POSTER: Detecting False-Data Injection Attacks on Air Traffic Control Protocols

Martin Strohmeier, Ivan Martinovic

University of Oxford, UK
martin.strohmeier, ivan.martinovic@cs.ox.ac.uk

ABSTRACT

The world's airspaces are becoming increasingly crowded as manned and unmanned aircraft must coexist in the future. To handle this growth, new and more efficient protocols are being rolled out in most countries. Automatic Dependent Surveillance-Broadcast (ADS-B) is one of the core pieces of next generation air traffic management. Several publications in the academic and hacker community have highlighted the vulnerabilities of ADS-B and consequently the need for improved security. We analyze means to detect such attacks and propose a transparent intrusion detection system.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*Security and protection*

Keywords

ADS-B, Security, Intrusion Detection, Air Traffic Control

1. INTRODUCTION

In order to meet future demands in increasingly congested airspaces, the world's aviation authorities are currently upgrading their air-traffic management systems. The ADS-B protocol is at the core of the Next Generation Air Transportation (NextGen). In contrast to traditional radar surveillance technologies that measure the range and bearing of an aircraft from a ground-based antenna, ADS-B allows aircraft to determine their own position using a Global Navigation Satellite System (GNSS) and to broadcast it periodically over a radio frequency to ground stations or other aircraft. The ability to continuously broadcast position, heading, velocity, and other information lowers the necessity for more expensive and less accurate radar technologies. This improves the situational awareness of pilots and air traffic controllers significantly while reducing surveillance costs.

Consequently, the Federal Aviation Administration (FAA) made ADS-B mandatory for all aircraft by 2020. Until then, many aspects of ADS-B need further evaluation to ensure a safe adoption. Originally open by design, ADS-B lacks any security or authentication mechanism, making every passive and active attack on the wireless communication channel possible. This includes but is not limited to the injection or flooding of a ground station with ghost aircraft, virtual trajectory modification, aircraft disappearance, and aircraft spoofing [1, 3]. As introducing cryptographic primitives is infeasible at the current stage of the ADS-B roll out, there is an urgent need for transparent countermeasures. We research different ways to detect fraudulently injected data.

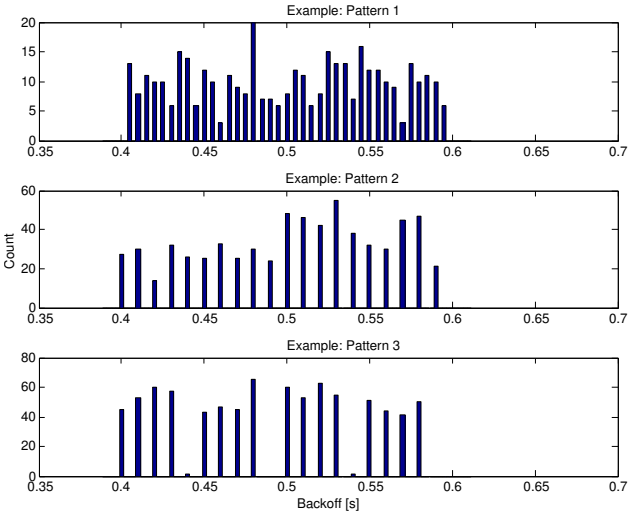

Figure 1: Example of varying random backoff distributions in positional ADS-B messages.

2. DETECTING FALSE DATA INJECTIONS

We identify two main scenarios that an attacker uses to inject data onto the wireless communication channel.

- The attacker injects a new ghost aircraft with a legitimate ICAO identifier and reasonable flight parameters.
- The attacker captures real ADS-B data in the area and plays it back at a later time without modification.

2.1 Detection of behavioural discontinuities

Our goal is to analyse patterns of aircraft messages to identify anomalous and malicious activity. For an attack to be successful, the attacker must prove sufficient knowledge of both the ADS-B protocol and customs in air traffic control. This includes the use of matching values for fields in the different message types but also sending all of the same message types the aircraft-specific transponder is broadcasting and doing so with the correct temporal spacing. A system that keeps track of an aircraft's historical message data can flag discontinuities when an attacker introduces forged messages without accurately copying the appropriate characteristics. A related possibility is to exploit the fact that there are a multitude of different transponder types in ADS-B-equipped aircraft today. These transponders exhibit a number of different behaviours on the data link level as well as the physical layer which can be utilized to validate incoming messages. We plan to identify such differences in both the

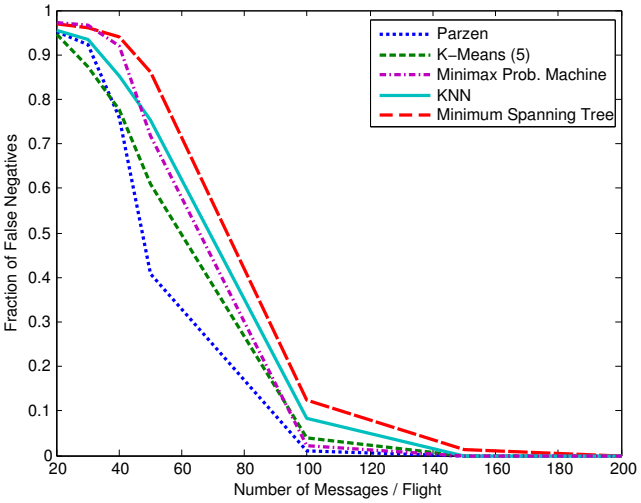

Figure 2: Anomaly detection using different classifiers. The graph shows the fraction of illegitimate flights missed by the classification.

time and the frequency domain, using manual and automatic features selection and classification, including principal component analysis and linear discriminant analysis. Copying these characteristics can be complex and difficult for an attacker and further restricts the available space in creating malicious ADS-B message content that will be considered legitimate by the intrusion detection system.

Feature Selection

We look at various approaches to fingerprint the transponders of ADS-B-equipped aircraft using distinct characteristics that are unique to the type of transponders or even the transponder itself such as clock skews or differences in the turn-on/off transient or the modulation of a radio signal. One example for a distinguishing feature on the data link layer would be the behavior of the random generator that generates the backoff between the periodically broadcasted ADS-B messages. The different transponders in commercial aircraft show very distinct behavior when we have a close look at the precise time periods between two subsequent messages (position, velocity, or call sign). One example of this backoff property between positional messages is shown in Fig. 1. We found various classes, each exhibited by between 5% and 30% of the monitored aircraft.

2.2 Pattern analysis of signal strength data

Further to matching aircraft and transponder behaviour, an attacker is required to know about signal propagation characteristics. When impersonating an existing aircraft, it is not sufficient to merely use or replay the correct 24-bit ICAO identifier in a fake message. The attacker must also be precise with their sending power to ensure there is no break from the expected aircraft behaviour in terms of received signal strength (RSS). When impersonating another aircraft that is either currently on a receiver's radar or has been previously, the receiver can measure the RSS of the captured messages. Due to the attacker's position on the ground and potential hardware differences, the measurements of the fake ADS-B messages are unlikely to match patterns of legitimate ones. They could e.g. be more constant over time compared to aircraft covering distances of hundreds of miles in relation to the receiver. With simple statistical methods such as mean, variance and correlation testing, it is possible to catch these breaks in RSS. Going one step further, by developing a free-space path loss model fitted to a receiver and its position, it is possible to estimate the RSS band that a distance claim sent by an aircraft over ADS-B should be in. Using standard hypothesis testing, an IDS can judge the probabilities if a sample of RSS collected from messages stems from legitimate aircraft or from a ground-based attacker.

Feature Selection

We analyse a number of features that can help to detect anomalies in flight data. For example, we calculate the Pearson correlation coefficient ρ between the distance and the RSS, expecting a strong negative relationship as predicted by the free-space path loss model. Any attacker that does not correctly adjust their sending strength in line with the claimed distance of a ghost aircraft to not show such a correlation. More formally, we test the Null Hypothesis H_0 which states that there is no association between the two variables in the population against the Alternative Hypothesis H_A, saying that there is a negative association between the two variables in the population. Similarly, other features such as the autocorrelation coefficient can be used to identify attackers that are stationary or do not adapt accordingly.

Anomaly Detection

We combine the selected features in a machine learning approach to make the intrusion detection more robust. We use MATLAB to create data descriptions of our collected air traffic data. Specifically, we define one-class data sets based on legitimate air traffic data and use various one-class classifiers to create descriptions which include the data. These classifiers decide for new samples if they fit the description or are rejected (i.e., are classified as an anomaly worth investigating). We tested a number of different classifiers, with an acceptance threshold of 99% and a 5-fold cross-validation. While the training sets were drawn from a sample of 16,000 legitimate flights collected with the OpenSky network [2], the test sets had 2% of falsely-injected flights that needed to be detected by the classifier. Fig. 2 shows the results of the comparison between some of the tested classifiers, depending on the number of samples collected per flight. We can see that the Parzen classifier performs best in our sample, showing the lowest false negative (FN) rate, i.e. misclassified attacker flights. At 100 collected messages per flight, it is followed by the Minimax classifier, but also K-Means, MST and KNN achieve a zero FN rate when collecting at least 200 messages, significantly improving on individual features such as hypothesis testing.

3. REFERENCES

[1] M. Schäfer, V. Lenders, and I. Martinovic. Experimental analysis of attacks on next generation air traffic communication. In *Applied Cryptography and Network Security*, number 7954 in LNCS. Springer, June 2013.

[2] M. Schäfer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm. Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research. In *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, Apr. 2014.

[3] M. Strohmeier, V. Lenders, and I. Martinovic. On the Security of the Automatic Dependent Surveillance-Broadcast Protocol. arXiv preprint arXiv:1307.3664, July 2013.