

Lightweight Location Verification in Air Traffic Surveillance Networks

Martin Strohmeier
University of Oxford, UK
martin.strohmeier@
cs.ox.ac.uk

Vincent Lenders
armasuisse, Switzerland
vincent.lenders@
armasuisse.ch

Ivan Martinovic
University of Oxford, UK
ivan.martinovic@
cs.ox.ac.uk

ABSTRACT

In this work, we develop a realistic threat model for attacks on modern air traffic communication networks and show that current state-of-the-art countermeasures such as multilateration are insufficient. We propose two alternatives, a statistical location verification technique and a grid-based location estimation approach, to deal with the identified threats. We evaluate our proposals using real-world flight data and quantify their effectiveness in terms of aircraft location accuracy, resilience to message injection attacks, attack detection speed, and surveillance coverage.

Our results show that the statistical verification approach can increase the effective air traffic surveillance coverage compared to multilateration by a factor of more than 100. Concerning our location estimation method, we find that the mean aircraft location accuracy can be increased by up to 41% in comparison with multilateration while also being able to pinpoint ground-based attackers with a mean error of 145 m for air-based attackers. Finally, we demonstrate that our proposal is lightweight as it does not require any changes to the existing air traffic protocols and transmitters, and is easily implemented using only low-cost hardware.

Keywords

ADS-B, air traffic control, air traffic security, aircraft localization, location verification

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: General—Security and protection

1. INTRODUCTION

Automatic Dependent Surveillance - Broadcast (ADS-B) is currently rolled out as part of next generation air traffic control (ATC) networks in most of the world's airspaces. The ADS-B protocol is intended to facilitate the safe and efficient transportation of more than 2 billion passengers per

year by 2020 and provide the future backbone for the regulated control of Unmanned Aerial Vehicles (UAV). By now, the major airlines are upgrading their fleets with this new technology, and ADS-B signals are broadcast by 70-80% of commercial aircraft in airspaces in Europe and America [1].

While the aviation industry has a long tradition of imposing strong safety requirements on top of any technical and operational design decisions, it has neglected to consider security requirements the ADS-B's protocol, which does not provide message authentication or data encryption. The same type of negligence in protecting networked systems is evident in industrial control systems [2] or power grids [3]. The omission of security primitives in the ADS-B communication protocol is, however, particularly problematic as ADS-B messages are broadcast over the wireless channel. The system is therefore susceptible to various kinds of well-known message injection and manipulation attacks.

Aviation authorities long argued that the security issues found did not constitute a real threat in ADS-B because efforts to launch critical attacks were considered too difficult and costly. This view on the security of ADS-B has changed recently after hackers at Black Hat and DEFCON [4, 5] and academic researchers [6] reported that they were able to successfully launch attacks on air traffic control networks using only low-cost software-defined radios. Since then, security has become a top priority on the authorities' agendas and work groups were installed to address this most urgent problem [7]. While there have been no proven (or disclosed) attacks in the wild, security issues in ADS-B have also played a role in discussions about the recent disappearance of Malaysian aircraft MH370 (e.g., [8]).

In this work, we investigate how the threat of false-data injections can be mitigated in air traffic surveillance networks. The attack scenario is a malicious party who wants to falsify the recognized air picture by injecting false messages about aircraft locations. We consider different threat models where the attacker is fixed, mobile, on the ground, or in the air, and propose and evaluate methods to detect these false messages. None of our methods require changes to current standard or to the aircraft's legacy hardware equipment. This lightweight approach is particularly important given aviation's long adoption and certification cycles.

We make the following contributions in this paper:

- We identify and discuss a relevant threat model comprising distinct types of attackers that threaten ADS-B surveillance networks and general air traffic security.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CPSS'15, April 14, 2015, Singapore.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3448-8/15/04 ...\$15.00.

<http://dx.doi.org/10.1145/2732198.2732202>.

- We propose lightweight countermeasures based on location fingerprints which are used for statistical verification of flight data, and to directly locate aircrafts.
- We evaluate our approach on real-world data and show that it performs better in wide area aviation than currently utilized countermeasures such as multilateration. Compared to the latter, our approach is cheaper and more scalable, and improves surveillance range, detection speed, and location accuracy for both legitimate aircraft and attackers in real-world environments.

The remainder of this paper is organized as follows. In Section 2 we provide a short introduction to ADS-B and air traffic control. Section 3 explains our threat model. Section 4 discusses the aircraft location problem and its characteristics. Section 5 describes the design of our approach, whereas Section 6 details the experimental setup. Section 7 evaluates the scheme against real-world flight data and injection attackers. Section 8 summarizes and concludes this work.

2. OVERVIEW OF ATC SECURITY

This section gives a brief overview of air traffic control security. We explain the vulnerabilities identified in the literature and examine the proposed countermeasures by comparing them with ATC system requirements and constraints.

2.1 ATC Protocols

To obtain the location, altitude and identity of an aircraft for navigation, today's ATC relies on traditional primary surveillance radar (PSR) and interrogation-based secondary surveillance radar (SSR), using so-called modes of which Mode A, C and S are currently in use [9]. Neither PSR nor SSR surveillance technologies are able to cope with the increased air traffic density due to their limited accuracy and coverage. The introduction of ADS-B constitutes a significant change in air traffic surveillance, replacing expensive PSR installations and SSR-based interrogations of aircraft. Every ADS-B equipped aircraft has an onboard Global Navigation Satellite System (GNSS) receiver to fetch their own location and velocity, which are then broadcast in an ADS-B message, typically twice per second each. These messages are processed by ATC stations on the ground. This type of surveillance is *cooperative* and *dependent*, since cooperation by the aircraft is needed for broadcasting and the data is retrieved by onboard sensors. While ADS-B is still rolled out, ADS-B data is already transmitted by most aircraft.

2.2 Vulnerabilities

Since there is no encryption of ADS-B message content, any passive adversary with a receiver listening on the 1090 MHz channel can eavesdrop on messages sent out by aircraft. While this may pose potential risks of privacy breaches (e.g., the possibility of tracking private planes), this is a by-product of ADS-B's open design and such honest-but-curious attackers are not considered further in this work. Similarly out of scope are non-selective jamming attacks, which are inherent to the wireless medium and must be dealt with through conventional anti-jamming techniques.

Outside these inherent vulnerabilities, an attacker that can actively interfere with ATC communication poses a much more severe threat to security. With the introduction of software-defined radios (SDR) and receiver implementations freely available on the Internet, a somewhat knowledgeable




Figure 1: Result of an ADS-B injection attack. The radar screen shows legitimate ADS-B equipped aircraft and aircraft detected by PSR/SSR surveillance alongside aircraft injected by an adversary. On the screen, the injected aircraft (ADSB5, ADSB6) are indistinguishable from real ones.

attacker can exercise full control over the ADS-B communication channel. This means that the attacker is able to modify and inject ADS-B messages into ATC systems and manipulate radar screens, affecting the situational awareness of pilots and controllers. There are a multitude of such active attacks (for an overview see [10, 11]) which only require standard off-the-shelf hardware to execute, including:

- **Ghost Aircraft Injection / Flooding:** As demonstrated in [4, 6], injected ADS-B messages, claiming to be non-existing aircraft (so-called ghosts), are hard to detect. Especially under difficult weather conditions, injecting one or many different ghost aircraft may lead to serious distress for ATC and pilots.
- **Aircraft Disappearance:** Selectively jamming (as described in [12]) all ADS-B messages by a single aircraft would make the aircraft vanish from the ADS-B channel, requiring ATC to use backup systems.
- **Aircraft Spoofing:** Every ADS-B message requires an identifier which simply be replaced with an arbitrary one. Copying a known and trusted aircraft identifier may, for example, reduce the likelihood for alarms when an unexpected object is detected on the radar [6].
- **Virtual Trajectory Modification:** This attack is executed by selectively jamming an aircraft's messages and replacing them with modified location and heading data. This leads to a discrepancy between the real aircraft position and the one received by ATC [6].

2.3 Why is securing ADS-B a problem?

In the following, we discuss the crucial aspects and requirements to achieve effective protection of NextGen air traffic communication networks.

Legacy requirements. It is important that proposed security designs should not require changes to the existing protocols. This legacy requirement is common to slow-changing industries such as aviation. ADS-B, for example, has been in development since the early 1990s and is only now being deployed, over two decades later. Hence, urgently needed countermeasures against ADS-B attacks ought to work alongside the current system without disrupting it [13, 14].

Cost effectiveness. Cost has regularly been named as a main driver of NextGen air traffic adoption [15]. Conventional PSR and SSR technologies are both more expensive to deploy and experience much higher wear and tear compared to ADS-B. The International Civil Aviation Organization (ICAO) specifies the technological cost of using PSR to monitor an en-route airspace at \$10-14 million, while Mode S surveillance is priced at \$6 million and ADS-B at a significantly cheaper \$380,000 [16]. Being able to rely only on secure and accurate ADS-B data would be very cost-effective and thus a crucial argument considering the massive investments already made during the transition to ADS-B.

Loss tolerance. As has been studied before (e.g., [13]), there is substantial message loss on the 1090 MHz channel shared by ADS-B and other ATC communication protocols. With loss rates often exceeding 50% and peaking at up to 90% in airspaces with fewer than 100 aircraft in transmission range,¹ the impact of any additional measure on the stability and reliability of the whole network must be considered. Recent studies have looked at the possibility of introducing cryptographic methods to the ADS-B protocol (e.g., [18]). While most of these works have considered theoretical message and communication overhead introduced by cryptography, real-world loss figures indicate a much deeper issue. While cryptography can compound the loss problem, message loss itself can also affect the reliability and efficiency of a cryptographic method used to secure ADS-B.

Fig. 2 shows an example of the relationship between loss and distance from our data. The regression equation is

$$y \sim 0.31451 + 0.000985 \cdot \text{distance [km]}$$

showing a baseline message loss of around 30% and approximately one additional percent of loss per 10 km. Such numbers constitute a severe problem for the operational procedures of ADS-B as air traffic density increases further world-wide. They also pose a problem for security approaches that rely on frequent message reception.

Openness. ADS-B was designed to be an open protocol, i.e., encryption of the message content was not considered desirable in the planning stages. Despite the availability of cheap and powerful SDRs which sparked the recent concerns over authentication and integrity, the flight authorities in Europe and the US have shown no interest in changing this open approach.² It is likely that even for future generations of similar data communication networks, passive listening will be possible and even desired, as legacy issues, compatibility and administrative differences across countries and airspaces will continue to play an important role.

¹In busy airspaces such as LA, aircraft quantity can easily exceed 200 within ADS-B transmission range, a number that is expected to grow further in the future [17].

²As stated by ICAO: <http://goo.gl/IhGvoB>




Figure 2: Relationship between loss and aircraft distance from a receiving ADS-B sensor. The (least squares) regression line shows that there is a basic message loss of more than 30% with approximately 1% of additional loss per 10 km.

2.4 Existing Countermeasures

That ADS-B security is lacking has been well-known for a long time, warnings by interested parties can be traced back as far as 1999.³ However, only in recent years has it become a prominent topic, with broad reporting in the mainstream press [19, 20] prompted by various talks at security conferences such as Black Hat and DEFCON [4, 5, 21].

In this section, we discuss the previously proposed countermeasures and explain why they are insufficient:⁴

1. **Multilateration:** Hyperbolic positioning, or multilateration (MLAT), has been proposed throughout the academic literature [22, 23, 24] and can certainly be considered the standard recommendation when it comes to the mitigation of threats to the integrity of ADS-B messages. An MLAT system features four or more receivers in several locations that pick up the same signal and measure the time difference of arrival. Many algorithms (for an overview, see [25]) exist to solve the resulting system of equations to find the sender's origin and hence establish its legitimacy. ICAO itself considers the use of navigational backup systems such as MLAT an important part of the security concept for NextGen systems [14]. MLAT has long been utilized for both civil and military surveillance as it is *co-operative* and *independent*, making it a viable method to verify the positional claims made by aircraft using their radio signals. Unfortunately, it is very expensive to deploy enough sensors to provide reliable and thorough coverage of an airspace with MLAT. Currently, it is only an option in limited airspaces around airports where it is in common use as of today.
2. **Data fusion:** Related, to the above point, the fusion of ADS-B data with other navigation systems has also been proposed widely in the academic literature, see e.g., [22, 23]. These systems can include Mode-S, PSR or other ATC systems that are currently in use. However, the use of these systems defeats the original purpose of ADS-B, cost savings and increased accuracy, both of which are fundamental for the next generation of air traffic control. Furthermore, ADS-B has been developed to be the sole means of surveillance in areas where current ATC is not sufficient, such as in large parts of Australia and Canada, or over oceans.

³See e.g. <http://www.airport-corp.com/adsb2.htm>.

⁴The authors in [9] provide a more detailed overview of the current state of ADS-B security.

3. **Cryptography:** Several authors (e.g., [10, 26]) have conducted a holistic security analysis of the NextGen implementation plan and discussed cryptographic means to deal with authentication issues in ADS-B. They acknowledge the non-trivial difficulty of solving the open questions about key management and distribution in a technically and politically complex environment. Even if this can be overcome in the future, message and communication overhead on a channel already suffering from severe message loss present another major concern for a traditional authentication protocol. Lightweight encryption has also been proposed in other works on ADS-B security (e.g., [4, 22, 27]) and some potential approaches have been analyzed more closely in [18, 28]. One main obstacle for cryptographic adoption is the small size of commercial ADS-B messages (only 56 bit are available in the 1090 Extended Squitter format) which makes a fundamental change to a new protocol necessary. This is not only undesirable but infeasible for the foreseeable future due to the legacy constraints of the industry, where a typical protocol development cycle takes decades.
4. **Kalman filters:** Kalman filters are an estimation technique, which predicts the future trajectory of a flight based on recent directional information. Several works have suggested Kalman filters for aircraft intent verification to detect obvious discontinuities in transmitted ADS-B data [22, 29]. However, this is mostly an approach to detect less sophisticated attackers injecting unrealistic flight data and would not protect against even a simple replay attack for example, or a valid-looking trajectory change (known as frog-boiling attack [30]). As described in the next section, we use a stronger adversary model in our work.

In summary, none of these countermeasures have proven to provide sufficient security for data communication networks used in ATC. Despite this, MLAT is enjoying widespread acceptance in modern aviation, hence we consider it in more detail later in this work and use it as a baseline comparison for our own verification scheme.

3. THREAT MODEL

Taking into account the vulnerabilities identified above, we develop a threat model for ADS-B false-data injection attacks. We describe a realistic set of possible attacker types, based mainly on their resources and their intentions. All attacker types cover distinct threats that NextGen ATC systems need to consider, driven most notably by technological advances such as cheap off-the-shelf SDRs and drones.

False-data injections

We base our threat model on false-data injection attacks, i.e., an adversary who seeks to inject outdated, fake or otherwise incorrect data into an ADS-B surveillance system. We consider this the most important problem to study since such injections are the basis of most of the attacks on the ADS-B system described in the literature and can have the most subtle, yet devastating effects.

In the scope of this work, we identify two main scenarios for an attacker injecting data onto the wireless communication channel: replay attacks and message injections.




Figure 3: Graphical overview of the four distinct attacker types. Attacker 1 is stationary on the ground, attacker 2 is mobile on the ground, attacker 3 has mobility up to a few hundred meter above ground. All three attackers inject a ghost aircraft onto the channel. Attacker 4 is a commercial aircraft using its legitimate transponder to send out wrong ADS-B messages to conceal its true position.

- **Replay Attack:** This attack captures real ADS-B data in the area and plays it back at a later time without modification. This is a traditional replay attack, which is trivial, considering the ADS-B protocol has no built-in authentication. Concretely, we assume that the attacker captures a given flight’s ADS-B messages (positional, velocity, identification and potentially others) and plays them back in the same order.
- **Message Injection:** This type of attack injects a new ghost aircraft created from scratch, by creating correctly formatted ADS-B messages according to the specified standards [31]. We also assume the attacker crafts his messages with a legitimate identifier and reasonable flight parameters (e.g., believable altitude and speed) to create an aircraft which on an ADS-B radar is indistinguishable from a legitimate one. This also forms the basis of virtual trajectory modification, virtual aircraft hijacking and aircraft spoofing attacks [9].

For both scenarios, we adopt a non-naive attacker that has a sufficient amount of knowledge to inject valid-looking position messages. In other words, we assume these ADS-B messages are well-formed and their content is reasonable and able to withstand a superficial sanity check. The attackers have different mobility models which can influence the temporal credibility of their positional claims as their physical positions and signal characteristics change. In this section, we provide a concrete description of the attackers’ characteristics (see also Fig. 3 for an illustration).

3.1 Ground-based and stationary

The typical ground-based and stationary attacker wants to exploit the well-known and publicized security holes in ADS-B with existing, easy-to-use attacks and typically possesses fewer technical means. Using a programmable ADS-B transponder such as a software-defined radio, the attacker listens in to legitimate radio communication on the 1090 MHz channel, modifies the aircraft identifier and/or information such as position and velocity and plays it back.

3.2 Ground-based and mobile

The second type of attacker also uses an SDR to inject data into the ADS-B system but is mobile. Concretely, we assume the attacker is using a battery-powered laptop and utilizes a ground-based vehicle to achieve (somewhat limited) mobility. This enables the attacker to change position with an assumed speed of 50km/h. While they are normally constrained by the given infrastructure, we assume they can freely roam on the ground within their speed limits.

3.3 Low airspace and mobile

Attacker 3 is mobile within the limits of a typical unmanned aerial vehicle. Without loss of generality, we assume a hand-held commercial UAV, for example a standard model working within 2km range, up to an altitude of approximately 600m and with a vertical top speed of 100 km/h. In general, a UAV is a versatile ADS-B sender and a much more flexible tool for an attacker than ground-based solutions. The airborne attacker seeks to emulate the physical characteristics of a commercial aircraft (or other UAV using ADS-B for navigation and collision avoidance in the future) much more closely than the previous threat models.

3.4 High airspace and mobile

Attacker 4 differs from the previous three types in the fact that the sender is actually a legitimate aircraft. While the other threat models assume that the messages are injected onto the ADS-B channel by outsiders seeking to cause confusion within air traffic control systems, we now consider the case where a malicious person has control over a commercial aircraft and its ADS-B transponder. The inside attacker tries to conceal the real position of the hijacked aircraft by sending out fake positional ADS-B data. When the aircraft is diverted from its original course, its messages claim that everything is normal, prompting no action from authorities relying on ADS-B. Even where this *virtual trajectory modification* variant is picked up by other systems such as PSR, this would delay detection, and consequently the initial response, in a situation where even seconds can be crucial.

4. CONSIDERATIONS ABOUT AIRCRAFT LOCALIZATION

In this section, we discuss several characteristics inherent to the air traffic location problem and distinguish between location verification and location estimation in this context. We also discuss MLAT in more detail, a well-established navigation technique in aviation that can independently detect injected ADS-B positions. As we argue in this section, while still a viable solution in some areas, the real-world applications of MLAT in the air traffic surveillance space are considerably limited and require significant improvement.




Figure 4: An example illustrating the calculation of expected TDOAs. The assumed distance of the sender to both receivers is multiplied by c . Subtracting the smallest time t_i from the other times gives the TDOAs relative to receiver i .

4.1 Problem Characteristics

We identify the following characteristics distinguishing the aircraft location verification problem from other wireless localization problems (e.g., in wireless sensor networks or vehicular ad hoc networks):

- **Outdoor line-of-sight environment:** Contrary to many location estimation and sender verification problems found in academic research, the aircraft location problem is naturally outdoors. On the 1090 MHz channel, the line of sight (LOS) is a crucial factor in receiving signals. We require an outdoor LOS propagation model for our work in terms of loss and propagation.
- **Vast distances:** In *wide area* surveillance, the distances covered are naturally much larger than in more local or indoor problems. Aircraft flying at cruising altitudes (typically 35000 feet or higher for commercial aircraft) can be observed up to the radio horizon of 400 km or more. This is orders of magnitude larger than typical indoor location problems. While we gain most over such large distances, our approach can easily be adapted for airport surveillance, too.
- **Few multipath effects:** At typical aircraft cruising altitudes, we experience comparably few diffractions leading to multipath effects that influence signal characteristics. This enables us to use simpler theoretical models than in more complex indoor and multipath-rich environments. Most importantly, the propagation timings between aircraft positions and sensors can be approximated easily by using the speed of light c .

4.2 Location Verification vs. Estimation

One popular means proposed to secure navigation protocols that do not inherently provide encryption has been **location verification**, which is an umbrella term for a set of methods that can be used by receivers to independently verify the location claim of a sender. This is crucial to detect an intruder who is replaying or injecting false data.

Location estimation constitutes a subset of location verification methods whereby the actual location of the sender is estimated and compared to the claimed location. By comparing the claimed data with our own, we can easily verify if the sender is near the claimed position and thus identify unlikely or impossible flight trajectories. Furthermore, having estimated the location of the signal directly provides its origin and opens up further options of dealing with an attack. There has been a large body of work in the popular area of




Figure 5: Geometric dilution of precision. The circles show the measurement errors of the respective receivers; the intersections demonstrate the area where the true location of the measured object can be found.

location estimation methods. Most relevant for the aircraft location verification problem are methods based on the time differences of arrival between the receivers of a signal.

The **time differences of arrival** (TDOA) of a received signal between multiple sensors are a primitive that can be used to establish the possible location(s) of a sender. Using an outdoor LOS propagation model suitable to the aircraft location verification problem, we can calculate the absolute propagation times of an ADS-B signal to the ground stations by dividing the distances $d_1, \dots, d_n$ between the sender and each of the stations by the speed of light c (see Fig. 4).⁵

4.3 The Drawbacks of Multilateration

MLAT is a proven and well-understood concept that is used in civil and military navigation and already serves as a backup for ATC around some airports. It is the consensus solution in academia and aviation circles regarding short- and medium-term security against injections of ADS-B position messages. However, there are potential pitfalls:

1. MLAT is highly susceptible to noisy environments and even small measurement errors outside a small area. An important quality metric for a deployment and its MLAT accuracy with respect to the target object's (the sending aircraft, in case of ADS-B) relative position is the *geometric dilution of precision*, or GDOP. It describes the effect of deployments on the relationship between the errors of the obtained time measurements and their resulting impact on the errors in the object's calculated position, or formally:

$$\Delta \text{Location Estimate} = \Delta \text{Measurements} \cdot \text{GDOP}$$

GDOP is widely used in positioning systems such as GPS, where good ratings for this multiplier are commonly considered to be below 6, with 10 to be fair and everything over 20 to be of poor quality [33].

2. Theoretically, four or more sensors are sufficient to

⁵As the propagation is not happening in a vacuum, this is an approximation, however, the difference is insignificant [32].

compute a position of an object in 3D space. However, it is very difficult to get the precise altitude of an aircraft when all the receivers are on the ground (i.e., in one plane) and do not provide sufficient elevation angle diversity. In that case, the *vertical* dilution of precision (VDOP) may be too large, so that only horizontal coordinates are calculated for aircraft surveillance and the altitude is obtained by other means [34].

3. While not a security challenge per se, MLAT systems are very expensive. ADS-B needs only one receiver for accurate wide area surveillance; MLAT requires every signal to be received by at least four stations with little noise. Geographical obstacles (e.g., mountain ranges, oceans) make it even more difficult to install a comprehensive wide area system at the desired service level.
4. A determined and resourceful attacker could spoof wireless signals such that using their TDOAs for localization would result in a position of the attacker's choice. This is shown in [35] for the case of GPS. While based on TDOAs, too, GPS is different as only a single receiver is attacked. The authors further discuss the case of spoofing a group of distributed GPS receivers similar to MLAT. They find that a system of multiple receivers severely restricts the attacker placement, each receiver making an attack exceedingly more difficult.

Considering some of these drawbacks and the fact that MLAT is currently the main security solution for unauthenticated ATC networks, we argue that there is an urgent need for other TDOA-based approaches that improve on these problems and provide an immediate practical increase in security.

5. DESIGNING LIGHTWEIGHT LOCATION VERIFICATION OF AIRCRAFT

We propose a solution to verify the location of aircraft based on the physical security properties of TDOA measurements but apply new methods to counter some of the real-world drawbacks of MLAT. By using a mix of deterministic location estimation techniques and statistical approaches, the utility of surveillance data can be vastly increased.

Scalability and coverage

One of the main goals of our design is to tackle MLAT's scalability and coverage problems. An ATC data communications network consists of a given number of sensors that are deployed outside, in a line of sight with the airspace they are expected to cover. Naturally, overlapping reception ranges between receivers are required to obtain TDOAs. If more sensors are to receive the same message, they need to be located closer together. While this increases the overlap, it also decreases the overall ADS-B coverage of the receivers. Worse even, only a small part of the MLAT coverage is usable, since GDOP causes its accuracy to deteriorate quickly. Methods not suffering from GDOP and working with fewer sensors could vastly improve security compared to MLAT.

To demonstrate this fact, we analyzed more than 50 million ADS-B messages from aircraft at cruising altitudes (ca. 38000 ft) with a network comprising 8 receivers (see Table 1). Fig. 6 shows the regions where messages are picked up by a given number of receivers. It also depicts the MLAT-capable area which makes up roughly 5% of the overall covered area.




Figure 6: The map shows the practical reception ranges of a real-world 8 sensor ADS-B system. The turquoise part is the MLAT-capable area, the purple center shows the area with acceptable (i.e., $DOP < 10$) accuracy.

The area where MLAT is reliably accurate is even smaller at around 0.37% of total coverage. When we look at the relative number of messages which can be used for verification purposes, this becomes even clearer. Less than 4% of all messages are seen by 4 or more sensors on the ground and can be used for MLAT. If we take into account dilution of precision, we are left with 0.36% of usable messages. While these numbers concern a natural deployment under real-world constraints, we found that even in simulations with near-optimal coverage (e.g., rectangular or triangular, as discussed in [36]) this does not change significantly.

Of all analyzed legitimate flights for which we received more than 100 messages, 87.7% had at least 10 messages received by 2+ sensors, 65.37% by 3+ sensors and only 9.73% were MLAT-capable. Taking these results into account, we propose other TDOA-based methods to verify aircraft location claims: a grid-based k-NN approach and statistical verification based on expected TDOAs at ground sensors. Both do not suffer from dilution of precision and work with as little as 2 sensors, increasing the effective coverage of our deployment by a factor of >100 , thus vastly reducing costs.

Location verification

For our statistical location verification, we collect TDOAs between at least two sensors that received the message and use them to verify the claimed position of the signal. In other words, we compute the expected TDOA as shown in Fig. 4 and compare it to the measured values.

We use the nonparametric Wilcoxon rank-sum test to continuously check if the received sample matches the expected distribution. By establishing the proximity to the expected data distribution, we can validate the sender. Through collecting more sample messages, we can gain more confidence over time, and prevent outliers without creating false positives. Of course, such a statistical verification approach only shows that it is not *impossible* that the sender is at the claimed location. However, failing this test is a certain indication that the sender is at a different position, at least

	Absolute	Relative	Area covered
All messages	53,551,672	100%	100%
# seen by ≥ 2 sensors	21,437,841	40.03%	45.83%
# seen by ≥ 3 sensors	7,191,209	13.43%	16.56%
# seen by ≥ 4 sensors	2,015,532	3.76%	5.07%
# seen by ≥ 5 sensors	321,719	0.60%	0.79%
# seen by ≥ 6 sensors	16,068	0.0003%	0.0004%
# seen by ≥ 7 sensors	104	$2 \times 10^{-6}\%$	$2.5 \times 10^{-6}\%$
# MLAT & $GDOP < 10$	191,072	0.36%	0.37%

Table 1: The table shows the absolute and relative number of messages collected by a given amount of sensors. The last column provides the relative area covered by the sensors.

outside measurement errors or multipath effects which can be eliminated through repeated application.

Location estimation

Our location estimation method provides a direct and quantitative estimate of a sender's position. These estimates can be used to verify an aircraft's positional claim, where an accuracy of a few hundred meters is typically enough to establish the authenticity of an aircraft for tracking purposes.

Indoor and outdoor localization problems have been studied extensively in the literature, often in the scope of sensor networks and radar applications. Liu et al. [37] give an overview of the techniques used in wireless indoor positioning including the different algorithms (k-Nearest Neighbor, lateration, least squares and Bayesian among others) and primitives such as received signal strength (RSS), TDOA, time of arrival (TOA) and angle of arrival (AoA). While TDOA systems are limited in indoor environments (due to multipath effects and non-availability of time synchronization and clocks fine-grained enough to provide good results at very short distances [38]), they offer very good performance in long-distance outdoor line-of-sight environments such as those encountered in the aircraft location problem.

In terms of algorithms, the k-Nearest Neighbors (k-NN) has proven to do very well in short-distance, indoor RSS fingerprinting compared to other methods [39], although it can become computationally expensive with large databases.

Putting these findings together, we design a novel approach to locate aircraft by creating a 2D grid that contains expected TDOA measurements for each position. For every incoming message, the nearest neighbors of the measured TDOAs are calculated, then the result is compared with the position given by the aircraft. When the estimate deviates too far from the claim, an attack is likely.

6. EXPERIMENTAL SETUP

Data collection and hardware

As ADS-B has been in the roll-out phase for years, we can use real-world data to estimate the propagation characteris-




Figure 7: Location estimation with 3-NN in an adversarial setting where actual and claimed trajectory diverge. Using TDOA data from 4 sensors $S_1, \dots, S_4$ the 3 nearest neighbors N_1, N_2, N_3 found in the lookup table are averaged to obtain the location estimate E . If the deviation between E and the ADS-B claim C exceeds a threshold, an alarm is sent.

tics of ADS-B messages. We do not make any assumptions on hardware features such as sending power or antennas as there are many configurations found in different aircraft.

For our evaluation, we rely on real-world ADS-B data which we obtained from the OpenSky project [40]. OpenSky is a participatory sensor network that collects ADS-B messages in a centralized database. In its current deployment, it receives data from 26 sensors, capturing more than 30% of the commercial air traffic over central Europe. The data is made freely available to researchers. For the present analysis, we use a dataset that spans the period between 26 June 2013 and 25 June 2014. This dataset contains 53,551,672 ADS-B messages received from SBS-3 sensors manufactured by Kinetic Avionics. Besides the message content, they provide a timestamp of the message reception. From this data, we use 5 sensors that are closely located together to be able to calculate their TDOA data. The timestamps have a clock resolution of 50 ns. All sensors have omnidirectional antennas and can receive signals from a distance of up to 400 km.

Synchronization

As our low-cost SBS-3 sensors do not provide built-in synchronization (e.g., via GPS), we synchronize our data a posteriori with the help of positional ADS-B messages sent by aircraft. By using the positional information in those messages and approximating their respective propagation time, we can recover the timing offset between our ground station sensors and achieve global synchronization. We also take into account the drift of the internal clocks to improve the results. Overall, this approach enables us to achieve synchronization that is low-cost and works well with minimal requirements. More accurate and efficient synchronization using GPS could help to further improve on the accuracy of our results. However, the increased security of GNSS-

Algorithm 1 Location estimation offline phase. Requires coordinates of sensors and grid as input and outputs the training sets for the online phase.

```

1: Input: gridcoords, sensors, squaresize
2:
3: trainingset  $\leftarrow []$ 
4: grid  $\leftarrow$  construct_grid(gridcoords, squaresize)
5: for  $\forall$  sensorcombinations do
6:   tdoa_training  $\leftarrow []$ 
7:   for  $\forall$  gridsquare  $\in$  grid do
8:     tdoas  $\leftarrow$  compute_tdoas(sensors.coords, gridsquare)
9:     tdoa_training.add(tdoas, gridsquare)
10:  end for
11: trainingset.add(tdoa_training, sensorcombination)
12: end for

```

free synchronization is another major advantage besides cost savings. It is obvious that in the attacker model with full access to the wireless channel, GPS-spoofing or jamming⁶ are further tools available to the attacker besides the mere injection of ADS-B messages and hence GPS does not necessarily improve the overall security of the system.

Grid design

We construct a 2D grid over a typical flight altitude of 38000 ft (ca. 11,582 m) with a size of 2 degrees longitude and 2 degrees latitude which, due to the Earth's spherical geometry, translates to an area of ca. $150 \text{ km} \cdot 220 \text{ km} = 33,000 \text{ km}^2$. We obtain evenly-spaced approximate squares where the number of squares (or square size) is a trade-off between performance and accuracy as elaborated in the evaluation section. Of course, computation time and accuracy also depend on the size of the surveillance area. 33,000 km are representative for wide area ATC surveillance, covering aircraft's en-route flight phase at cruising altitude.

6.1 Location Verification

Location verification as discussed in this section takes TDOA data as input and outputs whether the data matches pre-determined characteristics of the claimed position.

6.1.1 Offline phase

In the offline phase of the location verification approach, we create a lookup table with fingerprints for every grid position. In detail, we save the deviations between real and expected TDOA between two or more sensors and create a sensor-specific distribution, taking into account all the real-world noise introduced through propagation, synchronization etc. These distributions are leptokurtic with a mean of 0 and a standard deviation of approximately 1 microsecond.

6.1.2 Online phase

In the online phase, we continuously test the likelihood of the measured TDOA of a message. The deviation between the expected TDOA based on its positional claim and the actual TDOAs must conform to the distribution of our collected data for any receiving sensor j . This approach is especially useful when a message has been received by only 2

⁶A practical real-world threat, see e.g. [41].

Algorithm 2 Location estimation online phase. Requires the number of neighbors k and the *trainingsets* from the offline phase as input and calculates the distance between its location estimate and the message’s claim. If *threshold* is exceeded, an alarm is sent.

```

1: Input: threshold,  $k$ , trainingset, flight
2:
3: loop
4:    $m \leftarrow \text{new\_position\_message}(\text{flight})$ 
5:    $r \leftarrow \text{receivers}(m)$ 
6:   if  $\text{number\_of\_receivers}(m) > 2$  then
7:      $\text{tdoas} \leftarrow \text{calculate\_tdoas}(m)$ 
8:      $\text{trainingset} \leftarrow \text{get\_trainingset}(r)$ 
9:      $\text{knn} \leftarrow \text{run\_knn}(\text{trainingset}, \text{tdoas}, k)$ 
10:     $\text{estimate} \leftarrow \text{get\_center}(\text{knn})$ 
11:   end if
12:    $\text{deviation} \leftarrow m.\text{locationclaim} - \text{estimate}$ 
13:   if  $\text{deviation} > \text{threshold}$  then
14:     alarm
15:   end if
16: end loop

```

sensors (i.e. only a single TDOA measurement is available), so an accurate solution is not possible with traditional location estimation methods.

We can gain more confidence over time by collecting more samples and comparing their distribution to the expected one, effectively dealing with outliers without creating false positives. To check if the measurements match the expected distribution, we employ the nonparametric Wilcoxon rank-sum test to test the null hypothesis

H_0 : *The sample comes from the same distribution as our training data.*

against the alternative hypothesis

H_A : *The sample comes from a different distribution than our training data.*

(i.e., they are sent from a source not legitimately at this position) at a 99.99% significance level. The Wilcoxon test is more robust on non-normal distributions as we experience them, compared to other distribution or location tests.

If there is data from more than two receivers available, we increase the robustness of this approach by using a majority voting function to decide whether to classify a flight as legitimate or not. When more than 50% of sensors reject the hypothesis, we classify a flight as illegitimate.

6.2 Location Estimation

Our location estimation also uses an offline training phase while the online phase continuously verifies new aircraft.

6.2.1 Offline phase

Over an exemplary grid of $N \cdot M$ squares, we generate one fingerprint vector of TDOAs between the 5 sensors for every square. We then create a training set for every subset of combinations with at least 2 sensors ($\sum_{i=2}^n \binom{n}{i}$, with n being the number of sensors), i.e., 26 sets overall. This is required when a message is received by fewer than all 5 sensors. In

Attacker Type	Dist. from claim [start/end/avg]
Ground, stationary	74.772 / 90.439 / 78.176 km
Ground, mobile	74.897 / 88.682 / 77.535 km
UAV	74.287 / 87.417 / 77.417 km
Aircraft	0 / 27.778 / 7.191 km

Table 2: Averaged horizontal distances from the four attackers’ positions to their claimed aircraft positions during the time that flight data is injected.

that case the appropriate set is chosen to find the k nearest neighbors. Algorithm 1 details our approach.

6.2.2 Online phase

In the online phase, new message data is analyzed and the location verified (see Algorithm 2 for an overview of the whole process). Using the k-Nearest Neighbors algorithm, we find the closest points from our training grid that match the fingerprints of our test data.

Setting the number of nearest neighbors to k , we match the received fingerprint $R = TDOA_1, \dots, TDOA_n$ to the saved grid fingerprint F based on their Euclidean distance

$$D_{(R,F)} = \sqrt{\sum_{i=1}^n (R_{TDOA_i} - F_{TDOA_i})^2}$$

It is intuitive that in the spatial domain of our grid there are multiple neighbors that are approximately the same distance from our point of interest, hence k is an important parameter influencing the accuracy. If $k > 1$, the positions of all k neighbors are averaged by taking the mean of the longitude and the latitude. This constitutes the estimate of the aircraft position which is closer to the true location than any single neighbor (see Fig. 7 for an illustration).

7. EVALUATION

In this section, we use the collected flight data to verify our approach. Furthermore, we inject data from four different attackers to test the system’s resilience against intruders.

Test data

We use real-world flight data to test our scheme. Taking 10,443 legitimate flights with more than 100 collected messages each, we show that they are accurately verified by our system. Furthermore, we use data from various simulated attackers (due to ethical reasons, we do not implement real-world attacks) on the ground and in the air and check whether they will be verified or not. Table 2 shows the average simulated positions for all four attackers as described in Section 3. Using an omnidirectional antenna, each attacker injects 200 messages with the legitimate coordinates of a real flight from our sample and follows specific location patterns:

- **Attacker 1** has a fixed random horizontal position on the grid with an altitude between 0 and 500 m from which all 200 messages are sent.
- **Attacker 2** is defined by a random start position similar to attacker 1 and a random horizontal direction, moving on the ground with a speed of 50km/h.

# sensors	2				3				4				5			
# messages	1	10	30	100	1	10	30	100	1	10	30	100	1	10	30	100
Legit flight	0	<0.1	0	0	0	<0.1	0	0	0	<0.1	0	0	0	<0.1	0	0
Attacker 1	0	93.8	91.2	93.8	0	99.9	99.7	99.9	0	99.5	99.2	99.9	0	100	99.9	100
Attacker 2	0	98.6	95.9	94.0	0	99.8	99.5	99.9	0	99.9	99.8	99.9	0	99.9	100	100
Attacker 3	0	98.8	96.3	94.5	0	99.8	99.6	99.9	0	100	99.9	99.7	0	100	99.9	100
Attacker 4	0	74.4	80.6	89.5	0	74.4	80.88	94.1	0	70.56	79.36	90.3	0	79.1	92.2	95.6

Table 3: Results of the location verification approach dependent on number of received messages and number of sensors. The values signify the percentage of flights that have been classified as attackers.

- **Attacker 3** is defined by a random start position, a random altitude between 0 and 1100m and a random horizontal direction, moving with a speed of 200km/h.
- **Attacker 4**’s starting position is the same as the real aircraft but diverts horizontally at a random angle between 10 and 45 degrees (at cruising altitude), making attacker 4 the most difficult to detect.

The attacker’s TDOAs are calculated by dividing the 3D distance between the sensors by the speed of light c and adding some white Gaussian noise analogous to our real data to account for measurement and processing errors. We test each scenario 1000 times and analyze the detection rate.

7.1 Location Verification

Table 3 shows the results of testing our location verification method. As we can see, it is able to detect all attackers successfully, while minimizing false positives. For all legitimate flights, the null hypothesis is accepted when at least 30 samples are collected. For attackers 1-3, which are all relatively far away from their claimed distances (i.e., on the ground or in low airspace), H_0 is generally rejected after collecting 10 or more message samples. False negatives stay in the low single digits even with TDOAs gathered by only two sensors. For the most powerful attacker 4, who is acting very similar to the injected ADS-B claims, a sample size of 50 is needed to detect most of the injected flights. In a non-lossy environment, we can collect 50 messages in under 10s. Assuming 50% message loss, we are alerted within 20s after the aircraft has diverted from its claimed course.

7.2 Location Estimation

We first compare our location estimation method with the GPS-based ADS-B position claims of legitimate flight data, to ensure its accuracy. We use a data set of over 100,000 positional ADS-B messages from a two-week sample where every message has been seen by 5 sensors, providing us with the necessary TDOA measurements. All location claims are on the grid in terms of latitude and longitude, while the mean altitude is 11,148.8m ($\sigma = 687.59m$). Table 5 shows the location estimation quality using k-NN with squares of five different sizes over an area of 33,000 km² with $k = 5$ (see Fig. 8) for the optimal choice of k .

As expected, increasing the number of squares has a positive impact; the smaller the square, the more accurate location predictions become. For example, a reduction in grid square size from 600 m² to 300 m² improves mean accuracy by 37.5%. This naturally comes with a trade-off as the computational time to run the k-NN algorithm increases linearly by 400%. Overall, we found that 150 m² provides a good trade-off between accuracy and performance.

Error [m]	MLAT	2 sens.	3 sens.	4 sens.	5 sens.
Mean	199.5	26,956.7	311.8	147.3	122.3
Median	91.9	22,737.1	145.4	95.8	84.9
RMSE	334.5	33,380.4	761.3	237.6	190.3
99%ile	1306.7	63,500.2	2,469.6	983.7	870.6

Table 4: Average horizontal errors using k-NN ($k = 5$) with 150 m square size and different amounts of receivers. MLAT (5 sensors) is provided as comparison.

We also compared k-NN with a linearized MLAT algorithm using the same TDOA measurements from 5 sensors. The results show that with a 600 m² grid size, k-NN does 14.2% better than MLAT on mean errors, increasing to 41% for a 50 m² grid size. Overall, we find that k-NN does better than MLAT on noisy TDOA measurements such as those we experienced in our real-world data. Especially the more outlier-sensitive metrics RMSE and mean improve with k-NN while MLAT generally shows good median results. Since k-NN does not suffer from dilution of precision, this is to be expected as the mean GDOP in our dataset is 24.35 ($\sigma = 8.06$). Taking only “good” values below 10 into account, MLAT are bound to metrics improve vastly. However, doing this also decreases the number of usable messages by over 90%, reinforcing the fact that k-NN is useful in a much larger area. Of course, there is no reason why all considered TDOA approaches cannot easily complement each other.

The computational time is the trade-off for k-NN’s accuracy. Only with the largest square size of 600 m² it is comparable to MLAT. However, depending on the density of the airspace and the available equipment, even larger grids and longer computation times would not pose a problem in real-world settings.⁷ In scenarios where location estimation is run mainly to verify suspicious aircraft claims, it is entirely irrelevant as the examined amount of data is very small.

For a comprehensive security approach, it is furthermore important to compare the impact of sensor numbers on location estimation. Table 4 shows the results for the same dataset and a 150 m² grid size, if only a subset of the five sensors receives the messages. After analyzing all possible subsets and averaging the results, we find that with only three sensors sufficient horizontal accuracy can be achieved.

Attacker Detection

We analyze the results of our attacker models who inject false ADS-B data from a different location. From our ex-

⁷The complexity of the MLAT algorithm is constant, while k-NN depends on the number of squares, i.e., both the size of the monitored area and the desired accuracy.




Figure 8: Optimal choice of neighbors k for different square sizes (MLAT as comparison). We can see a large improvement until $k = 5$, further decreases in mean accuracy are small and less pronounced with smaller square sizes.

perimental analysis of the legitimate data, we find that the system should flag a given flight as illegitimate when the average deviation between ADS-B claim and k-NN estimate exceeds 1,000m over 12 messages received by 3 sensors. With this setting we encounter zero false positives in our test data, yet detect all false-data injections by attackers 1-3 within 12 messages as their location far exceeds the threshold. Attacker 4, who starts from the correct position, is detected in fewer than 38 messages on average, i.e. after about 20 seconds without loss or 40 seconds assuming 50% loss. MLAT is not able to function and detect the attackers with 3 sensors. Naturally, the precise thresholds depend on equipment and scenario and should be fitted accordingly.

Besides detection of false claims, location estimation can provide a guess of the attacker’s current location. Table 6 provides the results for all four attacker types. Our horizontal estimate for the origin of message signals fits within approximately 2,000 m for the ground-/low airspace-based attackers and for the aircraft attacker within the typical error range shown for legitimate flights of less than 200 m.

Accordingly, Table 6 also shows a major drawback of MLAT in the same scenario. While it is feasible (though costly) to build a system with good accuracy for larger areas in the sky, it is difficult to provide the same level of accuracy on the ground and within areas that are not expected to be used for commercial traffic but could be relevant in adversarial settings. Hence, MLAT offers a similar estimate quality for attacker 4 in our setup but is not able to provide the location of ground/low airspace attackers. Whereas their injections are also detected by MLAT, the estimates are too inaccurate to provide any information on the location of attackers 1-3.

7.3 Discussion: Sensitivity vs. Practicality

There is an important trade-off between false positives and false negatives which has to be considered when choosing thresholds. A sensor in a busy airspace can see thousands of flights per day, too sensitive settings may lead to a number of false alarms and cause users to disregard or deactivate the system. This is especially relevant considering the time-sensitivity and the general high-stress environment found in ATC. In evaluating our system, we have chosen thresholds that did not cause any false positives.

Yet, considering the potentially disastrous outcomes of a real attack, it seems likely that the user will have to accept at least the very rare occasional false alarm for increased safety. On top of this, there are cases where an IDS would rightly report a suspicious communication pattern by a le-

Estimate	Dist. to claim [km]		Dist. to attacker [km]	
	k-NN	MLAT	k-NN	MLAT
Attack 1	78.174	120.440	2.056	47.505
Attack 2	78.408	118.325	1.918	44.947
Attack 3	78.217	117.498	2.021	44.255
Attack 4	7.228	7.227	0.145	0.270

Table 6: Left: Mean distances between estimates and claimed location injected by an attacker. Right: Mean distances to actual horizontal location of an attacker. k-NN ($k = 5$) with 150m square size. k-NN accurately detects the distances between the attacker and the claim and gives a good guess about the real origin of the signal. MLAT also detects the deviations can only provide an accurate position of the aircraft-based attacker.

gitimate flight, for example when a transponder malfunction has occurred in an aircraft. The fine-tuning of the threshold in practice depends on the following factors, among others:

- The number of flights registered by the sensors per day.
- Availability of backup systems such as radar or MLAT.
- The quality of the collected data (e.g., number of sensors and channel quality).
- The desired time frame (i.e., number of collected samples) after which a decision by the IDS should be made.

8. CONCLUSION AND FUTURE WORK

In this work, we present a novel method to secure NextGen ATC surveillance systems. We develop a realistic threat model and show that existing and proposed countermeasures are insufficient to deal with these threats. Furthermore, we propose and evaluate two methods of location verification. The first one, statistical and based on collected time differences of arrival between as little as two ADS-B sensors, allows us to quickly detect injected data with high certainty. Using only low-cost ADS-B sensors, we find that it outperforms MLAT in terms of range and detection speed, increasing coverage by a factor of more than 100.

The second approach requires at least three sensors to not only detect false-data injection attackers even faster and more reliably than MLAT but also estimate their position. We evaluate our scheme with real-world flight data from a large-scale sensor network and test it against injected flights by simulated attackers. The results show that the mean aircraft location accuracy can be increased by up to 41% in comparison with MLAT and that ground-based attackers can be located with a mean horizontal error of 2,000 m.

TDOA-based security solutions remain the de facto standard as they are readily available without hardware or software changes. It is important to stay ahead in the security arms race often found in real-world systems until the aviation community works out fundamental long-term solutions for authentication in ATC communication networks. In the (foreseeably long) meantime, it is crucial to increase security to protect air traffic against potentially devastating events.

In future work, we plan to integrate other indicators such as heading and bearing of an aircraft or the received signal strength of ADS-B messages into the system to further improve on accuracy and detection metrics.

Horizontal Error [m]	MLAT	600 m ² Grid	300 m ² Grid	150 m ² Grid	75 m ² Grid	50 m ² Grid
Mean	199.46	171.01	134.37	122.31	118.14	116.454
Median	91.87	140.38	98.60	84.92	80.38	78.63
RMSE	334.47	225.51	198.14	190.29	187.31	185.79
99th percentile	1306.70	902.08	870.18	870.61	841.33	835.63
Relative comp. time	62.3%	100%	399%	1599%	7272%	16375%

Table 5: Horizontal errors in different grid square sizes using k-NN vs. MLAT, with 5 sensors and $k = 5$. k-NN shows a better mean accuracy than MLAT of up to 41% in our data set.

9. REFERENCES

- [1] L. Vidal, “ADS-B Out and In - Airbus Status,” ADS-B Taskforce - KOLKATA, Apr. 2013.
- [2] A. A. Cardenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry, “Attacks against process control systems: risk assessment, detection, and response,” in *Proceedings of the 6th ACM symposium on information, computer and communications security*. ACM, 2011, pp. 355–366.
- [3] D.-Y. Yu, A. Ranganathan, T. Locher, S. Capkun, and D. Basin, “Detection of GPS spoofing attacks in power grids,” in *Proc. of the 2014 ACM conference on Security and privacy in wireless & mobile networks*. ACM, 2014.
- [4] A. Costin and A. Francillon, “Ghost in the Air (Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices,” in *Black Hat USA*, 2012.
- [5] R. Kunkel, “Air traffic control insecurity 2.0,” in *DefCon 18*, 2010.
- [6] M. Schäfer, V. Lenders, and I. Martinovic, “Experimental analysis of attacks on next generation air traffic communication,” in *Applied Cryptography and Network Security*, ser. LNCS, no. 7954. Springer, Jun. 2013.
- [7] ICAO, “Cyber Security for Civil Aviation,” in *Twelfth Air Navigation Conference*, 2012, pp. 1–4.
- [8] M. Clayton. (2014, Mar.) Malaysia airlines flight mh370: Are planes vulnerable to cyber-attack? Christian Science Monitor.
- [9] M. Strohmeier, V. Lenders, and I. Martinovic, “Security of ADS-B: State of the Art and Beyond,” *arXiv preprint arXiv:1307.3664*, Jul. 2013.
- [10] D. McCallie, J. Butts, and R. Mills, “Security analysis of the ADS-B implementation in the next generation air transportation system,” *International Journal of Critical Infrastructure Protection*, vol. 4, no. 2, pp. 78–87, Aug. 2011.
- [11] L. Purton, H. Abbass, and S. Alam, “Identification of ADS-B System Vulnerabilities and Threats,” in *Australian Transport Research Forum*, Canberra, 2010.
- [12] A. Proano and L. Lazos, “Selective jamming attacks in wireless networks,” in *Communications (ICC), 2010 IEEE International Conference on*. IEEE, 2010.
- [13] M. Strohmeier, M. Schäfer, V. Lenders, and I. Martinovic, “Realities and Challenges of NextGen Air Traffic Management: The Case of ADS-B,” *Communications Magazine, IEEE*, vol. 52, no. 5, May 2014.
- [14] ICAO, “Guidance Material: Security issues associated with ADS-B,” Tech. Rep., 2014.
- [15] A. Smith, R. Cassell, T. Breen, R. Hulstrom, and C. Evers, “Methods to Provide System-wide ADS-B Back-Up, Validation and Security,” in *25th Digital Avionics Systems Conf.*, 2006.
- [16] ICAO, “Guidance Material on Comparison of Surveillance Technologies (GMST),” Tech. Rep. September, 2007.
- [17] M. Gariel and E. Feron, “Graceful degradation of air traffic operations: airspace sensitivity to degraded surveillance systems,” *Proceedings of the IEEE*, vol. 96, no. 12, 2008.
- [18] K. D. Wesson, T. E. Humphreys, and B. L. Evans, “Can cryptography secure next generation air traffic surveillance?” *IEEE Security and Privacy Magazine*, 2014.
- [19] A. Greenberg. (2012, Jul.) Next-gen air traffic control vulnerable to hackers spoofing planes out of thin air. Forbes.
- [20] K. Zetter. (2012, Jul.) Air traffic controllers pick the wrong week to quit using radar. Wired.
- [21] B. Haines, “Hacker + airplanes = no good can come of this,” in *Confidence X*, 2012.
- [22] B. Kovell, B. Mellish, T. Newman, and O. Kajopaiye, “Comparative Analysis of ADS-B Verification Tech.” 2012.
- [23] 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)*. IEEE, 2011.
- [24] B. Nuseibeh, C. B. Haley, and C. Foster, “Securing the Skies: In Requirements We Trust,” *Computer*, vol. 42, no. 9, 2009.
- [25] I. A. Mantilla-Gaviria, M. Leonardi, G. Galati, and J. V. Balbastre-Tejedor, “Localization algorithms for multilateration (MLAT) systems in airport surface surveillance,” *Signal, Im. and Video Processing*, 2014.
- [26] W. W. Li and P. Kamal, “Integrated Aviation Security for Defense-in-Depth of Next Generation Air Transportation System,” in *IEEE Conf. on Tech. for Homeland Sec.*, 2011.
- [27] M. Schäfer, V. Lenders, and I. Martinovic, “Experimental Analysis of Attacks on Next Generation Air Traffic Communication,” in *Applied Cryptography and Network Security*. Springer, 2013, pp. 253–271.
- [28] C. Finke, J. Butts, R. Mills, and M. Grimaila, “Enhancing the security of aircraft surveillance in the next generation air traffic control system,” *International Journal of Critical Infrastructure Protection*, vol. 6, no. 1, pp. 3–11, Mar. 2013.
- [29] K. Sampigethaya and L. Bushnell, “A Framework for Securing Future e-Enabled Aircraft Navigation and Surveillance,” in *AIAA Proceedings*, 2009, pp. 1–10.
- [30] E. Chan-Tin, V. Heorhiadi, N. Hopper, and Y. Kim, “The Frog-Boiling Attack: Limitations of Secure Network Coordinate Systems,” *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 3, p. 27, 2011.
- [31] RTCA Inc., “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 Corrig. 1, 2011.
- [32] N. J. Gomes, P. P. Monteiro, and A. Gameiro, *Next generation wireless communications using radio over fiber*. Wiley, 2012.
- [33] M. Mosavi and H. Azami, “Applying Neural Network Ensembles for Clustering of GPS Satellites,” *International Journal of Geoinformatics*, vol. 7, no. 3, 2011.
- [34] G. Galati, M. Leonardi, P. Magarò, and V. Paciucci, “Wide area surveillance using SSR mode S multilateration: advantages and limitations,” in *European Radar Conference (EURAD)*, 2005.
- [35] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Capkun, “On the requirements for successful GPS spoofing attacks,” in *Proceedings of the 18th ACM conference on Computer and Communications Security*. ACM, 2011.
- [36] W. Y. Poe, “Design problems in large-scale, time-sensitive wsns,” Ph.D. dissertation, TU Kaiserslautern, Germany, 2013.
- [37] H. Liu, H. Darabi, P. Banerjee, and J. Liu, “Survey of wireless indoor positioning techniques and systems,” *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, vol. 37, no. 6, 2007.
- [38] P. Bahl and V. N. Padmanabhan, “RADAR: An in-building RF-based user location and tracking system,” in *INFOCOM 2000. 19th Annual Joint Conf. of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2, 2000.
- [39] A. Rozyyev, H. Hasbullah, and F. Subhan, “Combined K-Nearest Neighbors and Fuzzy Logic Indoor Localization Technique for Wireless Sensor Network,” *Research Journal of Inform. Tech.*, vol. 4, no. 4, 2012.
- [40] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, “Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research,” in *ACM/IEEE International Conf. on Information Processing in Sensor Networks*, 2014.
- [41] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, “GPS vulnerability to spoofing threats and a review of antispooing techniques,” *International Journal of Navigation and Observation*, 2012.