LoVe is in the Air – Location Verification of ADS-B Signals using Distributed Public Sensors

1st Johanna Ansohn McDougall Security in Distributed Systems University of Hamburg Hamburg, Germany johanna.ansohn.mcdougall@uni-hamburg.de

4th Ben Ansohn McDougall Lübeck University of Applied Sciences Lübeck, Germany ben.mcdougall@th-luebeck.de 5th Joshua Stock Security in Distributed Systems University of Hamburg Hamburg, Germany joshua.stock@uni-hamburg.de

2nd Alessandro Brighente Department of Mathematics University of Padova Padua, Italy alessandro.brighente@unipd.it 3rd Willi Großmann Helmut-Schmidt-University Hamburg, Germany grossmann@hsu-hh.de

6th Hannes Federrath Security in Distributed Systems University of Hamburg Hamburg, Germany hannes.federrath@uni-hamburg.de

Abstract—The Automatic Dependant Surveillance-Broadcast (ADS-B) message scheme was designed without any authentication or encryption of messages in place. It is therefore easily possible to attack it, e.g., by injecting spoofed messages or modifying the transmitted Global Navigation Satellite System (GNSS) coordinates. In order to verify the integrity of the received information, various methods have been suggested, such as multilateration, the use of Kalman filters, group certification, and many others. However, solutions based on modifications of the standard may be difficult and too slow to be implemented due to legal and regulatory issues. A vantage far less explored is the location verification using public sensor data. In this paper, we propose LoVe, a lightweight message verification approach that uses a geospatial indexing scheme to evaluate the trustworthiness of publicly deployed sensors and the ADS-B messages they receive. With LoVe, new messages can be evaluated with respect to the plausibility of their reported coordinates in a location privacy-preserving manner, while using a datadriven and lightweight approach. By testing our approach on two open datasets, we show that LoVe achieves very low false positive rates (between 0 and 0.001 06) and very low false negative rates (between 0.00065 and 0.00334) while providing a real-time compatible approach that scales well even with a large sensor set. Compared to currently existing approaches, LoVe neither requires a large number of sensors, nor for messages to be recorded by as many sensors as possible simultaneously in order to verify location claims. Furthermore, it can be directly applied to currently deployed systems thus being backward compatible.

Index Terms—Wireless Security, Airplane Security, ADS-B, Location Verification, H3 Geospacial Index

I. INTRODUCTION

The ADS-B protocol is a system designed to monitor aircraft positions via constantly transmitted messages. Indeed, aircraft use ADS-B to communicate with ground control stations, other aircraft, or other ADS-B enabled receivers [1]. While used as a secondary source of real-time air traffic monitoring to radar systems in metropolitan areas, it is often the single source of information in sparsely populated areas, difficult terrain and over oceans. Thanks to its broadcast nature, any ADS-B enabled device can receive ADS-B messages

and obtain sensitive information on aircraft, such as their identifier and location. Also, commercial off-the-shelf (COTS) devices with limited resources can be used to receive and collect ADS-B messages and, hence, aircraft information. This lead to the creation of the OpenSky Network [2], where aircraft are monitored in a crowd-sourced fashion. The possibility of collecting aircraft information is further facilitated by the fact that, by design, ADS-B does not include any encryption or authentication measures [3]. While this provides fast access to aircraft data to be used for real-time monitoring, the lack of security measures provides attackers with the possibility for packet spoofing. In particular, an attacker might be able to create packets and inject false information on the presence of airplanes, or report false location claims. In fact, ADS-B utilises the GNSS system to retrieve an aircraft's location, but does not encrypt or sign this information, which can hence be modified by an attacker. The consequences of these attacks range from mild (e.g., distraction on the flight deck or in the control room) to more severe such as violations of mandatory safety separations, and eventually aircraft collisions [4]. The implementation of these attacks is not only an academic issue: it also captured the attention of a wider audience [5], [6]. Therefore, it is fundamental to develop solutions for real-time and scalable verification of ADS-B messages.

To verify the authenticity of ADS-B messages, the literature proposes three different approaches. The first approach is to redesign the ADS-B protocol itself to include cryptographic means to guarantee integrity and authentication. This could for example be done using public key-based, identity-based message encryption, or via Message Authentication Codes (MACs) [3]. However, ADS-B is slow to be deployed due to different countries' regulations and the scale on which a modification has to be rolled out. Therefore, a change in protocol is unlikely to happen any time soon [7].

The second approach entails fingerprinting the signals. This does not require a change in protocol; the idea is to perform

software-based, hardware-based or channel-based fingerprinting [3].

The third approach attempts to verify the signal origin: the use of Multilateration (MLAT) has been suggested to approximate the legitimate origin, and additionally the use of Kalman filtering or group certification using signals received by other airplanes. While this system is very reliable, it is costly as it requires a gigantic amount of sensors to be deployed to receive accurate measurements: Using a test set of 8 sensors distributed within a 100 km radius, Strohmeier et al. [8] show that MLAT is only applicable to 9.73 % of the messages due to the limited amount of sensors that record the same message. Another MLAT-based approach to determine the trustworthiness of ADS-B messages is MAVPro [9], a message verification protocol. The authors use MLAT in areas with large sensor coverage, and additionally rely on various other information like predicted trajectory information, preset anchors and information on flight tracks in areas with lower coverage.

Jansen et al. [7] propose a Machine Learning (ML) approach utilizing vectorsets and evaluating ADS-B sensor response patterns by applying a random-forest model. When evaluating ADS-B data with machine learning, two problems arise. First, the feature space does not scale: Each sensor corresponds to a feature, which means that the feature space grows quadratically with increasing local coverage. In addition, an algorithm trained with data collected from a certain area cannot be transferred to other areas, because the feature space would change. Second, the chosen feature space is only slightly characteristic. Especially in areas with low sensor coverage, the necessary redundancy to generate characteristic patterns is missing.

Other approaches of signal localisation are group localisation using a fixed set of aircraft as a group [10], distance bounding [4] or fingerprinting techniques on the physical layer [3], [11] or the use of Kalman filters to estimate the future trajectory of an aircraft via recent directional information [12].

In this paper, we propose Location Verification (LoVe), a simple yet effective approach based on data collected from public sensors for the verification of aircraft location claims sent via ADS-B messages. While our approach also uses distributed sensors, we choose to use an ML-free approach to verify signal authenticity: We take measurements on the range of coordinates a specific sensor can normally receive, and can thereby infer whether a newly received set of coordinates is within the range of typically received signals. We do this by mapping historic flight data into a representative map, through which current data can be evaluated with respect to plausibility of the signal. We are, to the best of our knowledge, the first to combine this with the use of a geospatial indexing system, resulting in a system that scales well even with a very high sensor count. Furthermore, it performs knowledge-based comparisons without resorting to costly ML-computations, while maintaining sensor location privacy. It achieves realtime compatible execution times while maintaining very low false positive rates between 0 and 0.00106 and very low false negative rates between 0.00065 and 0.00334. LoVe is also easily extendable and can be integrated into various existing schemes. To this end, we contribute the following:

- We propose a lightweight, privacy-preserving classification of sensors to evaluate the plausibility of received messages.
- We implement the classification scheme by calculating the message distribution per sensor for each hexagon over the hexagonal geospacial indexing system H3.
- We test and evaluate the proposed method using realworld data provided by the OpenSky Network and Flightradar24 and multiple different hexagonal resolutions.
- We provide our implementation in a public repository¹ for others to test and use.

This paper is structured as follows: In Section II, we first provide a background on ADS-B and the system and threat model. We then introduce our approach – LoVe – in Section III. This section contains our detailed approach, data acquisition, the experimental analysis and the verification evaluation. Section IV then proposes future work and concludes the paper.

II. PRELIMINARIES AND MODELS

In this section, we provide background information on ADS-B and describe the system and threat model.

A. Background on ADS-B

ADS-B is a broadcasting system in which aircraft transmit status information every 0.5 seconds [8]. ADS-B packets are 112 bit to 272 bit long packets, depending on the link standard (UAT or Extended Squitter) [4]. They consist of a preamble, the downlink format, capabilities, the aircraft address, an ADS-B data field and a parity check. The *ADS-B Data* field contains information on altitude, speed, destination, the airplane ID and coordinates that are determined via GNSS.

A distinction is made between ADS-B Out devices, which are ADS-B transmitters, and ADS-B In, the receiving side. Ground stations require ADS-B In functionality to receive the ADS-B signals. In European airspace, airplanes heavier than 5.7 tons or faster than 250 knots are required to transmit their ADS-B data via ADS-B Out, while the use of ADS-B In for airplanes is voluntary [13].

The ADS-B broadcast system does not include a collision detection mechanism for messages, and neither any means of security to ensure confidentiality or integrity. It is therefore possible to spoof messages, inject content into transmissions and to jam communications [3].

B. System Model

We consider a system where legitimate aircraft periodically send ADS-B messages, which are then recorded by geographically distributed ground sensors. Aircraft use both their local information (e.g., identifier) and information from GNSS satellites to create packets containing air traffic monitoring

¹https://github.com/heddha/LoVe

information according to the ADS-B standard. All ground sensors share their received messages in a crowdsourcing fashion, reporting their measurements to a central server. We assume that the central server implements our location verification approach to detect the presence of malicious aircraft with the aim of securing the air traffic control-based management operations.

C. Attacker and Threat Model

We consider an attacker able to monitor and send ADS-B messages via COTS equipment, e.g., a Software Defined Radio (SDR). We also assume that the attacker may control various stationary ADS-B transmitters. Based on the Federal Aviation Administration measurements², a strong enough ADS-B signal can be captured by sensors up to a maximum radius of 240 km. A stationary ADS-B transmitter transmitting messages with forged coordinates to simulate an airplane would therefore be in a legitimate coordinate range of the airplane for about 41 minutes when simulating an airplane travelling at a speed of about 600-800 km/h and at an altitude of 10.9 km. During this time, the signals would constantly be received by the same sensors, while the transmitted coordinates would place the aircraft in an area of less and less confidence, i.e. the received coordinates are implausible with respect to the position of the sensor that recorded it.

Our threat model considers various attacks, which we elaborate on in the following.

Message Injection and Ghost Plane Injection. As ADS-B messages are neither authenticated nor encrypted, it is easily possible to send legitimate looking messages. The attacker can, on the one hand, inject messages that appear like those of nearby planes, but possibly with altered details like coordinates or altitude. This way, it may look like the existing airplane is diverging from its original path and heading in a different direction. On the other hand, the attacker can also inject messages that appear like signals sent by a legitimate looking but nonexistent aircraft, so called *ghost plane*. Both modified messages and injected ghost planes would appear on surveillance monitors and could potentially disrupt normal flight procedures and can cause other airplanes to modify their course to avoid collisions unless a verification of origin is performed on the message.

Location Spoofing. Location spoofing of ADS-B messages can be done in two different ways, i.e., i) via GNSS or ii) via message modification. In the first case, programs such as gps-sdr-sim³ can be used to spoof GNSS receivers using a lowcost SDR like the HackRF and a current map of the satellite constellations. This way, a constellation containing up to 12 virtual satellites can be spoofed; these can then "send" spoofed signals which appear like legitimate signals and are used as the basis for the positional computations [14]. The second means of location spoofing is to inject different coordinates into existing ADS-B messages in a message modification attack. ADS-B messages contain a parity check of 24 bits, where an error correction of up to 5 bits is possible. To modify the message, the significant parts of the existing message are either *overshadowed* by a stronger signal or modified via bitflipping [15]. This way, an attacker could perform a virtual trajectory modification attack and thereby cause discrepancies between the actual positions and those sent via ADS-B and therefore received by air traffic surveillance.

Attacks on Sensors. The next threat considered in this paper is an attack using ADS-B sensors. In this scenario, an attacker attempts to inject messages into the system either by adding an own sensor and inserting illegitimate messages through it, or by flooding particular sensors with a large number of bogus messages.

III. LOCATION VERIFICATION - LOVE

While various approaches to location verification of ADS-B signals have been proposed, (c.f. Section I), we suggest a mask-based approach that can verify whether specific coordinates received from a sensor are legitimately within the range of coordinates that the sensor is able to sense.

In order to evaluate the recorded coordinates with respect to signals received by the same and other sensors in the area, we chose to use the H3⁴ hexagonal geospacial indexing system. It provides evenly distributed hexagons in various resolutions, spread over the whole world map. This approach is more applicable to the scenario than, e.g., a mercator projection or others, as the cell size is fixed and therefore delivers comparable results for each cell. Using H3, one can choose between 16 different resolutions⁵ per cell. The resolution, or cell size, defines the total amount of hexagons and the individual cell size, ranging from approximately four million square kilometres to one square metre. With respect to ADS-B signals, which often travel distances of a few hundred kilometres, several of the resolutions can be disregarded for being too big or too small. With our approach, we only focus on the resolutions 2-7 as shown in Table I, as their dimensions are most fitting with respect of the propagation of ADS-B signals.

The preparations of the LoVe-scheme were done as follows. We first acquire a dataset containing ADS-B messages of a whole day (see Section III-B for details), process them, and feed them into a Postgres database, which then contains the sensors that recorded them, and the latitude and longitude that was contained within each message. We assume the data set recorded on that day is reliable and a trustworthy basis. From it, we construct one table, in the following called the *amount table*, per resolution, as can be also observed in fig. 1: It contains the h3id (the H3-representation with respect to the chosen resolution), calculated from the latitude, longitude, and respective resolution, together with the sensor ID as the primary key and the respective amount of messages captured by the sensor within the area as the third column. The total

²https://www.faa.gov/air_traffic/technology/adsb

³https://github.com/osqzss/gps-sdr-sim

⁴https://h3geo.org

⁵https://h3geo.org/docs/core-library/restable/

 TABLE I

 Size and amount of hexagons per resolution in the H3 geospacial index and respective amount of entries in the database

Resolution	No. of	Avg. Hexagon	Sensor-location-pairs		Avg. No. of msg. per sensor-location-pair		Test time (s) for 200 000 entries	
	Hexagons	Area (km ²)	OpenSky	FlightRadar	OpenSky	FlightRadar	OpenSky	FlightRadar
2	5870	86 801.78	9963	26 133	208 721.51	535.43	0.549	0.586
3	41150	12393.43	32 339	65220	64302.93	214.54	0.602	0.678
4	288 110	1770.35	143143	223018	14527.38	62.74	0.717	0.806
5	2016830	252.90	776226	758857	2678.98	18.44	1.350	1.322
6	14117870	36.13	4442948	1948721	468.04	7.18	5.138	2.483
7	98825150	5.16	25086048	3741523	82.89	3.74	29.049	4.438

📲 h3id 🔹 👻	1 🧖 sensor_id 🗧	🔳 amount 🗘
843f76bfffffff	-1408233591	16024
843f769fffffff	1408236197	38
843f769fffffff	90310	11267
843f769fffffff	-1408233591	15831
843f767fffffff	90310	209
843f767fffffff	-1408233591	7147
843f765fffffff	90310	428
843f765fffffff	-1408233591	6472
843f763ffffffff	1408236197	119
843f763ffffffff	90310	1296
843f763ffffffff	-1408233591	13932
843f761fffffff	1408236197	20

Fig. 1. Sample data from the amount table in resolution 4, sorted by h3id.

amount of sensor-location pairs per resolution for our data sets can be seen in Table I, as well as the average number of messages per sensor-location pair. Since the amount of cells rises the smaller the hexagon area is, higher resolutions cause a higher amount of sensor-location pairs as well. To illustrate the message distribution further, Figure 2 shows a partial map of Europe distributed into H3-cells of resolution 4, with the colouring adapted to the maximum amount of messages captured by one sensor per hexagon.

A. Verification Approach

Our approach at origin verification of an incoming message M with coordinates lat and long begins by calculating the respective h3id for the two coordinates and the resolution. LoVe consists of the following two phases:

- 1) Check whether the coordinates correspond to a *h3id* in which the sensor has recorded messages before, using the *amount-table*. If so, the signal is classified as legitimate.
- 2) If not, check whether the specific sensor is known at all in our dataset.

If a sensor then receives spoofed coordinates that are not within the range it typically receives coordinates in, we assume that the message is an illegitimate one.

B. Data Acquisition and Structure

We test our approach using two different data sets recorded on the same day, July 23rd 2021: The first set was provided



Fig. 2. A map showing parts of Europe covered with H3 hexagons of resolution 4. The colours denote the amount of messages received by the sensor that received the most messages in the cell.

by the OpenSkyNetwork⁶. They gave us access to their historical database, containing the stored ADS-B messages from their crowdsourced network of sensors. The second set was supplied by FlightRadar24⁷, a commercial website offering live airtraffic monitoring. To make our analysis comparable to others, we limit the range of both data sets to only contain entries from Europe, with the latitude within the boundaries of 30 and 75 and the longitude between -25 and 45.

To test our implementation, we required a labelled test set containing both false and true data. For the OpenSky-data, we chose the following approach to construct a test set: For legitimate flight data, we retrieved several hours worth of ADS-B messages from the following day, July 24th, 2021 and limited the amount to 100 000 randomly chosen true entries. We additionally generated false test data using the following approach: For every single sensor in the set, all coordinates in which a signal has ever been received are collected. From this set, we determine the minimum and maximum latitudes and longitudes it can legitimately receive

⁶https://openskynetwork.github.io/opensky-api/

⁷https://www.flightradar24.com/



Fig. 3. Comparison of the test runs using labelled test data sets of OpenSky (left) and FlightRadar (right), with our LoVe algorithm in resolution 4.

information in. We then add or subtract a uniformly distributed random float between 0.1 and 10 to or from it, according to a random boolean⁸. While this is a very simple approach, it ensures that the coordinates are outside the range of expected coordinates and thereby simulates a GNSS-spoofing attack or an ADS-B message injection attack using false coordinates. For the FlightRadar24-test set, we were unfortunately unable to acquire a second data set from the company. Instead, we randomly chose 100 000 messages as true labelled test messages and removed them from the original set. For false labelled messages, we chose the same approach as described above for the OpenSky data set.

Both data sets are very different in content and structure, apart from the obvious differences in naming the sensors: The FlightRadar-set is first split up in flights surveilled and, for each flight, contains an additional CSV-file including all sensors that recorded the messages and for each message the received altitude, coordinates, velocity and others. Once limited to the aforementioned coordinate boundaries, it contains 14 092 420 messages in total recorded by 11 594 sensors. We additionally excluded another 100 000 random messages from the set in order to construct a legitimate test set, which leaves us with 13 992 420 messages in total, recorded by 11 582 sensors. On average, every sensor in this data set recorded approximately 1208.12 messages. In this data set, only about 0.18 % of the messages were recorded by more than one sensor.

The OpenSky data set contains far more messages recorded by far fewer sensors. This dataset contains 160 526 553 distinct ADS-B messages recorded by 971 sensors. If the same message was recorded by different sensors, it was stored with an array containing all receiving sensors; the actual amount of single messages is therefore 2 079 492 408, while the list of sensors that captured one particular message has an average length of 12.95. Every sensor recorded an average of 2 141 598.77 messages.



Fig. 4. Accuracy and relative time needed to verify location accuracy using different resolutions. FR denotes results obtained with FlightRadar dataset, whereas OS those obtained with the OpenSky dataset.

For every resolution, our verification test is performed on a data set with which a reliable classification can be performed, as will be demonstrated in the following.

C. Verification evaluation

To test our approach, we use both false and true labelled data, generated or acquired as described in Section III-B. Our evaluation achieves very low false positive rates between 0 and 0.001 06 and very low false negative rates between 0.000 65 and 0.003 34 for both the OpenSky and FlightRadar datasets, as can be seen in Figure 3. The time required to test 200 000 entries vary per resolution and span from 0.549 to 29.049 seconds for the OpenSky set and from 0.586 to 4.438 seconds for the FlightRadar set, as can also be seen in Table I. The surprisingly high maximum execution time of the tests on the OpenSky dataset is directly linked to the amount of sensor-location-pairs of 25 million in resolution 7.

The overall test results show that test results are best at resolution 4, while both higher and lower resolutions cause higher false negative and false positive rates. Larger table size simultaneously causes longer evaluation times. Both the relative execution time as well as the respective accuracy can be observed in Figure 4. The comparison between the FlightRadar and OpenSky data sets also shows that our solution scales very well even with a much larger amount of sensors present, as can be observed in Table I: even a data set containing 10 times the amount of sensors causes only slightly higher computation times, since our approach utilises the sensor-location-pairs as a basis for comparison and does not require to evaluate the message content.

Other approaches at location verification of ADS-Bmessages require multiple sensors to verify the message origin: Some attempt multilateration with a reduced amount of sensors [8], [9], others use ML approaches that require a large sensor set [7]. While their approaches only work with a large sensor set, our tests of the FlightRadar dataset show that our approach is also usable in the absence of a dense sensor network: The FlightRadar dataset does not contain any mentions on

⁸Notice that our approach in generating attacker data can not provide coherent spoofed values in successive time instants. This is not a problem, as our approach does not exploit any form of correlation or time dependency.

duplicate entries, e.g., single messages recorded by multiple sensors. To identify duplicates and see whether other data analysis strategies like those suggested in [7] are applicable to this data set, we used the following approach: For every message, we calculated a hash over several entries of the ADS-B data field (altitude, heading, latitude, longitude, speed). Only about 0.18% of the messages had duplicate hashes and had therefore been recorded by more than one sensor. While the approaches mentioned above are inapplicable to such a data set, our approach is not impeded by it.

D. Comparison to ML-Baseline

To establish a speed comparison between the LoVe implementation and a machine learning approach, we chose to use a Support Vector Machine (SVM), a well known, often used and easy to implement ML-classifier. As in the LoVe implementation, we use batch learning mode, i.e. an entire set of of training examples used offline. We use the same training data set as for the LoVe classifier but scale the features (sensor id, latitude and longitude) to a scale from 0 to 1. We parameterize the SVM with values obtained from a hyperparameter search: C=10 and gamma=4641. Training the SVM took approximately 48h, since the training time of SVMs grows quadratically in the number of training samples. Afterwards, we tested it against a FlightRadar data set with 200 000 data records. 832 of these records were misclassified (all of those are false negatives); compared to LoVe, the false negative rate is increased by 0.00491. The false positive rate and true negative rate stay at 0.0 respectively 1.0.

Testing the 200 000 samples took another 91.3 minutes, which is demonstrably worse than LoVe. While no training at all is needed for LoVe, the time to test the samples using SVM is increased by at least a factor of 6000.

IV. CONCLUSION

The LoVe setup is very flexible, since we do not require to retrain models but only rely on minimal modifications, e.g. a sensor can easily be added or deleted from the database manually. We neither depend on formal groupings of the messages: While the OpenSky messages are grouped by the set of sensors that received them, our approach works well with single messages that were observed only by a single sensor as well, as shown by the good results on the FlightRadar dataset.

Our suggested approach mainly consists of one database table per resolution, used to compare new ADS-B messages to previously recorded ones, and verify their plausibility with respect to the sender that recorded it. The underlying mapping scheme that works as the basis for the sensor map makes use of the H3 hexagonal indexing system that provides intrinsic functions to modify the resolution and measure the distance between different cells. This way, we are able to compare various resolutions and find the best consensus between low false-negative rate and comparison speed. Additionally, our approach scales very well, both with respect to a high amount of sensors, messages and a larger surface area, due to the following constraints:

- The number of possible cells is limited to a maximum for each resolution;
- A large amount of sensors can be represented without significantly impacting the computational cost;
- With a large amount of messages, only the number of messages recorded per sensor and thereby the trustwor-thiness of the sensor is increased;
- Since we omit the message content, our system is particularly lightweight.

LoVe is also privacy-preserving with respect to the location of the sensor: we don't require its position, but rather its reception area. The amount of messages captured does not necessarily provide information about the sensor location, rather about the direction the antenna is directed towards. While it is likely that the actual sensor location is within a reasonable distance of the area with most messages received, our calculation neither requires nor performs exact localisation.

Altogether, this makes our approach an easy-to-use basis that can be extended and integrated into already-established schemes.

REFERENCES

- I. C. A. Organization, "ADS-B Implementation and Operations Guidance Document," 2014.
- [2] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. 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.
- [3] Z. Wu, T. Shang, and A. Guo, "Security Issues in Automatic Dependent Surveillance - Broadcast (ADS-B): A Survey," *IEEE Access*, vol. 8, pp. 122 147–122 167, 2020.
- [4] M. Riahi Manesh and N. Kaabouch, "Analysis of vulnerabilities, attacks, countermeasures and overall risk of the Automatic Dependent Surveillance-Broadcast (ADS-B) system," *International Journal of Critical Infrastructure Protection*, vol. 19, pp. 16–31, 2017.
- [5] A. Greenberg, "Next-gen air traffic control vulnerable to hackers spoofing planes out of thin air," *Forbes Magazine. Retrieved September*, vol. 10, p. 2014, 2012.
- [6] A. Costin and A. Francillon, "Ghost in the Air(Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices," 07 2012.
- [7] K. Jansen, L. Niu, N. Xue, I. Martinovic, and C. Pöpper, "Trust the Crowd: Wireless Witnessing to Detect Attacks on ADS-B-Based Air-Traffic Surveillance," 02 2021.
- [8] M. Strohmeier, V. Lenders, and I. Martinovic, "Lightweight Location Verification in Air Traffic Surveillance Networks," in *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security*. New York, NY, USA: Association for Computing Machinery, 2015, p. 49–60.
- [9] A. Darabseh, H. AlKhzaimi, and C. Pöpper, "MAVPro: ADS-B message verification for aviation security with minimal numbers of on-ground sensors," in *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, 2020, pp. 53–64.
- [10] K. Sampigethaya and R. Poovendran, "Security and privacy of future aircraft wireless communications with offboard systems," in 2011 Third International Conference on Communication Systems and Networks (COMSNETS 2011), 2011, pp. 1–6.
- [11] B. Danev, D. Zanetti, and S. Capkun, "On Physical-Layer Identification of Wireless Devices," ACM Comput. Surv., vol. 45, no. 1, dec 2012.
- [12] J. Krozel, D. Andrisani, M. Ayoubi, T. Hoshizaki, and C. Schwalm, "Aircraft ADS-B Data Integrity Check," 09 2004.
- [13] P. Holland-Moritz. Längere Frist für ADS-B-Ausrüstungspflicht. [Online]. Available: https://www.aerokurier.de/business-aviation/ wegen-corona-laengere-frist-fuer-ads-b-ausruestungspflicht/
- [14] Wikipedia. GPS-Spoofing. [Online]. Available: https://de.wikipedia.org/ wiki/GPS-Spoofing
- [15] M. Strohmeier, V. Lenders, and I. Martinovic, "Security of ADS-B: State of the Art and Beyond," *IEEE Communications Surveys & Tutorials*, vol. 17, 07 2013.