Crowd-GPS-Sec: Leveraging Crowdsourcing to Detect and Localize GPS Spoofing Attacks

Kai Jansen*, Matthias Schäfer†, Daniel Moser‡, Vincent Lenders§, Christina Pöpper¶ and Jens Schmitt†

*Ruhr-University Bochum, Germany, kai.jansen-u16@rub.de
†University of Kaiserslautern, Germany, {schaefer, jschmitt}@cs.uni-kl.de
‡ETH Zurich, Switzerland, daniel.moser@inf.ethz.ch
§armasuisse, Switzerland, vincent.lenders@armasuisse.ch
¶New York University Abu Dhabi, United Arab Emirates, christina.poepper@nyu.edu

***Abstract*—The aviation industry's increasing reliance on GPS to facilitate navigation and air traffic monitoring opens new attack vectors with the purpose of hijacking UAVs or interfering with air safety. We propose *Crowd-GPS-Sec* to detect and localize GPS spoofing attacks on moving airborne targets such as UAVs or commercial airliners. Unlike previous attempts to secure GPS, *Crowd-GPS-Sec* neither requires any updates of the GPS infrastructure nor of the airborne GPS receivers, which are both unlikely to happen in the near future. In contrast, *Crowd-GPS-Sec* leverages crowdsourcing to monitor the air traffic from GPS-derived position advertisements that aircraft periodically broadcast for air traffic control purposes. Spoofing attacks are detected and localized by an independent infrastructure on the ground which continuously analyzes the contents and the times of arrival of these advertisements. We evaluate our system with real-world data from a crowdsourced air traffic monitoring sensor network and by simulations. We show that *Crowd-GPS-Sec* is able to globally detect GPS spoofing attacks in less than two seconds and to localize the attacker up to an accuracy of 150 meters after 15 minutes of monitoring time.**

I. Introduction

Today, more than a billion devices rely on the Global Positioning System (GPS) for various applications that require accurate positioning or precise time synchronization. With its ubiquitous coverage, GPS has become the *de facto* standard means of navigation and tracking services in outdoor environments, where it achieves an accuracy of up to three meters [1]. For navigation purposes, satellite systems such as GPS are mission-critical for Unmanned Aerial Vehicles (UAVs), ranging from consumer-class mini or micro drones to tactical and strategic UAVs.

Although GPS is commonly used in aviation, the system is not secure, i. e., civilian (public) GPS signals sent by the satellites are neither authenticated nor encrypted. As a consequence, aircraft and UAVs are vulnerable to GPS signal spoofing attacks, where a malicious transmitter emits signals similar to those from the satellites but at a higher power and, potentially, at slightly different time delays. The aircraft's GPS receiver will lock on to the spoofed signal as it arrives with a higher signal strength than the authentic signals.

By selectively varying the time offsets of the spoofed satellite signals, attackers are able to mimic arbitrary positions. These kinds of spoofing attacks are well-known [2]–[7] and have been shown to be feasible in the real-world [5], [8]. In fact, GPS spoofing has allegedly been used to hijack a CIA stealth drone (RQ-170) in Iran in 2011 [9] or luring ships off their course [4], [10]. Moreover, GPS spoofing has been used as a defense against GPS-controlled UAVs flying in the vicinity of the Kremlin in Russia [11].

Over the years, the price to perform GPS spoofing attacks has dramatically dropped. Mobile commercial off-the-shelf GPS spoofing devices are available for less than $1,000 [4] and publicly available software tools [12] allow the generation of arbitrary GPS signals. The price fall and low-expertise requirements raise the risk for applications relying on GPS for safety- or security-critical decisions and processes.

The democratization of GPS spoofing technologies has triggered the development of various countermeasures, which can be coarsely categorized into three classes: (i) cryptographic techniques, (ii) detection at signal level, and (iii) direction of arrival sensing. Cryptographic techniques [13]–[16] aim at authenticating signals from satellites with additional signals that are unpredictable to users that do not own a secret key. However, these techniques are not resistant to replay attacks and would require a costly upgrade of the GPS infrastructure. Spoofing detection at signal level are based either on anomaly checks in the physical signal waveform [17]–[19] or on measuring the angle of arrival from which the signal is originating [20], [21]. While these techniques do not require a change in the structure of GPS signals, they impose modifications on existing receivers and increase the complexity and computational requirements of those devices. We conclude that existing countermeasures are unlikely to be implemented in the near future since they all require far-reaching modifications of either the GPS infrastructure or the receiving devices.

Driven by the increasing threat and the lack of realistic short-term solutions, we propose *Crowd-GPS-Sec*, a system that detects and localizes GPS spoofing attacks on aerial vehicles without the need to update the structure of the GPS satellites' signals nor the logic of the airborne GPS receivers. *Crowd-GPS-Sec* leverages crowdsourcing to monitor the position advertisements derived from GPS that aircraft and UAVs periodically broadcast for air traffic surveillance. Using those advertisements, we devise a GPS spoofing *detection* and *localization* solution that analyzes the contents and the time of arrival of these surveillance messages as received by different sensors on the ground.

We have evaluated *Crowd-GPS-Sec* with simulations and real-world data from the OpenSky Network [22], a crowd-sourcing initiative which maintains a network of more than 700 air traffic communication sensors around the world. Our implementation of *Crowd-GPS-Sec* is able to globally detect GPS spoofing attacks in less than two seconds and to localize the attacker up to an accuracy of 150 meters after 15 minutes of monitoring time.

While the problem addressed in this work is related to spoofing detection and localization in classical direction finding [20], [21] and multilateration systems [23], there is one fundamental difference and unique advantage. Instead of trying to detect and localize the GPS spoofer through direct measurements of its own signals, we rely on indirect measurements from the position advertisements that the aircraft are broadcasting. This approach enables us to detect and localize the spoofer even when there is no direct line-of-sight connection from a sensor to the spoofer. Maintaining a line-of-sight connection to the *aircraft* is much simpler and thus more effective since the aircraft are in the sky and use high transmission power levels which render the signals receivable from the ground up to several hundred kilometers away. Another major advantage is that *Crowd-GPS-Sec* relies on data from air traffic monitoring sensors that are already widely deployed around the world. Thus, the solution does not require a dedicated GPS signal acquisition infrastructure for spoofing detection and localization. To the best of our knowledge, this paper is the first to propose a GPS spoofing countermeasure which takes advantage of considering indirect GPS-inferred data rather than raw GPS signals.

In summary, this paper makes the following contributions:

- We propose *Crowd-GPS-Sec* and elaborate on the idea to provide security via an existing infrastructure of crowd-sourcing sensors.
- We present algorithms for the *detection* of GPS spoofing attacks on airborne targets by using aircraft reports and multilateration.
- We provide a novel technique for the *localization* of GPS spoofers based on position differences between pairs of spoofed aircraft.
- We report on experiments with aircraft transponders and assess the performance of *Crowd-GPS-Sec* analyzing real-world air traffic control data.

II. The Global Positioning System

The GPS infrastructure is a satellite-based navigation network of over 30 satellites located in the medium Earth orbit, more than $20{,}000\,\mathrm{km}$ above the Earth's surface. GPS-capable receivers can determine their position and time by measuring the time of arrival (ToA) from at least four satellites. Based on the ToA and the transmission time embedded in the signals, receivers can calculate distances to each satellite. Multilateration of those distances yields the position and the local time of the receiver.

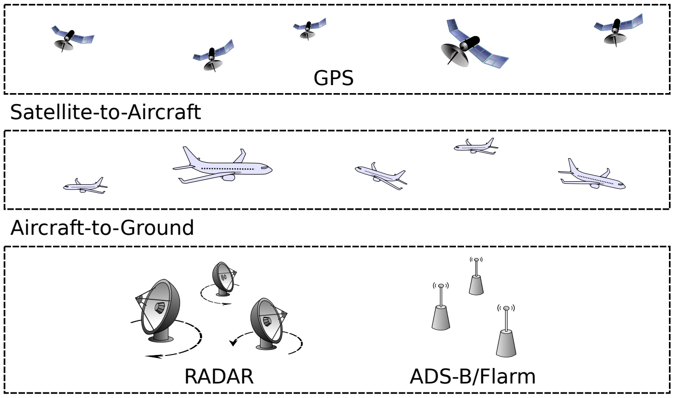

Fig. 1. Schematic overview of currently deployed technologies used to monitor air traffic including GPS, RADAR, and ADS-B/Flarm.

The ToA measurements are affected by a range of errors resulting in a typical localization uncertainty of $\sigma = 4\,\mathrm{m}$ (mean error of about $7\,\mathrm{m}$) [24]–[27]. While civilian (public) GPS signals can be decoded by everyone, including airplanes, drones, and other UAVs, military GPS signals are protected by (at least) secret spreading codes restricting their users to a selected group with additional knowledge. We focus on civilian GPS with non-authenticated signals, which is the standard in commercial and general aviation.

A. GPS Usage in Aviation

While in the past, radar and inertial systems used to be the two main localization technologies in aviation, GPS is today often the preferred solution due to its superior accuracy. Modern airliners, smaller aircraft, gliders, helicopters, or UAVs are almost all equipped with GPS receivers. GPS is typically used by pilots or UAVs for self-localization but the technology is also used for remote air-traffic surveillance and collision-avoidance applications. In the latter cases, aerial vehicles are required to periodically broadcast position and velocity advertisements to inform neighboring aircraft and ground controllers about their presence. Larger aerial vehicles generally transmit those messages over the Automatic Dependent Surveillance – Broadcast (ADS-B) system while smaller and slower vehicles rely on the Flarm [28] system. Irrespective of the used system, these advertisements contain a position that is directly derived from airborne GPS receivers as depicted in Figure 1.

In this work, we propose to leverage the position advertisement messages of ADS-B and Flarm in order to detect and localize GPS spoofers. While ADS-B and Flarm rely on different radio frequencies and message formats, the underlying concept is the same. On regular random intervals (around twice per second), aircraft broadcast their current position together with their unique addresses. Neighboring aerial vehicles and ground stations receive these messages to generate a recognized air picture. The advertisement messages can be received over long distances. In ADS-B, messages can be received up to distances of $700\,\mathrm{km}$ when there is a direct line-of-sight connection between the transmitter and the receiver. In Flarm, the range is smaller but reception ranges of up to $100\,\mathrm{km}$ are possible.

B. GPS Spoofing Attacks

GPS spoofing attacks exploit the lack of encryption and authentication of civilian GPS signals by imitating the legitimate signals with the purpose of modifying the localization or time result of a victim [3], [7], [25]. Technically, spoofing attacks are based on fake GPS signals manipulating the ToAs of signals that otherwise use the same payload as real signals.

In the past, incidents were reported [4], [9]–[11] where spoofers successfully interfered with the integrity of GPS-dependent systems, thus rendering the spoofing threat far from being only of theoretical nature. As a result, currently marketed drones, aircraft, helicopters, or any kind of vehicles that rely on GPS are prone to spoofing attacks and lack effective countermeasures.

Based on common assumptions on attacker capabilities and recent incidents, we assess the resulting threat model in this section. First, we clarify our considered adversary model. Second, we reason about key assumptions that *Crowd-GPS-Sec* is based on to detect and localize spoofing attacks.

1) Threat Model: The attacker's motivation to interfere with the air safety by injecting false positioning information into UAVs or aircraft can be manifold. An attacker may consider hijacking the targeted victim for an own benefit of acquiring goods or circumventing flying bans. Even more severe, an attacker may participate in terrorist attacks by manipulating the air-traffic control or the collision-avoidance systems, e. g., by spoofing fake position information to fool the safety logic of these systems.

In our adversary model, the attacker is able to transmit specially crafted signals identical to those broadcasted by GPS satellites but can achieve a higher power *at the target location*. The attacker aims at spoofing a moving aircraft or a UAV from a position on the ground. In order to conduct a stealthy and unnoticed attack, the spoofer may use a directional antenna[1] directed towards the victim in the sky. However, due to the target's movement, the attacker needs to transmit signals from a considerable distance, hundreds of meters to kilometers away. We note that typical operating altitudes of UAVs range from $60\,\mathrm{m}$ to $20{,}000\,\mathrm{m}$ and their mission radii vary from $5\,\mathrm{km}$ to $200\,\mathrm{km}$ and beyond [30]. Hence, if the route taken by the victim is not predictable, the attacker will be forced to use antennas with wide-beam propagation patterns. This forces the attacker to transmit signals of such a strength and propagation that the spoofing signals most likely will not only be received at a particular primary target location but also over a wider area, affecting other aircraft and UAVs in the neighborhood. Since the spoofer is targeting moving vehicles, we further assume that the spoofer is emulating a moving track such as a straight line or a curve with some potential acceleration.

(a) PowerFLARM Core (b) PowerFLARM Portable

Fig. 2. Two newest-generation Flarm transponder models. Both transponders have an integrated GPS receiver but do not provide any protection to GPS spoofing and advertise false positions when spoofed.

2) Validation of Assumptions: Crowd-GPS-Sec relies on two key assumptions which we validate in this section. The first assumption is that whenever a GPS receiver locks on to the spoofed signals, the position advertisements of the aircraft and UAVs will contain the spoofed GPS positions. While commercial GPS receivers are known to be vulnerable to spoofing attacks [2]–[5], [8], [10], [31]–[33], aviation transponders could have additional plausibility checks to prevent that spoofed GPS positions propagate to the broadcasted position advertisements. The second assumption is that the spoofed signals will not only affect the target victim of the spoofer but also neighboring aircraft and UAVs. We validate these two assumptions with controlled lab experiments and simulations with real-world air traffic data from the OpenSky Network.

GPS Spoofing Experiments. We perform GPS spoofing experiments with two Flarm transponders that are widely deployed. As we could not get formal approval from our national office of communications to perform GPS spoofing experiments in the wild with real aircraft, we rely on an isolated experimental setup inside a shielded lab environment. The goal of these experiments is to demonstrate that existing transponders do not perform any checks on the derived GPS position and that spoofers can precisely control the position and speed of victim receivers.

Our experimental setup consists of two new-generation Flarm transponder models from Flarm Technology: a *PowerFLARM Core* and a *PowerFLARM Portable* both with an integrated GPS receiver from u-blox, see Figure 2. More than $30{,}000$ manned aircraft, helicopters, and UAVs over the world are equipped today with these transponders [28]. As GPS spoofer, we rely on a USRP B200 from Ettus Research and the software-defined GPS signal simulator *gps-sdr-sim* [12]. To monitor the reported Flarm position advertisements by the transponders, we use a Raspberry Pi with an RTL-SDR software-defined radio dongle and the *flare* open-source Flarm decoder [34]. All devices are equipped with omnidirectional antennas.

[1] We focus on the common assumption that the attacker uses a single antenna for transmitting the spoofing signals, but the proposed technique could also be extended to multi-antenna attackers representing an emerging threat [29].

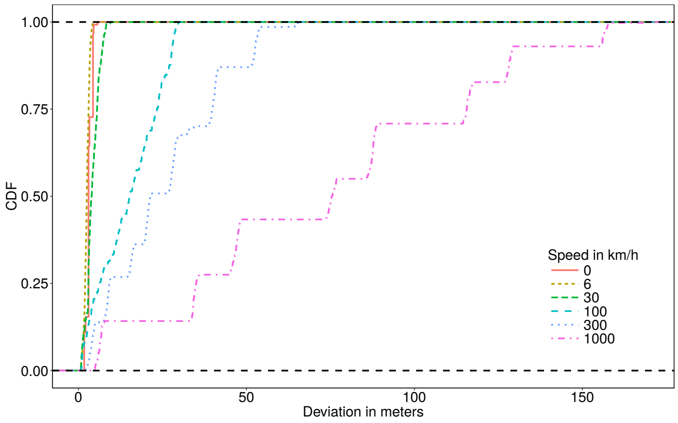

Fig. 3. Cumulative distribution function (CDF) of deviation between spoofed and reported position messages of the *PowerFLARM Core* transponder.

We put all devices in vicinity of each other and spoof tracks with speeds of 0, 6, 30, 100, 300, and $1{,}000\,\mathrm{km/h}$, respectively. The difference between the fake target positions emitted by the spoofer and the reported positions in the Flarm advertisements is plotted in Figure 3. While the deviation becomes larger with increasing speed, our experiments confirm that an attacker can exactly control the derived position and speed at the Flarm devices. Even for speeds up to $1{,}000\,\mathrm{km/h}$, the deviation of both spoofed devices is always smaller than $160\,\mathrm{m}$, and thus significantly smaller than the mandated separation minima in aviation [35]. These experiments also confirm that such commercial transponders as deployed in aerial vehicles do not perform plausibility checks on the GPS signal input and simply report the spoofed GPS data in the advertisement messages. This result is inline with air traffic communications not being protected against wireless attacks [36].

GPS Spoofing Coverage Estimation. To validate the assumption that a GPS spoofer will affect the GPS receivers of multiple aerial vehicles at the same time, we evaluate the reception range of a spoofer using the free-space path loss model and a typical airspace density model as observed by the OpenSky Network in the European airspace.

Since the power of GPS signals at the Earth's surface is very low (approx. $-160\,\mathrm{dBW}$), the necessary power to create adequate spoofing signals is accordingly low. We assume an attacker with standard equipment, who can reasonably achieve a generated signal power of $15\,\mathrm{dBm}$ (USRP2 [37]) coupled with an exemplary antenna gain of $12\,\mathrm{dBi}$ in the main lobe. We also consider an additional signal attenuation at aircraft of approx. $30\,\mathrm{dB}$ due to the fuselage and the downward direction. Based on these estimations, we can calculate the reception range with regard to the free-space path loss [38]:

$$L_{\mathrm{fs}} = 32.45 + 20\,log_{10}(d_{\mathrm{km}}) + 20\,log_{10}(f_{\mathrm{MHz}}), \qquad (1)$$

where d_{km} is the distance between the source of the signal and the receiver in kilometers and f_{MHz} is the signal frequency given in megahertz; the constant of 32.45 depends on the utilized units. The resulting reception range is based on the signal power impaired by all attenuation sources and the distance d from Equation (1):

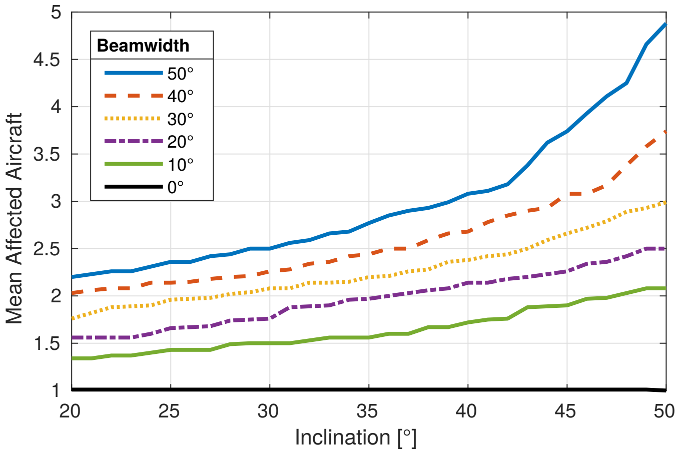

Fig. 4. The number of affected aircraft depends on the directional antenna beamwidth and the inclination angle. The figure uses a realistic airspace density sampled from OpenSky Network data.

$$\mathrm{Power} - L_{\mathrm{fs}}(d) - \mathrm{Attenuation} \geq -160\ [\mathrm{dBW}],$$

which results in a distance d of approx. $34\,\mathrm{km}$. Considering our parameter estimations, all aircraft within the main lobe closer than $34\,\mathrm{km}$ will receive the spoofing signal with at least $-160\,\mathrm{dBW}$.

In general, an attacker will be interested to exceed these power levels to ensure the takeover of the GPS lock at the intended targets. However, to remain as stealthy as possible, the attacker is likely to use an attack setup with directional antennas to avoid a wide signal broadcast detectable by, e. g., ground-based signal power sensors. A directional antenna setup is characterized by its beamwidth influencing the signal spread and the inclination angle determining how the main lobe of the signal beam is targeted. Notably, an attack on moving targets requires to increase the beamwidth and to use higher inclination angles, resulting in a certain proliferation of the affected area.

Based on data from the OpenSky Network of the European airspace, we perform a conservative estimate of the average number of aircraft affected by a spoofing attack targeting a randomly selected aircraft, as shown in Figure 4. The baseline (0° beamwidth) is an attacker that can perfectly pinpoint a victim, thus avoiding secondary targets. Such a small beamwidth is however impossible to achieve in practice and would further be very sensitive to small orientation errors of the antenna. As we can see, already small beamwidths and inclination angles span enough space to affect several aircraft around the intended target, making it highly likely to hit several additional aircraft. The assumption that our work relies on is therefore realistic for dense airspaces such as found in Europe.

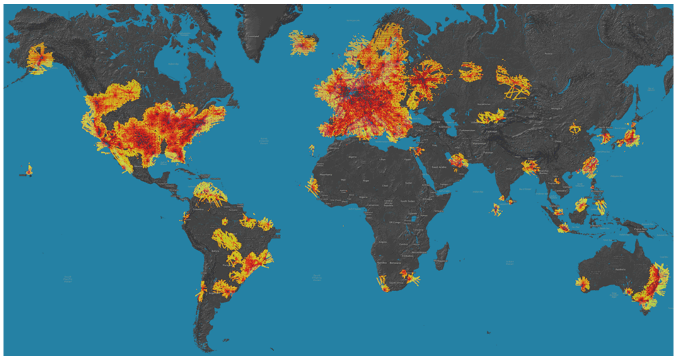

Fig. 5. Worldwide coverage of *Crowd-GPS-Sec* as of December 2017.

III. Crowd-GPS-Sec

We propose *Crowd-GPS-Sec* as an independent system infrastructure on the ground that continuously analyzes the contents and the time of arrival of Flarm and ADS-B position advertisements. As its name suggests, *Crowd-GPS-Sec* relies on crowdsourcing to monitor those messages at global scale. The sensors used for *Crowd-GPS-Sec* are part of the growing OpenSky Network [39], a crowdsourcing initiative with the purpose to make air traffic communication data available to the public.

The vast majority of the sensors are installed and operated by aviation enthusiasts and volunteers which support the cause of the network. As of this writing, it collects more than 200,000 messages per second at peak times from over 700 sensors which are distributed all over the world[2] as shown in Figure 5. Europe and the American continent exhibit a particular high density of sensors such that individual position advertisement messages are most likely being received by more than four sensors.

The goals of *Crowd-GPS-Sec* are to detect GPS spoofing attacks on aerial vehicles as quickly as possible and to localize the position of the spoofer(s). To achieve these goals, *Crowd-GPS-Sec* has three modules which continuously process all position advertisements that are received from the OpenSky Network, as shown in Figure 6. The *multilateration (MLAT)* module estimates the location of the aircraft based on the time difference of arrival (TDoA) of position advertisements between different sensors. This module is fundamental to *Crowd-GPS-Sec* as it allows us to determine the true position of the aircraft independently of the content of the advertised messages. The *spoofing detection* module checks for inconsistencies between multilaterated positions and GPS-derived positions in the advertisement messages as well as for inconsistencies between position advertisements from different aircraft (e. g., when two aircraft advertise the same position at the same time). The *spoofer localization* module, finally, is triggered only when the spoofing detection module has detected a GPS spoofer. It then estimates the position of the spoofer by analyzing time differences between received positions in advertisements from the aircraft and the true position as estimated by MLAT. We describe the modules in the next three subsections.

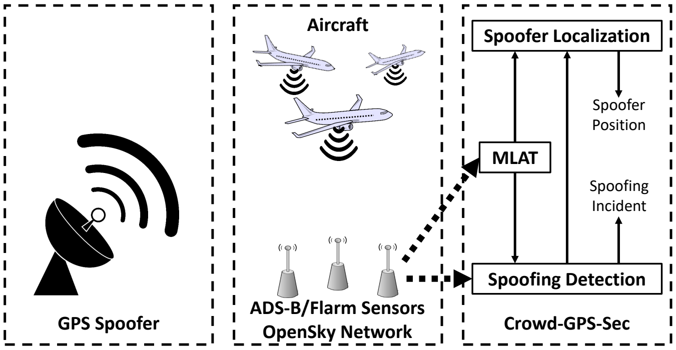

Fig. 6. *Crowd-GPS-Sec* system overview. A GPS spoofer transmits fake GPS signals that are received by multiple aircraft periodically broadcasting ADS-B/Flarm position reports. Ground-based sensors record these reports, which are then processed by *Crowd-GPS-Sec* for spoofing detection and spoofer localization.

A. Multilateration (MLAT)

The implementation of *MLAT* as an independent aircraft localization will serve as an auxiliary component for one of the spoofing detection tests and the subsequent spoofer localization. To implement such a system, we make use of the fact that in regions with high sensor density position advertisement messages are received by multiple geographically distributed sensors. Each message is timestamped at the receiver on arrival and can be represented as a simplified tuple of the reported position and the time of arrival:

$$\text{ADS-B/Flarm Report} := (\hat{a}_i, t_s), \tag{2}$$

where $\hat{a}_i$ denotes the reported position of aircraft i as derived by GPS and t_s is the timestamp as generated by sensor s.

Since the sensors are geographically distributed, propagation distances of the transmitted signals differ. Hence, the same broadcasted message is timestamped differently at diverse sensors. If the sensors are synchronized to the same global clock, e. g., by GPS time synchronization, and are deployed at known positions, we can formulate relations between the propagation distances and the differences in the time of arrival (TDoA):

$$dist(s_i, \mathrm{A}) - dist(s_j, \mathrm{A}) = \Delta t_{i,j} \cdot c, \tag{3}$$

where s_i, s_j denotes the position of sensor i and sensor j. The TDoA of the same message from reference aircraft A between these sensors is $\Delta t_{i,j} = t_i - t_j$, and c is the speed of light.

Equation (3) is fulfilled for all points that have the same distance difference to both considered sensors determined by the TDoA. By construction of at least four relations of this type, we perform multilateration to approximate the position of the targeted aircraft. Geometrically, each relation describes a hyperbola in 2D and a hyperboloid in 3D. The intersecting point of all relations indicates the aircraft position. Figure 7 provides a visual interpretation of this multilateration process.

[2] See https://opensky-network.org/network/facts for more statistics.

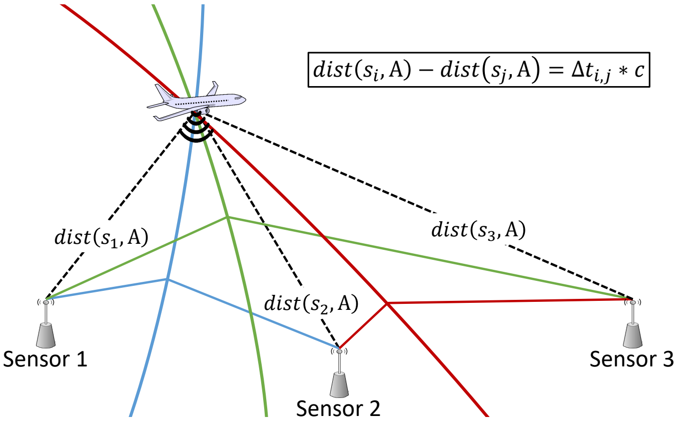

Fig. 7. Implementation of an independent aircraft localization scheme based on multilateration considering the TDoA of broadcasted ADS-B/Flarm messages.

B. GPS Spoofing Detection

Spoofing detection is the first step in a mitigation strategy to counter GPS spoofing attacks. The idea of *Crowd-GPS-Sec* to detect GPS spoofing attacks is based on the broadcasted ADS-B/Flarm reports containing potentially spoofed positioning information. We propose a verification process consisting of two complementary checks.

1) Time Alignment of Transmissions: Since ADS-B/Flarm messages are broadcasted at variable transmission times, we need to time-align those reports in order to make them comparable. This is achieved by incorporating the results from the MLAT computation. To align the position reports to a reference global time, two steps are performed subsequently.

The first step yields the transmission time t_{TX} at which the GPS-derived position was reported:

$$t_{\mathrm{TX}} = t_s - \frac{dist(s,a)}{c}, \tag{4}$$

with t_s being the time at which sensor s has timestamped the message, $dist(s, a)$ representing the Euclidean distance between the considered sensor and aircraft, and c being the speed of light.

The second step is an interpolation to approximate the aircraft position a_{REF} at a global reference time t_{REF}. We need to consider the following three cases:

$$a_{\mathrm{REF}} = \begin{cases} \frac{a_{\mathrm{TX}} \cdot (t_{\mathrm{TX}+1} - t_{\mathrm{REF}}) + a_{\mathrm{TX}+1} \cdot (t_{\mathrm{REF}} - t_{\mathrm{TX}})}{t_{\mathrm{TX}+1} - t_{\mathrm{TX}}} & t_{\mathrm{TX}} < t_{\mathrm{REF}} \\ a_{\mathrm{TX}} & t_{\mathrm{TX}} = t_{\mathrm{REF}} \\ \frac{a_{\mathrm{TX}} \cdot (t_{\mathrm{REF}} - t_{\mathrm{TX}-1}) + a_{\mathrm{TX}-1} \cdot (t_{\mathrm{TX}} - t_{\mathrm{REF}})}{t_{\mathrm{TX}} - t_{\mathrm{TX}-1}} & t_{\mathrm{TX}} > t_{\mathrm{REF}} \end{cases}$$

with $a_{\mathrm{TX}} = \hat{a}$ denoting the aircraft position at transmission time, TX−1, TX, and TX+1 being the previous, current, and next transmission event, respectively. After this interpolation, all reported positions are time-aligned and can be compared with respect to the same time basis. In the remainder of this paper, we assume time-aligned positions.

2) Test 1 (Cross-Checks with MLAT): We propose the implementation of two complementary tests. The first test performs a cross-check between the reported positions and the estimated real positions from the previously described MLAT approach. We check for each incoming position report whether

$$dist(a_i, \hat{a}_i) \overset{?}{<} \mathcal{T}_1 \tag{5}$$

holds, where a_i is the real position of aircraft i determined by MLAT, $\hat{a}_i$ is the position reported by aircraft i using ADS-B/Flarm, $dist()$ is the Euclidean distance function, and $\mathcal{T}_1$ denotes a predefined threshold which tolerates measurement errors in a_i and $\hat{a}_i$. Choosing the right threshold $\mathcal{T}_1$ depends on the accuracy of the underlying secondary localization method (here MLAT). Smaller $\mathcal{T}_1$ lead to higher false positive rates, while larger $\mathcal{T}_1$ create more room for undetected manipulations.

Complexity. Let n be the number of aircraft. Equation (5) needs to be checked once for each aircraft, i. e., n times, resulting in a complexity of $\mathcal{O}(n)$. For each sampling time, we require the positioning information from ADS-B/Flarm and MLAT. The comparisons of both positioning sources can be parallelized, since the checks for each aircraft are independent of all other aircraft. As a result, the first test of GPS spoofing detection scales linearly with the number of simultaneously tracked aircraft.

3) Test 2 (Multiple Aircraft Comparison): The second test makes use of the information provided by other aircraft. In particular, we perform a comparison between reported positions of multiple aircraft. When multiple aircraft receive the signals from the same spoofer device, they will appear at the same location [7] since the time differences between individual satellites are emulated on the radio of the spoofer prior transmission. Due to mandatory separation minima [35], i. e., minimum required distances between en-route aircraft, similar positions are critical and are caused either by a serious incident, e. g., near-collision, or a GPS spoofing attack. Eventually, the multiple aircraft comparison test is defined as:

$$dist(\hat{a}_i, \hat{a}_j) = d_{i,j} \overset{?}{>} \mathcal{T}_2, \tag{6}$$

where i and j denote two different aircraft, $\hat{a}_i$ and $\hat{a}_j$ are the GPS-derived positions of aircraft i and aircraft j, $dist()$ is the Euclidean distance function, and $\mathcal{T}_2$ is a threshold tolerating the GPS positioning noise. Choosing an appropriate $\mathcal{T}_2$ depends on the mandated separation minima in the considered airspace and the accuracy of the GPS information provided via position reports. However, as accuracy is one of the design goals of ADS-B and Flarm and the separation minima are usually in the order of kilometers, a threshold as small as a few hundreds of meters is appropriate.

Complexity. Let n be the number of aircraft. Since Equation (6) considers pairs of aircraft, a naive implementation would require $\binom{n}{2} = \frac{n^2-n}{2}$ comparisons resulting in a complexity of $\mathcal{O}(n^2)$. However, since Test 2 considers spatial data only, the complexity can be reduced by implementing nearest neighbor searches based on k-d trees and cover trees.

TABLE I
SPOOFING DETECTION TESTS COMPARISON

Feature	Test 1	Test 2
Equation	$dist(a_i, \hat{a}_i) \overset{?}{<} \mathcal{T}_1$	$dist(\hat{a}_i, \hat{a}_j) \overset{?}{>} \mathcal{T}_2$
Complexity	$\mathcal{O}(n)$	$\mathcal{O}(n \cdot \log n)$
Requirement	MLAT positioning	Multiple aircraft
Advantages	Single spoofed aircraft detection	Independent of MLAT Separation of attacks

In fact, since Test 2 fails if there is any neighbor closer than $\mathcal{T}_2$, solving the 1-nearest-neighbor problem for each aircraft is sufficient. Using the aforementioned data structures, this can be accomplished at a complexity of $\mathcal{O}(\log n)$ for each aircraft [40], resulting in a global complexity of $\mathcal{O}(n \cdot \log n)$.

4) Complementary Design: We propose a complementary design consisting of both tests in parallel. Table I contains a comparison of the spoofing detection tests. While the first test based on the cross-check of Equation (5) is independent of other flights, the second test based on the comparison of multiple aircraft of Equation (6) is independent of the MLAT positioning and can thus tolerate bad MLAT performance (e. g., when sensors have a bad geometric distribution leading to high dilution of precision). Furthermore, the second test is able to separate multiple spoofing attacks occurring at the same time as there will be independent sets of coinciding aircraft. The combination of both tests can overcome the pitfalls of the other and we can achieve a more versatile and robust spoofing detection.

C. GPS Spoofer Localization

After spoofing detection, *Crowd-GPS-Sec* aims at localizing spoofer devices. This is the next step in tracing an attacker in order to take appropriate action for shutting down an attack. We present a novel localization approach to remotely pinpoint such devices using already available ADS-B/Flarm reports broadcasted by aircraft. We start by describing the high-level idea and then detail on the functionality of the crowdsourced localization system.

1) Localization Model: If a malicious device emits GPS spoofing signals, aircraft within the effective range will broadcast spoofed positions as contained in their ADS-B/Flarm reports. All aircraft that receive the same fake GPS signals will report positions on the same track but timely shifted as a result of the propagation delay from different distances to the spoofing source [7]. In particular, at the same global time, the aircraft have different synchronizations on the spoofing signals based on how long it takes the signals to arrive at the aircraft's GPS receiver, i. e., aircraft that receive the fake signals earlier are ahead on the spoofed track, whereas aircraft that are further away from the spoofer receive the signals at a later point in time and are thus behind on the track. We extract the resulting position differences from the ADS-B/Flarm reports and backtrace these deviations to the location of the spoofing device.

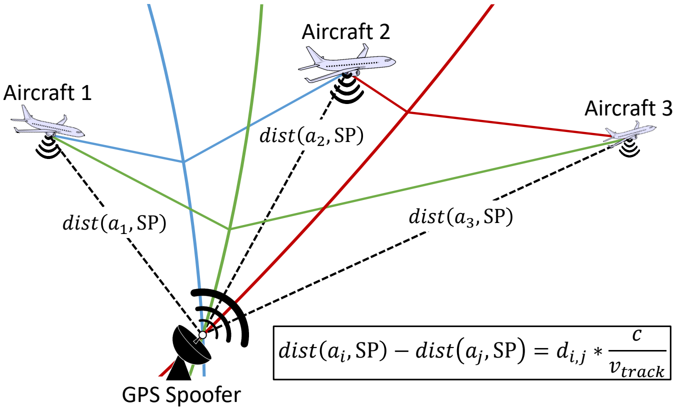

Fig. 8. Each relation forms a hyperboloid representing all points with the same distance differences. For the shown 2D projection, we can construct three distinct relations considering three different aircraft.

Our starting point is the identification of the currently spoofed aircraft, which is the outcome of the GPS spoofing detection module. For those identified aircraft, we forward related information to the spoofer localization module. We further require the actual aircraft positions a_i, a_j from MLAT and the mutual distances $d_{i,j}$ with $i, j \in \{\text{spoofed aircraft}\}$.

As next step, we put the aircraft distance into relation with the propagation distances and the rate of position change, i. e., the spoofed track velocity. We can formulate this as follows:

$$dist(a_i, \mathrm{SP}) - dist(a_j, \mathrm{SP}) = d_{i,j} \cdot \frac{c}{v_{track}}, \qquad (7)$$

where a_i, a_j indicate the actual position of aircraft i, j as given by MLAT, SP is the unknown spoofer location, $d_{i,j}$ the respective aircraft distance, and v_{track} the velocity of the spoofed GPS track. The factor $\frac{c}{v_{track}}$ relates the position change rate to the signal propagation speed (close to the speed of light). We note that we need to assure $v_{track} \neq 0$ and hence require a track of changing positions. Having related the reported positions to the spoofer location, we solve each equation towards this location. In particular, each equation describes all points that have the same mutual distance differences.

Geometric Interpretation. Considering the solutions of one relation of the type given by Equation (7), all potential solutions geometrically describe a hyperbola in 2D and a hyperboloid in 3D with foci a_i, a_j and distance difference $d_{i,j} \cdot \frac{c}{v_{track}}$. With two different relations, the possible solutions describe a curve, which is the intersection between the hyperboloids. Eventually, three hyperboloids intersect in at most two points, whereas four or more hyperboloids narrow down the location of the spoofer to a single point. The general functionality of this approach is depicted in Figure 8 (2D projection).

Requirements. In order to get at least four different relations, we need to fulfill one of the cases shown in Table III. In particular, we either require four or more different reference aircraft or, in the case we have less, we need to gather reports from the same reference aircraft but from different locations on their tracks. In other words, position reports sent by only two

TABLE II
MLAT VS. SPOOFER LOCALIZATION

Approach	Scenario	Equation	Reference	Target	Measure	Scaling Factor
MLAT	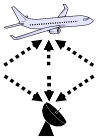	$dist(s_i, \mathrm{A}) - dist(s_j, \mathrm{A}) = \Delta t_{i,j} \cdot c$	Sensors	Aircraft	Time	c
Spoofer Localization		$dist(a_i, \mathrm{SP}) - dist(a_j, \mathrm{SP}) = d_{i,j} \cdot \frac{c}{v_{track}}$	Aircraft	Spoofer	Position	$\frac{c}{v_{track}}$

TABLE III
LOCALIZATION REQUIREMENTS

Affected Aircraft	Possibility of Localization
1	Localization not possible
2	At least 4 different locations
3	At least 2 different locations
4+	Localization possible

aircraft but from four different locations are already sufficient to perform spoofer localization. Since we consider moving targets, the transmission origins will also change likewise. Hence, we are able to trade the number of spoofed aircraft with the required observation time, which we can formulate as follows:

$$\binom{m}{2} \cdot t_s \geq 4, \tag{8}$$

where m is the number of spoofed aircraft and t_s denotes the number of observed samples from different aircraft positions. The binomial coefficient provides the number of possible relations. Equation (8) defines the minimum requirements for our spoofer localization. If fulfilled, we can construct at least four equations and eventually determine a distinct solution.

Comparison with MLAT. The described localization approach exhibits similarities to the MLAT process of Section III-A but is characterized by decisive differences as compared in Table II. Our approach uses the position information included in the ADS-B/Flarm reports, whereas MLAT is based on differences in the time of arrivals at multiple sensors. We want to highlight that it is not possible to trace the location of spoofing devices with MLAT. In our approach, we thus exploit a characteristic that is attacker-controlled such as the spoofed positions in the advertisements. As a result, we obtain a multilateration with switched roles, i. e., the references are moving aircraft as compared to the stationary ADS-B/Flarm sensors. Since the considered measure is shifted from time to positioning information, we need to adjust the scaling factor with the velocity of the spoofed track. As a beneficial side effect, this diminishes the factor with which the uncertainties in the GPS-derived positions are multiplied and consequently minimizes the noise impact on the localization accuracy.

2) Error Minimization: In contrast to a definite analytic solution considering relations based on Equation (7), real-world signal reception and measurements suffer from several error sources and hence prevent a distinct solution for the spoofer position. Both the positions from MLAT as well as the reported spoofed GPS positions are affected by noise. Notably, the interpolation process for time-alignment induces even more noise into the system. Consequently, compared to the theoretical analysis, the constructed hyperboloids do not intersect in a distinct point but rather mark an area.

In order to find the optimal solution for the spoofer position SP, we formulate the following error function $\mathrm{E}_t(\cdot)$:

$$\mathrm{E}_t(\mathrm{SP}, i, j) = dist(a_i, \mathrm{SP}) - dist(a_j, \mathrm{SP}) - d_{i,j} \cdot \frac{c}{v_{track}}, \tag{9}$$

where $d_{i,j}$ is the distance in the reported ADS-B/Flarm positions and t is the current sample time. The real aircraft positions are denoted by a_i, a_j and c is the speed of light.

All resulting errors add up to the overall error, which we try to minimize by computing the root mean square error (RMSE). Eventually, our algorithm outputs the most likely spoofer position:

$$\underset{\mathrm{SP}}{\arg\min} \sqrt{\frac{\sum_{t=1}^{\infty} \sum_{i=1}^{m} \sum_{j=1}^{i-1} \mathrm{E_t}(\mathrm{SP}, i, j)^2}{t \cdot \left(\frac{m^2}{2} - m\right)}}, \tag{10}$$

with t indicating the sample time corresponding to Equation (9). The inner two sums aggregate the errors of relations between all spoofed aircraft, whereas the outer sum aggregates the errors over all sample times. The argument with the minimum error is calculated to be the best approximation for the spoofer position.

When time progresses, the total number of relations considering different references increases. This also affects the error minimization process by expanding the system of equations that are simultaneously evaluated. However, the complexity increase is only linear and, as we will show, this process stabilizes quickly. As all measurements are affected by noise, more relations are beneficial to reduce the system-intrinsic errors and the localization is predicted to gain precision.

3) Improved Filtering: For GPS spoofing targeting multiple aircraft, we identify an additional optimization technique that helps to lower the impact of uncertainty in the reported positions even further. As all affected aircraft receive the same spoofing signals, they report positions on the *same* track irrelevant of timing information. This allows to better predict the underlying track by incorporating *all* available reports. Consequently, we can apply a subsequent filtering of the spoofed aircraft positions.

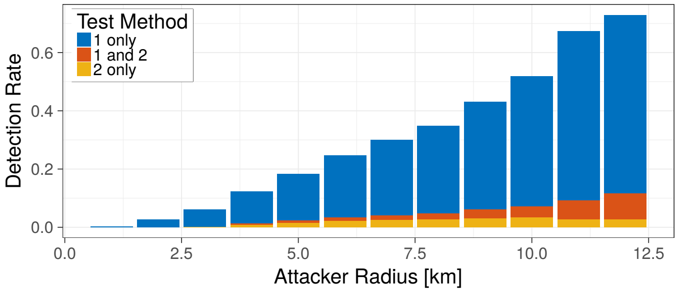

Fig. 9. Detection rates and coverage of Test 1 and Test 2 in the considered OpenSky Network data set depending on the attacker's range.

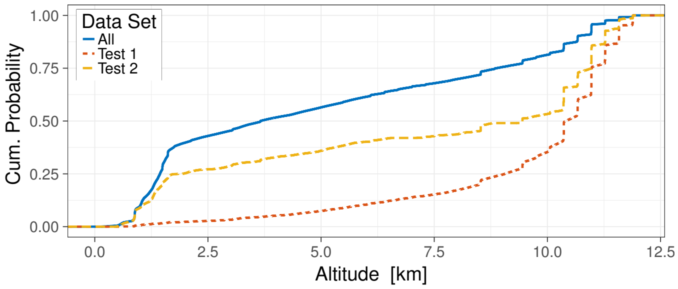

Fig. 10. Comparison of the detection rates of Test 1 and Test 2 in the OpenSky Network data set depending on the target's altitude.

In particular, we apply a projection of the reported positions on the combined estimated track. Notably, with this projection we cannot correct timing inaccuracies, but we can better estimate the most likely position at the current measurement time. The (orthogonal) projection provides the least error with respect to the estimated track and can be described as:

$$\hat{a}_i - \hat{a}_i{}' \perp track, \tag{11}$$

where $\hat{a}_i$ is the noisy GPS position and $\hat{a}_i{}'$ is the projected point with $\hat{a}_i - \hat{a}_i{}'$ being orthogonal on the estimated track. Moreover, we do not necessarily require a continuous straight line but the track can also contain separated segments, which are then evaluated separately to apply the projection.

IV. Evaluation

To evaluate the applicability of *Crowd-GPS-Sec* to real-world air traffic, we assess its performance in terms of spoofing detection and accuracy of the spoofer localization. In particular, we have implemented *Crowd-GPS-Sec* and applied it to real-world data from the OpenSky Network. Moreover, we have built a simulation framework to generate results with respect to spoofing scenarios.

A. Spoofing Detection Performance

We compare our two spoofing detection tests with regard to their coverage, detection delay, and detection rate. The tests are applied to air traffic data of Central Europe as received by the OpenSky Network over a period of $1\,\mathrm{h}$. The data set contains 141,693 unique positions of 142 aircraft.

Coverage. We define the coverage of a test as the percentage of aircraft positions that is protected by a test. Protection means that a test indicates a spoofing attack if the aircraft is indeed spoofed. For simplicity, we assume that the attacker is using an omnidirectional antenna and is positioned right underneath the target using exactly the required transmission power to have the target aircraft lock on the spoofer. This results in an attack range in the form of a sphere with a radius of the altitude of the aircraft. Note that this setup models an unrealistically optimal attacker since in reality, the attacker may not be able to stay exactly underneath the target aircraft as the aircraft is moving and it may use higher transmission powers than the minimal required power.

Since both tests rely on different features, the sets of positions covered by one test is different from the one covered by the other test, but there are overlaps. We therefore analyze how many aircraft in our data set are covered by which test. Figure 9 shows the fractions of aircraft in the data set covered by Test 1, Test 2, or both depending on the target's altitude. The results show that Test 1 clearly outperforms Test 2. Overall, $61.2\,\%$ of the aircraft are covered only by Test 1 while $2.9\,\%$ are covered only by Test 2. In addition, $8.9\,\%$ are covered by both tests. This result is not surprising since the receiver density of the OpenSky Network is high (which benefits Test 1), while the aircraft density (which Test 2 relies on) is limited due to separation minima. In total, we can summarize that if the spoofer's target is at an altitude above $11\,\mathrm{km}$ and the spoofer is directly underneath the target, the detection rate is about $75\,\%$ using both tests. If the spoofer uses higher transmission powers or if it is not directly underneath the target, the detection rate increases quickly towards $100\,\%$ (not shown in the Figure).

As mentioned above, Test 1 directly depends on multilateration coverage and should therefore work better at high altitudes where aircraft are tracked by more sensors. In contrast, Test 2 benefits from dense airspaces since close aircraft "protect" one another. To further investigate this effect, we considered the cumulative distribution of the altitudes of all aircraft and compared it to those of the aircraft protected by either of the tests. The results are shown in Figure 10. As expected, Test 2 has a distribution similar to all altitudes. The steep inclines in its distribution confirm that it is most effective at the common altitudes above $10\,\mathrm{km}$ (en route flights) and at around $1\,\mathrm{km}$ (approach areas). Most aircraft detected by Test 1, on the other hand, were higher than $10\,\mathrm{km}$ which also complies with the above hypothesis.

Detection Delay. We define the detection delay as the delay between the point in time when the attack takes effect, i. e., when the aircraft's GPS sensor locks on to the spoofed signal until the detection test will detect the attack. As for Test 1, this corresponds to the delay between receiving the ADS-B position and the MLAT position update. To evaluate this, we used the open-source MLAT implementation [41] with the OpenSky Network's real-time data stream and measured the time between the reception of an ADS-B position and the

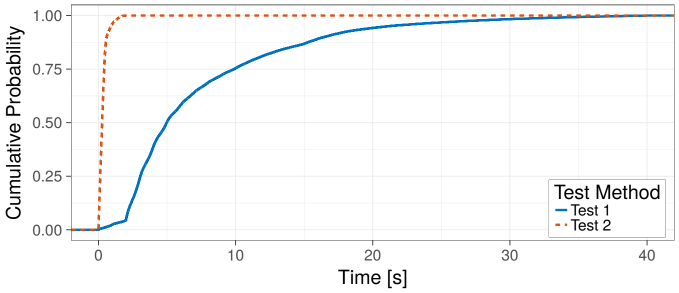

Fig. 11. Comparison of the detection times of Test 1 and Test 2 in the OpenSky Network data set.

TABLE IV
SIMULATION FRAMEWORK PARAMETERS

Parameter	Parameter Range	Default
Sensor Density	$10 \dots 100 \left[\frac{1}{(100\,\mathrm{km})^2}\right]$	OpenSky
Airspace Density	$10 \dots 100 \left[\frac{1}{(100\,\mathrm{km})^2}\right]$	OpenSky
Flightpath	random	OpenSky
Flight Altitude	0 ...10,000 [m]	OpenSky
Airspeed	0 ...1,000 [km/h]	OpenSky
Spoofer Position	random	random
Spoofing Range	10 ...200 [km]	100 km
Spoofed Track Velocity	0 ...10,000 [km/h]	1,000 km/h
GPS Noise (std)	0.01 ...4 [m]	4 m
MLAT Noise (std)	1 ...100 [m]	10 m

emission of the respective position by the MLAT implementation. As for Test 2, the delay can be reduced to the inter-arrival times between spoofed position reports. Figure 11 shows the distributions for the delays of the two tests. The delay of Test 1 is a result of the delay of the relatively long MLAT calculations. Test 2, on the other hand, can detect an attack as soon as a false position report is received from two different aircraft. Note that the position broadcast interval of ADS-B is random within an interval of $0.4\,\mathrm{s}$ to $0.6\,\mathrm{s}$, explaining the average detection delay close to $0.5\,\mathrm{s}$.

Conclusion. The results of our evaluation show that with realistic air traffic and implementation characteristics, the two tests can reach a detection rate of up to $75\,\%$ when the attacker is directly underneath the target. While Test 1 performs much better in terms of coverage and detection rate, the detection delay is much smaller for Test 2. These results encourage a complementary implementation as proposed in Section III-B4.

B. Spoofer Localization Performance

To evaluate *Crowd-GPS-Sec* in terms of GPS spoofer localization accuracy, we have built a simulation framework in MATLAB, which allows us to analyze spoofing scenarios in a controlled environment without having to spoof real aircraft. In particular, we assess the impact of noise in the GPS-derived position reports, MLAT positioning noise, and spoofed track velocity.

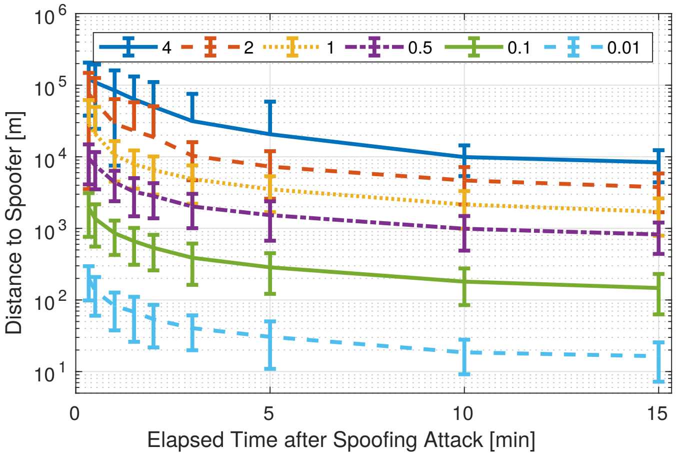

Fig. 12. The impact of GPS noise models ranging from $\sigma_{\mathrm{GPS}} = 4\,\mathrm{m}$ to $0.01\,\mathrm{m}$ on the spoofer localization, depicted including standard deviation errorbars. The MLAT positioning accuracy is fixed to $\sigma_{\mathrm{MLAT}} = 10\,\mathrm{m}$.

Simulation Framework. While we are interested in results from varying parameter sets, we otherwise incorporate realistic data observed by the sensor infrastructure of the OpenSky Network. Table IV contains an overview of the utilized simulation parameters. In the default case, our simulation samples aircraft from the OpenSky Network including reported positions, altitudes, airspeeds, and headings. The spoofer is randomly positioned in an exemplary area of $(400\,\mathrm{km})^2$ and its range is set to $100\,\mathrm{km}$ spoofing a track of $1{,}000\,\mathrm{km/h}$. On the other hand, we are able to simulate different airspace constellations, attacker configurations, and noise impacts of MLAT and GPS. In particular, we consider standard assumptions taken from specifications [1] and technical reports [42] as well as more optimistic assumptions that could be achieved with more sophisticated equipment.

To simulate the impact of GPS spoofing on aircraft, we imitate position reports from already spoofed aircraft by incorporating the attacker-controlled position and adding Gaussian noise according to the considered noise model. Subsequently, we apply standard noise correction techniques based on a Kalman filter [43]. For the error minimization considering distance relations, we implement a numerical solver. To cope with an increasing number of equations, we only evaluate the relations at discrete time intervals which are defined as the time that has elapsed since the spoofing attack was launched, ranging from a few seconds up to 15 minutes.

Metrics. In order to quantify our results we define two metrics. First, we consider the distance between the actual spoofer position and our estimation. Second, we construct a circle around our estimated position with a radius equal to the distance to the actual spoofer. We consider this to be the search space to find the attacker and we compare it to the observed area of $(400\,\mathrm{km})^2$, on which the spoofer was randomly positioned. For each of the analyzed parameter sets, we performed 200 randomized simulation runs and averaged the results.

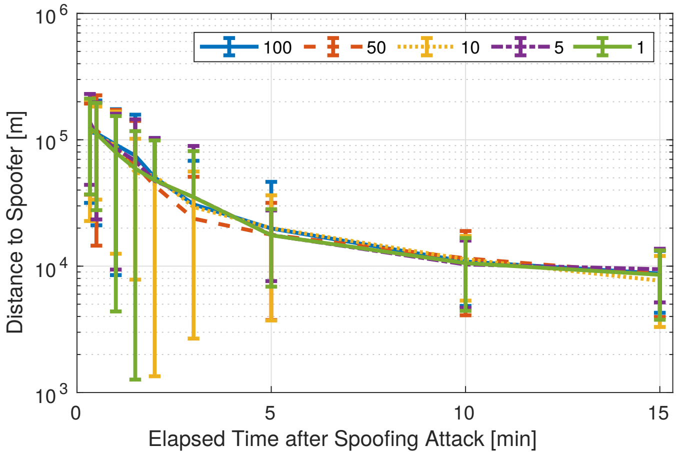

Fig. 13. The considered MLAT positioning noise models in the range of $\sigma_{\text{MLAT}} = 100\,\text{m}$ to $1\,\text{m}$ do not show any significant impact on the localization accuracy. The results are based on a high GPS noise of $\sigma_{\text{GPS}} = 4\,\text{m}$.

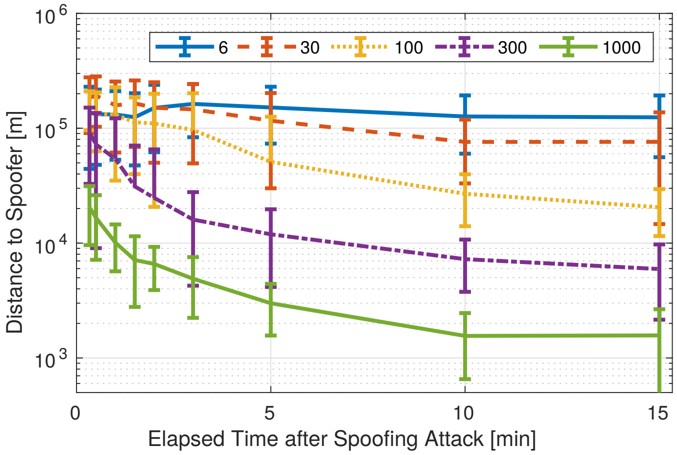

Fig. 14. The velocity of the spoofed track is analyzed for speeds between $v_{track} = 6\,\text{km/h}$ to $1{,}000\,\text{km/h}$. The results consider a GPS noise level of $\sigma_{\text{GPS}} = 1\,\text{m}$ and an MLAT positioning accuracy error of $\sigma_{\text{MLAT}} = 10\,\text{m}$.

1) Impact of GPS Accuracy: Figure 12 depicts the impact of high GPS noise ($\sigma = 4\,\text{m}$) to low GPS noise ($\sigma = 0.01\,\text{m}$) applied to the latitude and longitude direction. We do not require altitude information for spoofer localization and can therefore neglect altitude inaccuracies. We conclude that the extent of noise in the reported GPS positions is a dominating factor that can make the difference between a few kilometers and merely tens of meters in spoofer localization. In particular, we achieve an average localization accuracy of approx. $8.2\,\text{km}$ for $\sigma_{\text{GPS}} = 4\,\text{m}$, approx. $1.7\,\text{km}$ for $\sigma_{\text{GPS}} = 1\,\text{m}$, and approx. $149\,\text{m}$ for $\sigma_{\text{GPS}} = 0.1\,\text{m}$, each after 15 minutes. Considering the search space reduction, we need to scan approx. $0.13\,\%$ for $\sigma_{\text{GPS}} = 4\,\text{m}$, approx. 5.8×10^{-5} for $\sigma_{\text{GPS}} = 1\,\text{m}$, and approx. 4.4×10^{-7} for $\sigma_{\text{GPS}} = 0.1\,\text{m}$, again after 15 minutes. Furthermore, we can observe that the localization accuracy increases rapidly within the first few minutes, whereas after $5\,\text{min}$ the accuracy only improves slowly. From $5\,\text{min}$ to $15\,\text{min}$, the distance roughly halves. As a result, we can already give a good spoofer position estimation in a timely manner after the spoofing attack is launched and narrow it down to a more exact position after a few minutes.

2) Impact of MLAT Accuracy: Another uncertainty of our localization approach is the accuracy of the MLAT positioning that we require to determine the actual (unspoofed) aircraft positions. We choose to vary the MLAT accuracy between high noise ($\sigma_{\text{MLAT}} = 100\,\text{m}$) and lower noise levels ($\sigma_{\text{MLAT}} = 1\,\text{m}$), each representing the standard deviation in latitude, longitude, and altitude. Figure 13 contains the impact on the localization of different MLAT noise levels. In contrast to the strong dependence on the GPS noise in the spoofed measurements, the MLAT noise has little impact on the accuracy of the spoofer localization. As a result, our localization approach does not rely on highly accurate MLAT measurements of the actual aircraft position and can still perform decently on relatively noisy data.

3) Impact of Spoofed Track Velocity: As the spoofed track velocity v_{track} is part of the scaling factor in the distance relations, we identify it to be another important parameter. The results for varying spoofed track velocities are depicted in Figure 14. For a spoofed track velocity of $v_{track} = 300\,\text{km/h}$, the accuracy decreases by nearly one fourth. The accuracy decreases further for a track velocity of $v_{track} = 100\,\text{km/h}$. Eventually, for track speeds lower than $v_{track} = 30\,\text{km/h}$, the spoofer localization fails to narrow down a useful search radius. However, considering less GPS noise, we expect to see better results even for lower track velocities. The strong dependence on the track velocity is due to the scaling factor, which relates the observed distances to the spoofed track velocity and the speed of light. Hence, low velocities result in smaller distance differences among the spoofed aircraft and are relatively more affected by system-intrinsic noise.

V. Discussion

Combined Error Effects. The spoofer localization accuracy of *Crowd-GPS-Sec* depends on the GPS error, the MLAT error, and the spoofed track velocity. These three parameters are all components of the relations defined in Equation (7) and thus impact the accuracy. While the MLAT noise is less decisive, the GPS noise and the spoofed track velocity are significantly affecting the achievable accuracy. This is due to the small differences in spoofed aircraft positions with respect to the speed of light divided by the spoofed track velocity. In general, we expose the following relationship between the localization error E, the GPS noise σ_{GPS}, and the spoofed track velocity v_{track}:

$$E \propto \frac{\sqrt{2} \cdot \sigma_{\text{GPS}}}{v_{track}}, \tag{12}$$

with σ_{GPS} being scaled with $\sqrt{2}$ due to the Euclidean distance based on two normally distributed points in space. Hence, we can expect to see similar results for low track velocities with low GPS noise and high track velocities with high GPS noise.

Localizing Spoofers of Stationary Targets. The attacker model considered in this paper assumes that the spoofer's target is a moving object. If instead the target is stationary, the attacker could also spoof constant positions. While the detection would still work, the localization would fail since the differences in propagation delays between spoofer and aircraft would not be reflected in the reported position differences (compare $d_{i,j}$ in Equation (9)). One way to cope with such attackers is to additionally propagate GPS time synchronization information to the ground infrastructure. As time is evolving, the spoofer would have to imitate a progressing GPS time to remain undetected by the target. Having information about the time synchronization of affected aircraft would allow performing a localization by analogy. More specifically, if t denotes the real GPS time and $\hat{t}_i$ the reported time of aircraft i, the relation from Equation (7) can be rewritten to:

$$\mathrm{dist}(a_i, \mathrm{SP}) - \mathrm{dist}(a_j, \mathrm{SP}) = (\hat{t}_i - \hat{t}_j) \cdot \frac{c}{\delta}, \qquad (13)$$

where δ denotes a factor representing the spoofed GPS clock's speed. Equation (13) is independent from the spoofed position and therefore allows localizing spoofers, even if the target is stationary.

Applicability to Other Networks. The underlying idea of *Crowd-GPS-Sec* does not only apply to aircraft but can also be relevant to GPS spoofing attacks on cars, trucks, ships, or other vehicles on ground. Similar to the broadcasting of avionic position reports via ADS-B or Flarm, vehicular systems could also report state information to, e. g., roadside units. The combined reports can then be used to run our spoofing detection and localization scheme. Even though the speeds of vehicles are comparably low, the density of affected targets is much higher and the GPS filtering is expected to be more conditioned. Eventually, we envision the merging of information provided by different networks. In particular, each spoofed system, such as aircraft, vehicles, vessel, etc., can collaborate by sharing their information in a crowdsourcing manner.

VI. Related Work

As GPS is known to be vulnerable to spoofing attacks [2], [5], [8], [31], several works demonstrated their feasibility [3], [4], [10], [32], [33]. Attacks can target different domains such as vehicle navigation systems [4], [10], [32] or critical infrastructures [6]. The requirements for successful GPS spoofing attacks are analyzed in [7]. Attacks that also change the data content of the signals are discussed in [44]. It is worth noting that GPS spoofing has also been proposed as a countermeasure, e. g., to defend against hostile UAVs [8], [11], [32] by means of hijacking or misguidance.

General techniques for detecting and localizing wireless spoofing attacks (not specific to GPS satellite signals) are proposed by Chen et al. [45]. The authors use received signal strength (RSS) readings from different locations and compare them against RSS maps built during an offline calibration phase to locate the spoofer. They evaluated their scheme in 802.11 and 802.15.4 networks. Later, Yang et al. [46] extended the scheme to deal with attackers which vary their transmission power. Rather than using direct RSS values, they consider RSS differences at multiple locations.

A rich body of countermeasures specific to GPS exists in the literature which can be categorized into *prevention* and *detection* measures. In order to prevent spoofing of GPS signals, several works propose the use of cryptographic techniques to authenticate satellite signals [13]–[16]. This is similar to how military GPS signals are protected. However, cryptographic techniques require profound modifications of the GPS infrastructure as well as a key distribution system which is challenging to implement for applications with disconnected receivers. Further, the use of encryption alone does not protect against signal replaying attacks [33].

The detection of GPS spoofing attacks has also received considerable attention in the literature. Overviews can be found in [18] and [31]. Akos [19] suggests to monitor the incoming signal power and the state of the automatic gain control. Another technique called SPREE relies on auxiliary peak tracking [17] to detect suspicious peaks from signals with weaker acquisition correlation peaks. Psiaki et al. [47] propose a detection scheme which uses an additional reference receiver to correlate the received signal with authentic signals assuming the inclusion of the encrypted military signal. A spoofed signal does not correlate with the reference node's received signal and the attack can be detected.

A different class of detection approaches deploys multiple antennas. Tippenhauer et al. [7], [48] use multiple co-located GPS receivers whose calculated positions and times are compared; coinciding locations indicate an attack. A dual antenna receiver setup to determine the angle of arrival of incoming signals is proposed by Montgomery et al. [20] and extended by Psiaki et al. [49] to include differential carrier phase measurements. Magiera and Katulski [21] suggest even the use of arrays of antennas showing that antenna diversity is effective at detecting single antenna spoofers without knowledge of the target's position. Although these detection approaches do not require changes to the GPS infrastructure, they assume more sophisticated GPS receivers which would significantly increase the complexity, size, costs, and power requirements. This, however, is contradictory to the objectives of GPS. It is also worth noting that, in principle, almost any passive localization technique (such as multilateration) could be used to locate GPS spoofers. However, in contrast to our approach, these methods assume a direct line-of-sight between the localization system and the attacker. As a consequence, this requires a dedicated infrastructure which covers all potential attacker positions.

The authors of [23], [50]–[52] have proposed techniques to detect spoofing attacks in ADS-B. However, the threat model in these works is different as they consider spoofed ADS-B signals and not spoofed GPS signals. These techniques are therefore not capable of localizing GPS spoofers such as in *Crowd-GPS-Sec.*

VII. Conclusion

This work presented *Crowd-GPS-Sec*, an independent system to detect and localize GPS spoofing attacks targeted at aircraft and UAVs. *Crowd-GPS-Sec* is lightweight and leverages existing wireless air traffic broadcast infrastructures, the ADS-B and Flarm systems, to identify spoofing attacks from a remote location—possibly far from where the attack is happening. We have shown that our approach is effective at localizing spoofing devices by using differences in reported positions by multiple aircraft. Using simulations based on real-world input from the OpenSky Network, we have demonstrated that *Crowd-GPS-Sec* achieves attack detection delays below two seconds and an attacker localization accuracy of around 150 meters after 15 minutes of monitoring time.

Acknowledgment

The authors would like to thank Markus Fuchs and Giorgio Tresoldi for their help with the experiments and the evaluations. This work was supported in parts by the DFG Research Training Group UbiCrypt GRK 1817/1 and the BMBF SysKit project 16KIS0664.

References

[1] *Global Positioning System Standard Positioning Service Performance Standard*, 4th ed., U.S. Department of Defense, Sep. 2008.

[2] Anon., "Vulnerability Assessment of the Transportation Infrastructure Relying on the Global Positioning System," John A. Volpe National Transportation Systems Center, Tech. Rep. Final Report, Aug. 2001.

[3] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner Jr., "Assessing the Spoofing Threat: Development of a Portable GPS Civilian Spoofer," in *International Technical Meeting of the Satellite Division of The Institute of Navigation*, ser. ION GNSS '08, Savannah, GA, USA, Sep. 2008, pp. 2314–2325.

[4] M. L. Psiaki and T. E. Humphreys, "Attackers can spoof navigation signals without our knowledge. Here's how to fight back GPS lies," *IEEE Spectrum*, vol. 53, no. 8, pp. 26–53, Aug. 2016.

[5] T. E. Humphreys, "Statement on the Vulnerability of Civil Unmanned Aerial Vehicles and Other Systems to Civil GPS Spoofing," The University of Texas at Austin, Tech. Rep., Jul. 2012, submitted to the Subcommittee on Oversight, Investigations, and Management of the House Committee on Homeland Security.

[6] D.-Y. Yu, A. Ranganathan, T. Locher, S. Čapkun, and D. Basin, "Short Paper: Detection of GPS Spoofing Attacks in Power Grids," in *ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '14. Oxford, United Kingdom: ACM, Jul. 2014, pp. 99–104.

[7] N. O. Tippenhauer, C. Pöpper, K. B. Rasmussen, and S. Čapkun, "On the Requirements for Successful GPS Spoofing Attacks," in *ACM Conference on Computer and Communications Security*, ser. CCS '11. Chicago, IL, USA: ACM, Oct. 2011, pp. 75–86.

[8] T. E. Humphreys, "Statement on the Security Threat Posed by Unmanned Aerial Systems and Posssible Countermeasures," The University of Texas at Austin, Tech. Rep., Mar. 2015, submitted to the Subcommittee on Oversight and Management Efficiency of the House Committee on Homeland Security.

[9] M.-A. Russon. (2015, May) Wondering how to hack a military drone? It's all on Google. International Business Times. [Online]. Available: http://www.ibtimes.co.uk/wondering-how-hack-military-drone-its-all-google-1500326

[10] J. A. Bhatti and T. E. Humphreys, "Hostile Control of Ships via False GPS Signals: Demonstration and Detection," The University of Texas at Austin, Tech. Rep., 2014.

[11] C. Sebastian. (2016, Dec.) Getting lost near the Kremlin? Russia could be 'GPS spoofing'. [Online]. Available: http://money.cnn.com/2016/12/02/technology/kremlin-gps-signals

[12] OSQZSS. (2017) Software-Defined GPS Signal Simulator. [Online]. Available: https://github.com/osqzss/gps-sdr-sim

[13] L. Scott, "Anti-Spoofing & Authenticated Signal Architectures for Civil Navigation Systems," in *International Technical Meeting of the Satellite Division of The Institute of Navigation*, ser. ION GPS/GNSS '03, Portland, OR, USA, Sep. 2003, pp. 1543–1552.

[14] M. G. Kuhn, "An Asymmetric Security Mechanism for Navigation Signals," in *International Conference on Information Hiding*, ser. IH '04. Toronto, Ontario, Canada: Springer, May 2004, pp. 239–252.

[15] G. W. Hein, F. Kneissl, J.-A. Ávila Rodríguez, and S. Wallner, "Authenticating GNSS: Proofs against Spoofs," *Inside GNSS*, vol. 2, no. 5/6, pp. 58–63/71–78, 2007.

[16] K. Wesson, M. Rothlisberger, and T. E. Humphreys, "Practical Cryptographic Civil GPS Signal Authentication," in *International Technical Meeting of The Satellite Division of the Institute of Navigation*, ser. ION GNSS '11, Portland, OR, USA, Sep. 2011, pp. 3335–3345.

[17] A. Ranganathan, H. Ólafsdóttir, and S. Čapkun, "SPREE: A Spoofing Resistant GPS Receiver," in *ACM Conference on Mobile Computing and Networking*, ser. MobiCom '16. New York, USA: ACM, Oct. 2016, pp. 348–360.

[18] M. L. Psiaki and T. E. Humphreys, "GNSS Spoofing and Detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, Apr. 2016.

[19] D. M. Akos, "Who's Afraid of the Spoofer? GPS/GNSS Spoofing Detection via Automatic Gain Control (AGC)," *NAVIGATION, Journal of the Institute of Navigation*, vol. 59, no. 4, pp. 281–290, Dec. 2012.

[20] P. Y. Montgomery, T. E. Humphreys, and B. M. Ledvina, "Receiver-Autonomous Spoofing Detection: Experimental Results of a Multi-Antenna Receiver Defense Against a Portable Civil GPS Spoofer," in *International Technical Meeting of The Institute of Navigation*, ser. ION '09, Anaheim, CA, USA, Jan. 2009, pp. 124–130.

[21] J. Magiera and R. J. Katulski, "Detection and Mitigation of GPS Spoofing Based on Antenna Array Processing," *Journal of Applied Research and Technology*, vol. 13, no. 1, pp. 45–57, Feb. 2015.

[22] OpenSky Network. (2017) OpenSky Network. [Online]. Available: https://opensky-network.org

[23] D. Moser, P. Leu, V. Lenders, A. Ranganathan, F. Ricciato, and S. Čapkun, "Investigation of Multi-device Location Spoofing Attacks on Air Traffic Control and Possible Countermeasures," in *ACM Conference on Mobile Computing and Networking*, ser. MobiCom '16. New York, USA: ACM, Oct. 2016, pp. 375–386.

[24] B. Hofmann-Wellenhof, H. Lichtenegger, and J. Collins, *Global Positioning System: Theory and Practice*, 5th ed. Springer, 2001.

[25] P. F. Swaszek and R. J. Hartnett, "Spoof Detection Using Multiple COTS Receivers in Safety Critical Applications," in *International Technical Meeting of The Satellite Division of the Institute of Navigation*, ser. ION GNSS+ '13, Nashville, TN, USA, Sep. 2013, pp. 2921–2930.

[26] ——, "A Multiple COTS Receiver GNSS Spoof Detector – Extensions," in *International Technical Meeting of The Institute of Navigation*, ser. ION '14, San Diego, CA, USA, Jan. 2014, pp. 316–326.

[27] P. F. Swaszek, R. J. Hartnett, M. V. Kempe, and G. W. Johnson, "Analysis of a Simple, Multi-Receiver GPS Spoof Detector," in *International Technical Meeting of The Institute of Navigation*, ser. ION '13, San Diego, CA, USA, Jan. 2013, pp. 884–892.

[28] Flarm. (2017) Flarm. [Online]. Available: https://flarm.com

[29] K. Jansen and C. Pöpper, "Opinion: Advancing Attacker Models of Satellite-based Localization Systems—The Case of Multi-device Attackers," in *ACM Conference on Security and Privacy in Wireless and Mobile Networks*, ser. WiSec '17. Boston, MA, USA: ACM, Jul. 2017, pp. 156–159.

[30] The Executive Director of the Joint Air Power Competence Center (JAPCC), "JAPCC Strategic Concept of Employment for Unmanned Aircraft Systems in NATO," Joint Air Power Competence Center (JAPCC), Tech. Rep. UAS CONEMP Report, Jan. 2010.

[31] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "GPS Vulnerability to Spoofing Threats and a Review of Antispoofing Techniques," *International Journal of Navigation and Observation*, vol. 2012, May 2012.

[32] A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Unmanned Aircraft Capture and Control via GPS Spoofing," *Journal of Field Robotics*, vol. 31, no. 4, pp. 617–636, Jul. 2014.

[33] P. Papadimitratos and A. Jovanovic, "GNSS-based Positioning: Attacks and Countermeasures," in *IEEE Military Communications Conference*, ser. MILCOM '08. San Diego, CA, USA: IEEE, Nov. 2008, pp. 1–7.

[34] creaktive. (2017) nRF905 demodulator/FLARM decoder. [Online]. Available: https://github.com/creaktive/flare

[35] *Air Traffic Control - Order JO 7110.65W*, U.S. Department of Transportation, Dec. 2015.

[36] M. Strohmeier, M. Schäfer, R. Pinheiro, V. Lenders, and I. Martinovic, "On Perception and Reality in Wireless Air Traffic Communication Security," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 6, pp. 1338–1357, Jun. 2017.

[37] Ettus Research. (2017) Universal Software Radio Peripheral (USRP). [Online]. Available: https://www.ettus.com

[38] D. L. Adamy, *EW 101: A First Course in Electronic Warfare*, 1st ed. Artech House, 2001.

[39] 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 Conference on Information Processing in Sensor Networks*, ser. IPSN '14. Berlin, Germany: IEEE, Apr. 2014, pp. 83–94.

[40] A. Beygelzimer, S. Kakade, and J. Langford, "Cover Trees for Nearest Neighbor," in *ACM International Conference on Machine Learning*, ser. ICML '06. Orlando, FL, USA: ACM, Jun. 2006, pp. 97–104.

[41] O. Jowett. (2016) Mode S Multilateration Server. [Online]. Available: https://github.com/mutability/mlat-server

[42] W. H. L. Neven, T. J. Quilter, R. Weedon, and R. A. Hogendoorn, "Wide Area Multilateration," Eurocontrol, Tech. Rep. NLR-CR-2004-472, Aug. 2005.

[43] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[44] T. Nighswander, B. Ledvina, J. Diamond, R. Brumley, and D. Brumley, "GPS Software Attacks," in *ACM Conference on Computer and Communications Security*, ser. CCS '12. Raleigh, NC, USA: ACM, Oct. 2012, pp. 450–461.

[45] Y. Chen, W. Trappe, and R. P. Martin, "Detecting and Localizing Wireless Spoofing Attacks," in *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, ser. SECON '07. San Diego, CA, USA: IEEE, Jun. 2007, pp. 193–202.

[46] J. Yang, Y. Chen, W. Trappe, and J. Cheng, "Detection and Localization of Multiple Spoofing Attackers in Wireless Networks," *IEEE Transactions on Parallel and Distributed Systems*, vol. 24, no. 1, pp. 44–58, Jan. 2013.

[47] M. L. Psiaki, B. W. O'Hanlon, J. A. Bhatti, D. P. Shepard, and T. E. Humphreys, "GPS Spoofing Detection via Dual-Receiver Correlation of Military Signals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2250–2267, Oct. 2013.

[48] K. Jansen, N. O. Tippenhauer, and C. Pöpper, "Multi-Receiver GPS Spoofing Detection: Error Models and Realization," in *Annual Computer Security Applications Conference*, ser. ACSAC '16. Los Angeles, CA, USA: ACM, Dec. 2016, pp. 237–250.

[49] M. L. Psiaki, B. W. O'Hanlon, S. P. Powell, J. A. Bhatti, K. D. Wesson, T. E. Humphreys, and A. Schofield, "GNSS Spoofing Detection using Two-Antenna Differential Carrier Phase," in *International Technical Meeting of The Satellite Division of the Institute of Navigation*, ser. ION GNSS+ '14, Tampa, FL, USA, Sep. 2014, pp. 2776–2800.

[50] M. Schäfer, V. Lenders, and J. Schmitt, "Secure Track Verification," in *IEEE Symposium on Security and Privacy*, ser. SP '15. San Jose, CA, USA: IEEE, May 2015, pp. 199–213.

[51] M. Strohmeier, M. Smitt, M. Schäfer, V. Lenders, and I. Martinovic, "Crowdsourcing Security for Wireless Air Traffic Communications," in *International Conference on Cyber Conflict*. Tallinn, Estonia: IEEE, Jun. 2017.

[52] M. Strohmeier, V. Lenders, and I. Martinovic, "Lightweight Location Verification in Air Traffic Surveillance Networks," in *ACM Cyber-Physical System Security Workshop*. Singapore: ACM, Apr. 2015.