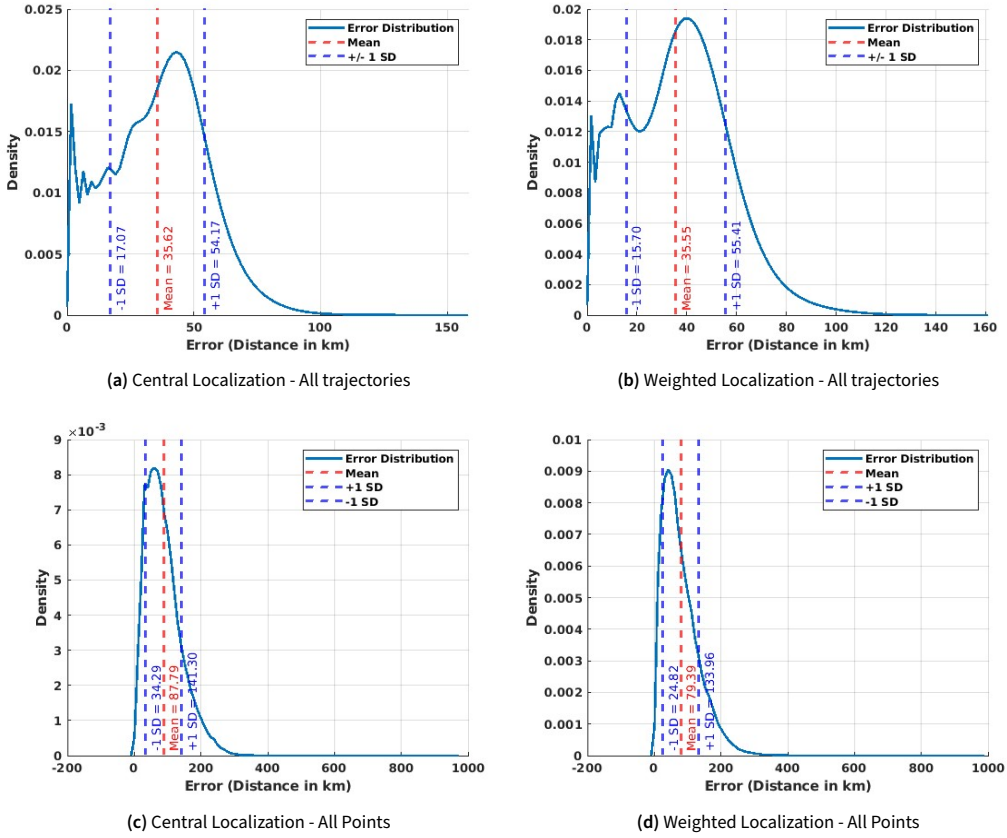

(a) Central Localization - All trajectories

(b) Weighted Localization - All trajectories

(c) Central Localization - All Points

(d) Weighted Localization - All Points

Figure 6. Error distributions of point based evaluation for the two localization methods

Table 1. Standard Deviation and Range of Errors of the Proposed Methods

	Trajectory based	Point based
Central Localization Method	0.8 (Coupling Measure)	20 km
Weighted Localization Method	0.4 (Coupling Measure)	19 km

distance between the reported and estimated locations for entire trajectories using the Central and Weighted Localization methods, respectively. Since this evaluation focuses on individual points, we combined all received locations to derive the overall error distribution based on the two localization methods, which are reflected in Figures 6c and 6d. Table 1 summarizes the comparison between the two proposed methods.

5. Discussion

5.1 Proposed Solution and Attack Detection

The proposed approach involves estimating the predicted location and subsequently cross-referencing it with the received location to assist in whether the location within the ADS-B message has been altered. The predicted value serves as an indicator of the effectiveness of this approach in identifying potential attackers.

The attacker's location is a key factor in this context. The attacker may eavesdrop on communication between the aircraft and ground sensors, modify the location, and then re-transmit the new message. Several considerations arise:

1. The attacker cannot change which sensor receives the message; their manipulation is confined to modifying the location and re-transmitting the message.
2. Detecting a stable attacker becomes more straightforward as the aircraft moves and encounters different sets of sensors along its route. A stable attacker, however, will consistently involve the same set of receivers receiving the message, providing an opportunity for a better detection process by emphasizing sensor diversity across the trajectory. This is because as the aircraft moves, sensor coverage changes, making it difficult for the stable attacker to consistently target the same aircraft.
3. If the attacker is mobile on the ground via a terrestrial vehicle, tracking its speed may not match the aircraft's speed. Detection might take longer, but it will eventually occur as the aircraft enters different geographical zones.
4. An attacker on the same aircraft, moving at the same speed, relies on the attacker's ability to modify packets within the aircraft and her speed to modify the location.

As the current accuracy of our method hovers around 20 km for both prediction methods, We cannot use these results for attack detection at this moment, as more accurate predictions are required. However, we demonstrate that our approach introduces a new method of detection that is worth considering. Ongoing efforts focus on refining accuracy, consolidating results, and exploring early-stage detection capabilities for subtle drifting scenarios.

5.2 Data Acquisition Challenges

Precise Location: The accuracy of our method hinges on the precise locations of sensors, particularly for method two. However, obtaining the ground truth for these locations poses a significant challenge. Some owners of ADS-B receivers consider this information private, which hinders our access to this crucial data. We managed to obtain the location of a limited set of existing receivers from the OpenSky network, though without knowledge of the owners' identities. The accuracy of these shared locations remains uncertain, as it depends on the owners' willingness to disclose accurate information, leaving us without a reliable ground truth for these locations.

5.3 Challenge in Achieving Prediction Accuracy

Our goal is to utilize the coordinates of ADS-B ground receivers to formulate an anticipated aircraft trajectory, allowing for a comparison with the trajectory derived from received data.

Ironically, increasing the number of receivers results in a smaller area of intersection which enhances the precision of trajectory estimation. However, this depends on how the receivers are spread and placed on the ground. Relying solely on the set of locations where sensors receive messages proves insufficient for attaining high accuracy.

To enhance prediction accuracy and overcome the challenges discussed, it becomes imperative to consider additional parameters, like sensor pattern, sensor ranking, ToA, and signal propagation speed. Future efforts will involve analyzing the impact of these parameters on prediction accuracy, with the aim of improving the success rate of attack detection.

5.4 Challenge in Attack Detection

Our primary focus in this work is on presenting the prediction concept and assessing potential ways for accuracy enhancement and we are not actively conducting the attack. This phase of the research is preliminary, and our emphasis is on refining accuracy and evaluating the resilience of

the proposed method against the outlined attacks. Future work will involve further testing and validation to ensure the robustness of the approach in the face of potential attacks.

6. Conclusion

The inherent openness of ADS-B messages renders them susceptible to various forms of drifting attacks. Consequently, there is a pressing need to either safeguard against or identify such attack vectors. In this study, we introduce a novel approach that involves cross-verifying the received location extracted from ADS-B messages with the predicted location generated by our methodology.

Our prediction technique is built on the premise that the message is received by a set of receivers whose coverage areas intersect. As a result, the predicted location should fall within this intersecting region. Two distinct prediction methods have been put forth: the central localization method and the weighted localization method. While the latter exhibits greater accuracy, there remains a requirement to enhance the precision of both methods by incorporating additional parameters. Our forthcoming efforts will focus on analyzing the impact of these parameters, to refine prediction accuracy and fortify the overall attack detection process.

Reproducibility statement

To ensure the reproducibility of our results, we provide the complete codebase and information about the dataset used in this study. The code, written in Matlab, is fully documented and includes detailed instructions for installation and usage. We have utilized OpenSky data, which is freely available online for researchers. The dataset can be downloaded from the OpenSky website (https://opensky-network.org/). For our test, we used these data, which are available for download [26]. However, information about sensor locations was shared with us privately upon request, and we received sample data to test our approach. Additionally, we provide the exact functions and evaluation tests used in our experiments. The code is available for download in the GitHub repository (https://github.com/afd1479/OpenSky2023). All scripts necessary to reproduce the figures presented in this paper are included.

References

[1] FAA. *Automatic Dependent Surveillance-Broadcast (ADS-B).* https://www.faa.gov/air_traffic/technology/adsb. Accessed 09.07.2023. 2023.

[2] FAA. *No Kidding: ADS-B Deadline of Jan. 1, 2020, is Firm.* https://www.faa.gov/news/updates/?newsId=90008. Accessed 06.07.2023. 2018.

[3] EASA. *EASA seasonal technical commission.* https://www.easa.europa.eu/sites/default/files/dfu/EASA_STC_NEWS_JUNE_2018.pdf. Accessed 02.07.2023. 2018.

[4] Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. "Experimental Analysis of Attacks on Next-Generation Air Traffic Communication". In: *Applied Cryptography and Network Security*. Springer Berlin Heidelberg, 2013, pp. 253–271.

[5] M. Strohmeier, M. Schäfer, V. Lenders, and I. Martinovic. "Realities and challenges of nextgen air traffic management: the case of ADS-B". In: *IEEE Communications Magazine* 52 (2014), pp. 111–118.

[6] Andrei Costin and Aurélien Francillon. "Ghost in the Air (Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices". In: *BlackHat USA* (2012), pp. 1–10.

[7] Syed Khandker, Hannu Turtiainen, Andrei Costin, and Timo Hämäläinen. "Cybersecurity Attacks on Software Logic and Error Handling Within ADS-B Implementations: Systematic Testing of Resilience and Countermeasures". In: *IEEE Transactions on Aerospace and Electronic Systems* 58.4 (2022), pp. 2702–2719. DOI: 10.1109/TAES.2021.3139559.

[8] Zhijun Wu, Anxin Guo, Meng Yue, and Liang Liu. "An ADS-B Message Authentication Method Based on Certificateless Short Signature". In: *IEEE Transactions on Aerospace and Electronic Systems* 56.3 (2020), pp. 1742–1753. DOI: 10.1109/TAES.2019.2933957.

[9] Haoran Zha, Qiao Tian, and Yun Lin. "Real-World ADS-B signal recognition based on Radio Frequency Fingerprinting". In: *2020 IEEE 28th International Conference on Network Protocols (ICNP)*. Madrid, Spain: IEEE, 2020, pp. 1–6. DOI: 10.1109/ICNP49622.2020.9259404.

[10] Kai Jansen, Liang Niu, Nian Xue, Ivan Martinovic, and Christina Pöpper. "Trust the Crowd: Wireless Witnessing to Detect Attacks on ADS-B-Based Air-Traffic Surveillance". In: *Proceesings of the Network and Distributed System Security Symposium (NDSS)*. Vol. 2021. Virtual Conference: The Internet Society, 2021.

[11] Ivan A. Mantilla-Gaviria, Mauro Leonardi, Gaspare Galati, and Juan V. Balbastre-Tejedor. "Localization algorithms for multilateration (MLAT) systems in airport surface surveillance". In: *Signal, Image and Video Processing* 9.7 (Oct. 2015), pp. 1549–1558. ISSN: 1863-1711. DOI: 10.1007/s11760-013-0608-1. URL: https://doi.org/10.1007/s11760-013-0608-1.

[12] Matthias Schäfer, Vincent Lenders, and Jens Schmitt. "Secure Track Verification". In: *2015 IEEE Symposium on Security and Privacy*. 2015, pp. 199–213. DOI: 10.1109/SP.2015.20.

[13] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. "Lightweight Location Verification in Air Traffic Surveillance Networks". In: *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security*. CPSS '15. Singapore, Republic of Singapore: Association for Computing Machinery, 2015, pp. 49–60. ISBN: 9781450334488. DOI: 10.1145/2732198.2732202. URL: https://doi.org/10.1145/2732198.2732202.

[14] Matthias Schäfer, Patrick Leu, Vincent Lenders, and Jens Schmitt. "Secure Motion Verification using the Doppler Effect". In: *Proceedings of the 9th ACM Conference on Security & Privacy in Wireless and Mobile Networks*. WiSec '16. Darmstadt, Germany: Association for Computing Machinery, 2016, pp. 135–145. ISBN: 9781450342704. DOI: 10.1145/2939918.2939920. URL: https://doi.org/10.1145/2939918.2939920.

[15] Daniel Moser, Patrick Leu, Vincent Lenders, Aanjhan Ranganathan, Fabio Ricciato, and Srdjan Capkun. "Investigation of multi-device location spoofing attacks on air traffic control and possible countermeasures". In: *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. MobiCom '16. New York, NY, USA: Association for Computing Machinery, 2016, pp. 375–386. ISBN: 9781450342261. DOI: 10.1145/2973750.2973763. URL: https://doi.org/10.1145/2973750.2973763.

[16] Martin Strohmeier, Ivan Martinovic, and Vincent Lenders. "A k-NN-Based Localization Approach for Crowdsourced Air Traffic Communication Networks". In: *IEEE Transactions on Aerospace and Electronic Systems* 54.3 (2018), pp. 1519–1529. DOI: 10.1109/TAES.2018.2797760.

[17] Martin Strohmeier, Mauro Leonardi, Sergei Markochev, Fabio Ricciato, Matthias Schäfer, and Vincent Lenders. "In Pursuit of Aviation Cybersecurity: Experiences and Lessons From a Competitive Approach". In: *IEEE Security & Privacy* 21.4 (2023), pp. 61–73. DOI: 10.1109/MSEC.2023.3265523.

[18] Thabet Kacem, Duminda Wijesekera, Paulo Costa, Jeronymo Carvalho, Marcio Monteiro, and Alexandre Barreto. "Secure ADS-B design and evaluation". In: *2015 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*. Yokohama, Japan: IEEE, 2015, pp. 213–218. DOI: 10.1109/ICVES.2015.7396920.

[19] Haomiao Yang, Qixian Zhou, Mingxuan Yao, Rongxing Lu, Hongwei Li, and Xiaosong Zhang. "A Practical and Compatible Cryptographic Solution to ADS-B Security". In: *IEEE Internet of Things Journal* 6.2 (2019), pp. 3322–3334. DOI: 10.1109/JIOT.2018.2882633.

[20] Haomiao Yang, Qixian Zhou, Dongxiao Liu, Hongwei Li, and Xuemin Shen. "AEALV: Accurate and Efficient Aircraft Location Verification for ADS-B". In: *IEEE Transactions on Cognitive Communications and Networking* 7.4 (2021), pp. 1399–1411. DOI: 10.1109/TCCN.2021.3072853.

[21] Johanna Ansohn McDougall, Alessandro Brighente, Willi Grobmann, Ben Ansohn McDougall, Joshua Stock, and Hannes Federrath. "LoVe is in the Air - Location Verification of ADS-B Signals using Distributed Public Sensors". In: *ICC 2023 - IEEE International Conference on Communications*. 2023, pp. 6040–6045. DOI: 10.1109/ICC45041.2023.10278848.

[22] Yoohwan Kim, Ju-Yeon Jo, and Sungchul Lee. "A secure location verification method for ADS-B". In: *2016 IEEE/AIAA 35th Digital Avionics Systems Conference (DASC)*. 2016, pp. 1–10. DOI: 10.1109/DASC.2016.7778003.

[23] Matthias Schäfer, Martin Strohmeier, Vincent Lenders, Ivan Martinovic, and Matthias Wilhelm. "Bringing up OpenSky: A large-scale ADS-B sensor network for research". In: *IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*. IEEE. 2014, pp. 83–94.

[24] Kevin Buchin, Maike Buchin, and Carola Wenk. "Computing the Fréchet distance between simple polygons". In: *Computational Geometry* 41.1 (2008). Special Issue on the 22nd European Workshop on Computational Geometry (EuroCG), pp. 2–20. ISSN: 0925-7721. DOI: https://doi.org/10.1016/j.comgeo.2007.08.003.

[25] Yago Lizarribar, Domenico Giustiniano, Gérôme Bovet, and Vincent Lenders. "SkyPos: Real-World Evaluation of Self-Positioning With Aircraft Signals for IoT Devices". In: *IEEE Journal on Selected Areas in Communications* 42.1 (2024), pp. 134–145. DOI: 10.1109/JSAC.2023.3322829.

[26] *Dataset for our test*. Accessed: 2024-06-21. URL: https://zenodo.org/records/12204932.