18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Detecting Malicious ADS-B Transmitters Using a Low-Bandwidth Sensor Network

Márcio Monteiro
Alexandre Barreto
Research Division
Instituto de Controle do Espaço Aéreo
São José dos Campos, São Paulo 12228–903
http://www.icea.gov.br
(contemmcm,barretoabb)@icea.gov.br

Thabet Kacem
Duminda Wijesekera
Paulo Costa
Radio and Radar Engineering Lab
George Mason University
Fairfax, Virginia 22030–4444
http://radio.vse.gmu.edu
(tkacem,dwijesek,pcosta)@gmu.edu

***Abstract*—Automatic Dependent Surveillance Broadcast (ADS-B) has been proposed as both an extension and an alternative to Primary and Secondary Surveillance RADAR. Although ADS-B has many advantages, security was not a key issue in its design. Packets are sent in clear text, making it vulnerable to many attacks. A main concern is that some of these attacks can be easily implemented using inexpensive ADS-B transmitters. For instance, ADS-B is vulnerable to injection attacks, which introduce ghost aircraft into an surveillance system. In this paper we propose a method that addresses these issues by detecting malicious ADS-B transmitters using a network of sensors with an associated data fusion center. A key aspect of our solution is that the data fusion process has a low message exchange overhead, enabling its applicability to current constrains of ADS-B sensors. Our method can be implemented using commercial off-the-shelf ADS-B receivers.**

I. Introduction

Airspace surveillance systems rely heavily on Primary Surveillance RADAR (PSR) and Secondary Surveillance RADAR (SSR), which are effective but tend to be expensive as well as difficult to deploy and maintain in remote areas (sea, mountains, etc.). Automatic Dependent Surveillance Broadcast (ADS-B) is a technology designed for extending or replacing RADARs for air traffic management and control. ADS-B systems can be deployed in remote areas at a fraction of the cost of a RADAR system, and are much easier to maintain.

Nevertheless, several researchers, such as Costin and Francillon [1], have shown that ADS-B was not designed with security as a key objective, and consequently present vulnerabilities that can be exploited to disrupt the operation of air surveillance systems using the technology. The main exploitable vulnerability of ADS-B is their clear-text transmissions, which makes the technology susceptible to attacks based on packet injection, eavesdropping, and denial of service attacks. Various researchers (e.g. [1], [2], [3]) have already demonstrated that ADS-B based surveillance systems can be disrupted using a combination of relatively inexpensive equipment and advances in open source software defined radios platforms. Many solutions have been proposed to address these vulnerabilities, but the vast majority of these require modifications to the ADS-B protocol that are unfeasible from both the technical and commercial aspects. For instance, adding security features (e.g. encryption) requiring changes in the message size would require costly modifications to the currently deployed systems.

An alternative is to use so-called multilateration, a technique that locates the source of a wireless transmitter to validate the authenticity of ADS-B messages by comparing the declared and detected positions [4]. Multilateration - also known as hyperbolic positioning, is a technique that takes the time of arrival (TOA) of a known wireless signal at a set of sensors with known locations. Consequently, given that the signal travels at the speed of light, locating the source of the signal from the time of arrival at several locations becomes a purely geometric problem. However, because the transmitter location is initially unknown, time difference of arrival (TDOA) is used to derive the position of the transmitter.

One of the main advantages of multilateration is the possibility of reusing existing infrastructure and without modifying the ADS-B packet format. If successful, this technique can be directly applied for mitigating attacks such as ghost aircraft injection. One caveat of this solution is that transmitter location accuracy is very sensitive to timing [5]. Multilateration was firstly envisioned as a backup for ADS-B, as well as a potential intermediate technology for supporting the deployment of ADS-B systems. Consequently, many complex architectures were proposed to achieve wide-area multilateration. We show that by constraining ADS-B signals over smaller areas, multilateration can be successfully deployed with simpler architectures that use low network bandwidth and commercial off-the-shelf (passive) receivers.

The rest of the paper is organized as follows. Section II presents previous work on multilateration using ADS-B signals. Section III describes the method we propose to detect and locate malicious ADS-B transmitters. Section IV describe our preliminary simulation results. Section V conveys the conclusions derived from the promising simulation results we have obtained in this initial research.

II. Related Work

Multilateration have been successfully used for airport surface movement surveillance [6], which involves short distances and ground altitudes (i.e. up to 60 m height). However, research on wide area multilateration (WMLAT) has been gaining momentum, mostly due to the many advantages of lower cost and less complexity. Previous work by Strohmeier [4] and Neven [7] show how multilateration can be used to validate ADS-B target reports based on the ADS-B signals itself. This same technique can be used to prevent cyber-attacks such as aircraft ghost injection, spoofing/impersonation and replays.

Several research efforts on WAMLAT and ADS-B have been proposed that differ on how to estimate the time of arrival (TOA) and how to synchronize the receivers. In particular, the most common architecture uses cross-correlation to estimate TDOA and Global Navigational Satellite Systems (GNSS) to synchronize sensors [8], [9].

Steffers et al. [9] proposed a method to estimate the time of arrival (TOA) at sensors with an accuracy of the order of nanoseconds. They use local correlation between the received signal and an ideal ADS-B preamble while the sensors are synchronized using GPS. In [10], the authors extend those ideas by applying a tracking method using non-linear Kalman filters. In [11], the authors proposed another time of arrival (TOA) estimator based on matched filters and maximum likelihood estimation.

Daskalakis and Martone [12] presented a product called HITS that performs both ADS-B and multilateration. For synchronization, a reference transmitter with known position and transmission rate is used. They present practical results in the Gulf of Mexico area with very promising results.

Johnson et al. [13] presented a study case for another WAMLAT product, Thales, that was deployed in Afghanistan in order to improve situation awareness and safety for overflight traffic.

Niles et al. [5] proposed WAMLAT as an alternative navigation system. It contains a discussion about how to optimize the number of sensors and its positions, in order for the system to meet some minimum performance requirements.

Our approach, presented in the next section, extends these research efforts by proposing a low complexity method to validate ADS-B target reports through multilateration, making it possible to track unreliable transmitters.

III. The Proposed Method

Using correlation between two signals is an effective method to obtain a precise measurement of the time difference of arrival (TDOA). However, this requires a digitized version of the signal to be stored, timestamped and transmitted to a central correlation unit from all sensors at some fixed (high) rate. Hence network delays becomes an issue for the wide area multilateration. For example, low cost receivers with sampling rate of 3.2 MS/s generate a substantial amount of data to be transmitted, requiring a fast network for producing accurate results. Although wireless networks are the preferred choice because of their lower deployment cost, bandwidth requirements becomes a challenge.

In this work we propose to timestamp the decoded messages that contain only 112 bits for SSR Mode-S extended squitter. This is less than 1% of the digitized version of an ADS-B message, given that every sample is encoded using a single-precision 32-bit floating-point number. This would greatly reduce the network bandwidth problem. Additional uncertainties that need to be addressed for successful deployment are described shortly.

A. Time Difference of Timestamps (TDOT)

A signal that reaches a sensor antenna must be demodulated, decoded and errors-checked prior to being timestamped. We define the time interval between the demodulation and timestamping at the n^{th} sensor as the processing time p_n. Although p_n is expected to be small, it is not negligible. Because ADS-B messages travels at the speed of light, any microsecond error in timing leads to an error of 300 meters in the distance measurement.

In this paper, we model the processing time p_n as a random variable with a Gaussian distribution, with mean μ_n and variance σ_n^2. Then, the timestamp τ_n of a message at sensor n can be expressed as

$$\tau_n = t_n + p_n \tag{1}$$

where t_n is the time of arrival of the message at the n^{th} sensor - a function of the distance between the emitter and the sensor.

$$t_n = \frac{\sqrt{(x_n - x_E)^2 + (y_n - y_E)^2 + (z_n - z_E)^2}}{c} \tag{2}$$

where c is the speed of light. The coordinates (x_n, y_n, z_n) and (x_E, y_E, z_E) are, respectively, the Cartesian coordinates of sensor n and the emitter E as shown in Fig. 1.

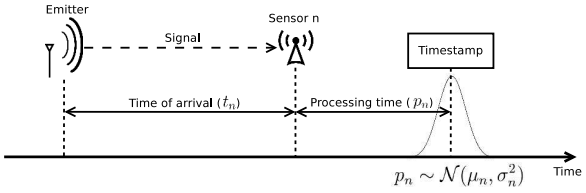

Fig. 1: Multilateration without correlation.

Hence the measured signal arrival time difference of timestamps (TDOT) between sensors i and j are:

$$\tau_i - \tau_j = (t_i - t_j) + (p_i - p_j) \tag{3}$$

Under the premise that all sensors have equivalent processing delay characteristics (because we assume them to use the same hardware, drivers and software), the mean processing time of all sensors would be the same. This also applies to the variances of the processing times. Consequently, we take $\mu_n = \mu$ and $\sigma_n^2 = \sigma^2 \ \forall n$ and rewrite the processing time p_n with a statistically equivalent expression, using a zero-mean Gaussian random variable:

$$p_n = \mu_n + \mathcal{N}(0, \sigma_n^2) \tag{4}$$

Hence, the time difference of timestamps can be expressed as follows:

$$\begin{aligned} \tau_i - \tau_j &= (t_i - t_j) + \left[\mu + \mathcal{N}(0,\sigma^2) - \mu - \mathcal{N}(0,\sigma^2)\right] \\ &= (t_i - t_j) + \left[\mathcal{N}(0,\sigma^2) - \mathcal{N}(0,\sigma^2)\right] \end{aligned} \tag{5}$$

Notice the mean processing time μ cancels. In addition, because a Gaussian distribution with a zero mean is symmetric around zero, $\mathcal{N}(0,\sigma^2)$ is statistically equivalent to $-\mathcal{N}(0,\sigma^2)$. Now we add the two Gaussian distributions and express the difference of timestamps (TDOT) as:

$$\tau_i - \tau_j = t_i - t_j + \mathcal{N}(0, 2\sigma^2) \tag{6}$$

where the term $\mathcal{N}(0, 2\sigma^2)$ is, hereafter, referred to as the noise of the estimation process, which we sometimes describe using the standard deviation σ. For example, a $\sigma = 1$ μsec at sensors leads to a noise standard deviation of $\sqrt{2}$ μsec, resulting in a distance estimation error of 423.97 meters at the speed of light. So, even a small variances in the processing time between sensors can lead to a considerable differences in the resulting estimation of the emitter's location.

Fig. 2 shows an example of the actual time difference of arrival and the measured time difference of timestamps between two sensors. The scenario is described in details in Section IV. Fig. 3 shows the effect of the small time variances at sensors on the location estimates of emitters computed using multilateration.

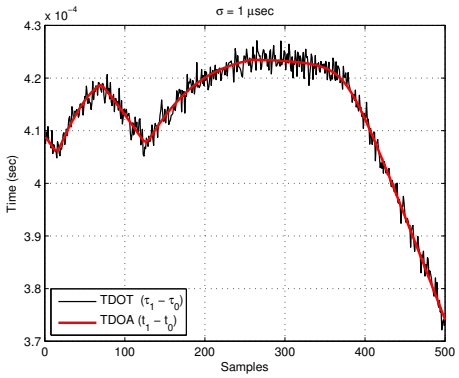

Fig. 2: Time difference of arrival and timestamp between two sensors.

However, because of the zero-mean property of the noise, the accuracy of emitter's location can be augmented by filtering a sequence of measurements with a time series smoother. We can do so using a simple moving average or a more sophisticated Kalman filter. Consequently, we can reduce the effects of the variance of TDOT.

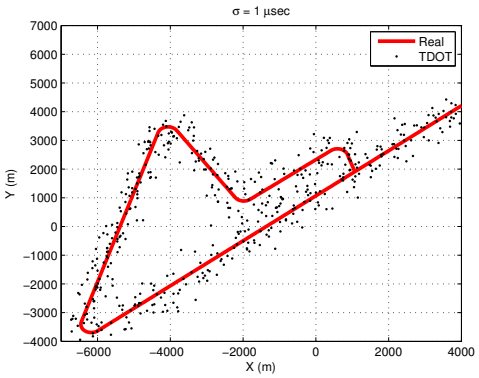

Fig. 3: Estimating position using time difference of timestamp (TDOT).

B. Validating ADS-B Messages

In summary, the multilateration algorithm detects the source of the signal with a quantifiable error. Hence the difference between the reported position and the estimated source position can be used to determine whether the ADS-B message is forged or can be attributed to an unauthorized relay. In either case, we have sufficient information to drop unreliable messages and protect the air traffic surveillance system while providing means to track the source of the signal and possibly identify the transmitter.

Fig. 4 shows two illustrations of performing ADS-B validation using multilateration. In the first example, the source is considered reliable because the horizontal distance between the reported and the estimated position is sufficiently short (below a threshold). In the second example, however, the source is considered unreliable because it does not meet that requirement.

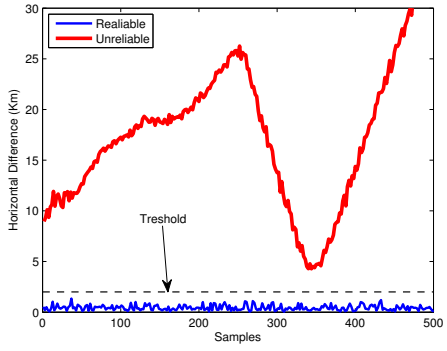

Fig. 4: Horizontal Distance between Estimated Position and the Signal Source.

Although the uncertainty introduced by the TDOT could affect the precision of the wide area multilateration algorithms, it provides an effective method for classifying ADS-B messages as unreliable and then tracking the transmitter under suspicion

over time.

IV. Simulation Results

This section presents some simulation results of the usage of time difference of timestamps (TDOT) in wide area multilateration. We assume that all sensors' clocks are perfectly synchronized, all antennas are omnidirectional, and channels are ideal (i.e. no energy loss during signals propagation). Our estimation of emitter positions uses a a flat model of the Earth, which suffices for the range of distances we are currently using but that will be substituted by a tridimensional model as we progress in our experiments.

The simulation presented in this paper used four sensors that were placed in a squared disposition. In order to simplify the TDOA calculations, a fifth sensor is placed at the origin of the square. This scheme can be seen in Fig. 7, which covers an area of 200×200 Km.

In subsection IV-A we discuss the evaluation of the impact of the standard deviation (σ) of sensors' processing time to the accuracy of the estimated positions, and use a known trajectory of an aircraft as a case study. The findings also show how the time series smoothers greatly reduce the noise on TDOT measurements, thereby providing better estimates. Subsection IV-B shows the effective range of the TDOA system based on sensors' positions and the standard deviation of sensor's processing time. Finally, subsection IV-C presents a practical example of how the proposed method can be used to detect an advanced ADS-B cyber attack.

A. Horizontal Error Due to Noise

In order to analyze how the noise on TDOT measurements translates into a positioning error while estimating the horizontal position of ADS-B emitters, we did numerical simulations using a known trajectory of an aircraft taking off from Brasília Airport (Brazil). During the simulations, this aircraft broadcasted Mode S extended squitter messages every 0.5 seconds. For every transmitted message, the time of arrival at all sensors is computed using (2). After that, we randomly generate the processing time $p_n \sim \mathcal{N}(\mu, \sigma^2)$ for every sensor n, then record the timestamps $\tau_n = t_n + p_n$ in order to compute the TDOT among sensors. Finally, the TDOT values are provided to the to the TDOA solver, which returns an estimate for the position of the emitter.

With sufficient measurements, we were able to filter (or smooth) the horizontal emitter position estimates in order to mitigate noise in timestamps. The filtering process can happen on two different domains: on the TDOT measurements itself (time domain), or on the estimates of the emitters position (Cartesian domain). Simulations have shown that both approaches provides very similar results. However, we preferred to filter the output of TDOA solver because it involved less computations (e.g., for a five sensors scenario, position filtering requires computations of two horizontal dimensions, X and Y, while filtering TDOT requires computing four pairwise TDOT measurements).

Fig. 5 shows a simulation run for $\sigma = 10^{-6}$ sec. Notice that for such variation in the processing time the raw TDOT measurements provide estimates that make it very difficult to visualize a trajectory of the aircraft. Conversely, filtering those results provides outcome that are much closer to the actual trajectory of the aircraft.

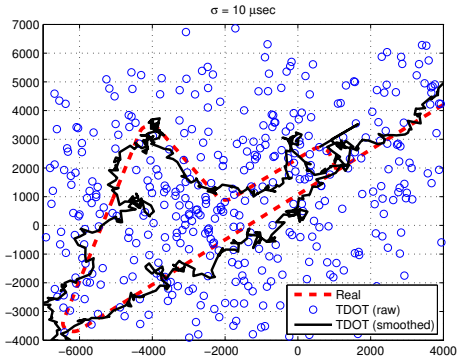

Fig. 5: Comparison between unfiltered and filtered position estimate.

Fig. 6 presents the mean horizontal error as a function of the standard deviation σ. Those curves were obtained by varying the σ from zero to 50 microseconds and recording the horizontal error for each iteration. This process was repeated a number of times in order to allow the mean horizontal error for different σ to stabilize. Notice that the error increases linearly with the noise and the effectiveness of the filtering process in improving the accuracy of TDOT measurements.

These results expose the practical limit for the effective variance of the timestamping process, which indicates the need for bounding it in deployed systems. For example, in order to guarantee a mean error less than 1 Km, system engineers must ensure that the variations in the processing time do not exceed an average of 19 microseconds according to Fig. 6.

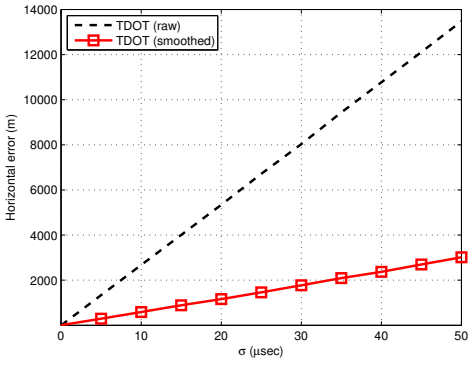

Fig. 6: Mean horizontal error for different values of σ using a known trajectory. Smoothing effectively improves accuracy.

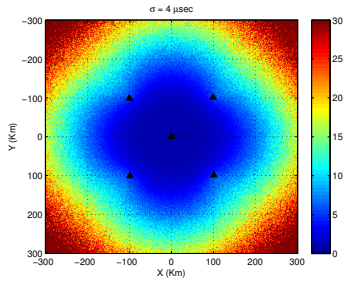

(a) $\sigma = 4$ μsec

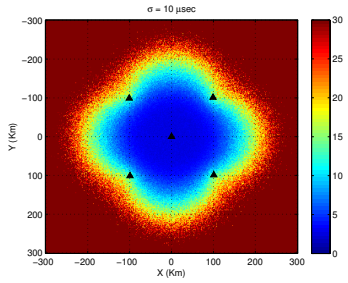

(b) $\sigma = 10$ μsec

Fig. 7: Effective range for different variations in the processing time.

B. *Effective Range*

Fig. 7 shows the effective range for wide area multilateration based on TDOT measurements. The sensors are arranged in a configuration known as "Square-5" [7], with a 140 Km base line. This result was obtained by varying the emitter's horizontal position (X and Y) on a grid of 600×600 Km. The altitude of the emitter was fixed at 30 Km. The simulation was repeated a sufficient number of times. Fig. 7 presents the mean value of the horizontal error, which is shown by the color bar (in Km). As expected, the region within the sensor's range has a lower error (represented by the color bar), while the performance outside this region quickly degrades. These results were obtained using a emitter located at 30 Km of altitude by repeating the experiment 40 times.

C. *Practical Application*

We now consider the scenario of having two aircraft. The first, the victim or the attack, is a commercial aircraft taking off from an airport with a large number of civilians on board. The second aircraft, the attack perpetrator, is a military fighter with stealth technology performing two types of cyber attacks:

1) Injecting ghost aircraft to surrounding ADS-B receivers, in order to disguise its main attack.
2) Spoofing the victim's ICAO address, in order to deceive ground stations with respect to the position of the victim.

As the enemy possesses stealth technology, it cannot be detected using primary surveillance RADARs. The main objective of the enemy is to intercept and gun down the victim, but making its actions to look like an accident by forging a crash of its victim using fake ADS-B messages.

Fig. 8 shows the actual trajectory of both aircraft. The victim is taking off from an airport and the enemy is flying in the vicinities of the terminal at a very high altitude, while sending forged ADS-B messages. The outcome of the attack is shown in Fig. 9. In the figure, it is possible to see the location of the victim constantly jumping from its actual position (reported by the victim aircraft) to the forged position, which is crafted by the enemy.

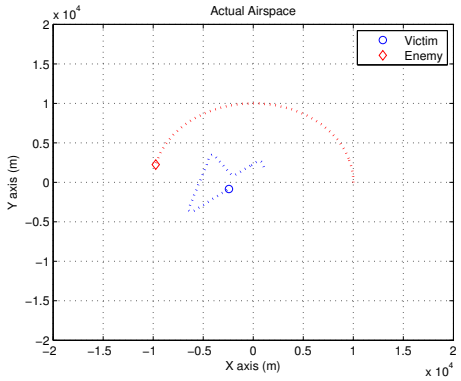

Fig. 8: The stealthy enemy is flying at the airport vicinity at high altitude, injecting ghost tracks and spoofing the victim's ICAO address.

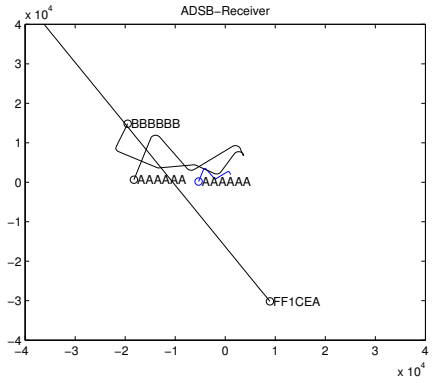

Fig. 9: The view of the air traffic controller.

By using the difference between timestamps (TDOT) for all the ADS-B messages, it is possible to calculate the position of the emitters of every message with quantifiable uncertainty. As described in subsection III-B, if the horizontal distance between the position reported in the ADS-B message's content and the one measured via the TDOA is too long, then the message can be classified as unreliable. Instead of discarding this message, the system can store it and use for tracking the source of the unreliable messages.

Fig. 10 shows the final result of the proposed low-bandwidth method. Using only ADS-B transponders we were able to

distinguish between actual and forged ADS-B messages (with respect to the transmitters location) and locate the source of unreliable ADS-B messages. In an actual system, it is conceivable to use a string of similar messages as the trigger for activating the Air Defense System to intercept the enemy and perform the appropriate actions.

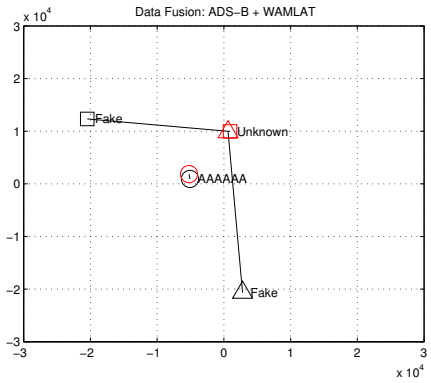

Fig. 10: Unreliable tracks can be easily spotted. The black markers indicate the reported position using ADS-B protocol and the red markers are the resulting of TDOA algorithm.

V. Final Remarks

We show a computationally feasible method for estimating a transmitter position using multilateration algorithms. Our method can be used in an architecture that combines inexpensive ADS-B receivers and a data fusion center that are capable of computing the location for each pair of receivers. By reducing the content of the transmitted messages between the ADS-B receivers and the fusion center, the data rates can be handled by a regular wireless network.

The uncertainty introduced due to the processing time makes this technique not suitable for being employed as the sole source for air surveillance systems. However, it can be successfully used to protect air traffic control systems from ghost injection attacks, while providing a mean for locating sources of suspect signals with computable error estimates. Considering the above-cited vulnerabilities of ADS-B systems against low cost transmitters available in the market today, a large number of potential attacks using malicious transmitters to disrupt air traffic control systems can be prevented, since air surveillance systems would have a much higher probability of detecting such attacks.

In order for this technique to be feasible and effective in practical deployments, system engineers must ensure that the processing time has no more than 19 microseconds of standard deviation. This can be achieved, for instance, by using known techniques for optimizing the computational time, such as avoiding alternate flows within code, removing unnecessary concurrent processes using the same CPU, or using a real-time operating system.

Acknowledgment

Márcio Monteiro and Alexandre Barreto would like to thank the financial support of the Brazilian agencies MCTI and FINEP (Ref. 04/2013/12).

References

[1] A. Costin and A. Francillon, "Ghost in the Air(Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices," in *BLACKHAT 2012, July 21-26, 2012, Las Vegas, NV, USA*, Las Vegas, ÉTATS-UNIS, 07 2012.

[2] M. Schfer, V. Lenders, and I. Martinovic, "Experimental analysis of attacks on nextgeneration air traffic communication," in *Applied Cryptography and Network Security*, ser. Lecture Notes in Computer Science, M. Jacobson, M. Locasto, P. Mohassel, and R. Safavi-Naini, Eds. Springer Berlin Heidelberg, 2013, vol. 7954, pp. 253–271.

[3] D. Magazu, R. F. Mills, J. W. Butts, and D. J. Robinson, "Exploiting the automatic dependent surveillance-broadcast system via false target injection," *Journal of Aviation and Aerospace Perspectives*, vol. 2, no. 2, pp. 5–19, 2014.

[4] M. Strohmeier, V. Lenders, and I. Martinovic, "On the security of the automatic dependent surveillance-broadcast protocol," *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1–1, 2014.

[5] F. A. Niles, R. S. Conker, M. B. El-Arini, D. G. O'Laighlin, and D. V. Baraban, "Wide Area Multilateration for Alternate Position, Navigation, and Timing (APNT)," MITRE-CAASD, Tech. Rep., 2012.

[6] J. Herrero, J. Besada Portas, F. Rodriguez, and J. Corredera, "ASDE and multilateration mode-S data fusion for location and identification on airport surface," in *Radar Conference, 1999. The Record of the 1999 IEEE*, 1999, pp. 315–320.

[7] W. Neven, T. Quitter, R. Weedon, and R. Hogendoorn, "Wide Area Multilateration Report on EATMP TRS 131/04 Version 1.1," National Aerospace Laboratory NLR, Tech. Rep., 2005.

[8] N. El Gemayel, S. Koslowski, F. Jondral, and J. Tschan, "A low cost TDOA localization system: Setup, challenges and results," in *Positioning Navigation and Communication (WPNC), 2013 10th Workshop on*, March 2013, pp. 1–4.

[9] C. Steffes, R. Kaune, and S. Rau, "Determining times of arrival of transponder signals in a sensor network using GPS time synchronization," in *Informatik 2011 - 6th Workshop Sensor Data Fusion: Trends, Solutions, Applications*, October 2011, pp. 481–481.

[10] R. Kaune, C. Steffes, S. Rau, W. Konle, and J. Pagel, "Wide area multilateration using ADS-B transponder signals," in *Information Fusion (FUSION), 2012 15th International Conference on*, July 2012, pp. 727–734.

[11] G. Galati, M. Leonardi, P. Magaro, and V. Paciucci, "Wide area surveillance using SSR mode S multilateration: advantages and limitations," in *Radar Conference, 2005. EURAD 2005. European*, Oct 2005, pp. 225–229.

[12] A. Daskalakis and P. Martone, "A technical assessment of ADS-B and multilateration technology in the gulf of mexico," in *Radar Conference, 2003. Proceedings of the 2003 IEEE*, May 2003, pp. 370–378.

[13] J. Johnson, H. Neufeldt, and J. Beyer, "Wide area multilateration and ADS-B proves resilient in afghanistan," in *Integrated Communications, Navigation and Surveillance Conference (ICNS), 2012*, April 2012, pp. A6–1–A6–8.