SECURITY IMPLICATIONS AND AUTHENTICATION METHODS OF AUTOMATIC DEPENDENT SURVEILLANCE-BROADCAST IN THE NATIONAL AIRSPACE SYSTEM

BY

BRIAN JASON WRIGHT

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2013

Urbana, Illinois

Adviser:

Associate Professor Yih-Chun Hu

ABSTRACT

With the growing popularity and variety of unmanned aircraft technologies and increasing number of aircraft present in the skies, the Federal Aviation Administration has followed the lead of the European Cascade program in implementing a new system of air traffic management in the United States. Mandating full compliance by all aircraft in busy airspace by 2020, the technology being implemented has many promising applications, including improved collision avoidance, more efficient travel patterns with reduced aircraft spacing, and reduced dependency on error-prone humans. Automatic Dependent Surveillance-Broadcast is a technology that relies on the continuous broadcast of satellite-derived location information from each aircraft.

By considering the possible attack taxonomy of the national airspace system with respect to Automatic Dependent Surveillance-Broadcast (ADS-B), methods of preventing such attacks can be formulated. Three primary communication channels have been identified for deployment of an authentication method. Each channel has unique characteristics that drastically alter the authentication scheme that can be applied. In this thesis, the proposed solutions can help mitigate the security issues identified with the ADS-B system.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS

ADS-B	Automatic Dependent Surveillance-Broadcast
AGL	Above Ground Level
ATC	Air Traffic Control
ATCRBS	Air Traffic Control Radar Beacon System
ATM	Air Traffic Management
DoS	Denial of Service
DSA	Digital Signature Algorithm
FAA	Federal Aviation Administration
FIS-B	Flight Information Service-Broadcast
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
IFR	Instrument Flight Rules
ILS	Instrument Landing System
METAR	Meteorological Aerodrome Report
MSL	Mean Sea Level
NAS	National Airspace System
NextGen	Next Generation Air Transportation System
NOTAM	Notice To Airmen
PKI	Public Key Infrastructure
PSR	Primary Surveillance Radar
SSR	Secondary Surveillance Radar
TAF	Terminal Aerodrome Forecast
TCAS	Traffic Collision Avoidance System
TIS-B	Traffic Information Service-Broadcast
UAV	Unmanned Aerial Vehicle
VFR	Visual Flight Rules
VOR	Very High Frequency (VHF) Omnidirectional Range

CHAPTER 1 INTRODUCTION

Since the Wright brothers' first flights in the early twentieth century, the skies have changed drastically. With over 70,000 flights occurring daily in the United States, the air traffic control (ATC) system is growing beyond its capacity [1]. As a growing global economy increases the demand for lighter, faster, and more efficient aircraft, including increased deployment of unmanned aerial vehicles (UAVs), the skies have become overwhelmed with aircraft. The extremely complex traffic control problem is a pressing topic for a variety of different research areas. Many labor surveys have already determined air traffic controller to be one of the most stressful jobs available [2]. The Federal Aviation Administration (FAA) has decided to move to a system that promises more automatic, efficient, and safer skies [3].

The FAA has chosen to implement the Next Generation Air Transportation System (NextGen) to help mitigate the congestion problem. The NextGen system calls for a satellite-based, as opposed to the current radar-based, tracking system. The system will depend on all participating vehicles to report their current location accurately and honestly. This information is provided to ATC, as well as nearby aircraft. As discussed in this research, the integrity of the system is dependent on the user's honesty.

The major contributions of this work include:

1. Provides a look at the taxonomy of attacks in ADS-B
2. Addresses the lack of authentication among aircraft and ground stations

The work is organized as follows: Chapter 2 provides background and implementation of the technologies currently used by the aviation community. In a more detailed analysis, Chapter 3 describes the NextGen system and Automatic Dependent Surveillance-Broadcast technology. Chapter 4 discusses many of the security concerns, threat vectors and shortcomings with the current and NextGen systems. In Chapter 5, authentication methods are discussed to address the issues discussed in the previous chapter. Chapter 6 emphasizes the previous chapter with simulation design. Chapter 7 discusses the results from simulations. Chapter 8 concludes the proposed research.

CHAPTER 2 BACKGROUND AND CURRENT INFRASTRUCTURE IN THE NATIONAL AIRSPACE SYSTEM

2.1 Overview

In order to understand the motivation behind the FAA's mandate to adopt new technologies in the National Airspace System (NAS), the current system must be overviewed. The current infrastructure of the United States airspace is much like what it was decades ago. Incremental improvements, such as more accurate and abundant radars, weather stations and reporting, and more informative displays for air traffic controllers, have been implemented. The status quo is entirely based on radio voice communication, with the information flow heavily biased in the ATC-to-pilot direction. The pilot is expected to be compliant and provide all information to the ATC personnel.

2.2 Flight Rules

Navigation while in flight can be conducted with many different methods and technologies, determined by the flight rules that are assigned to the particular flight's duties. Visual flight rules (VFR) and instrument flight rules (IFR) are the two schemes that all pilots must follow [4]. VFR, as the name suggests, requires that the pilot is visually unencumbered such that they can navigate via ground references free of thick cloud cover or fog. IFR flight requires more training by the pilot, technologies to assist with navigation, and strict flight patterns and spacing. This mode of flight relies almost exclusively on the sensors and gauges onboard the aircraft, with the trust of ATC to keep surrounding traffic at a safe distance.

2.3 Airspaces

Aircraft are subject to various levels of control in airspaces, designated airspace classes A through G, based on their proximity to airfields and to ground level. Overall, the airspace designation is determined by the traffic in the area, type of flight activity, necessary safety level and the interest of the public [4]. The busier the airspace, such as that close to an international airport (Class B), the more stringently it is controlled. Smaller airports have class C or D

designations. In areas closer to the ground, generally less than 1200 feet above ground level (AGL), with no nearby airports, airspace is uncontrolled (Class G). Class A exists at high altitudes, while class E encompasses the rest of the airspace. A general diagram of these airspaces can be seen in Figure 1.

The areas around airports are classified as controlled airspace. *Controlled airspace* indicates that there is often radar surveillance, control towers, or, at the minimum, established radio communication procedures.

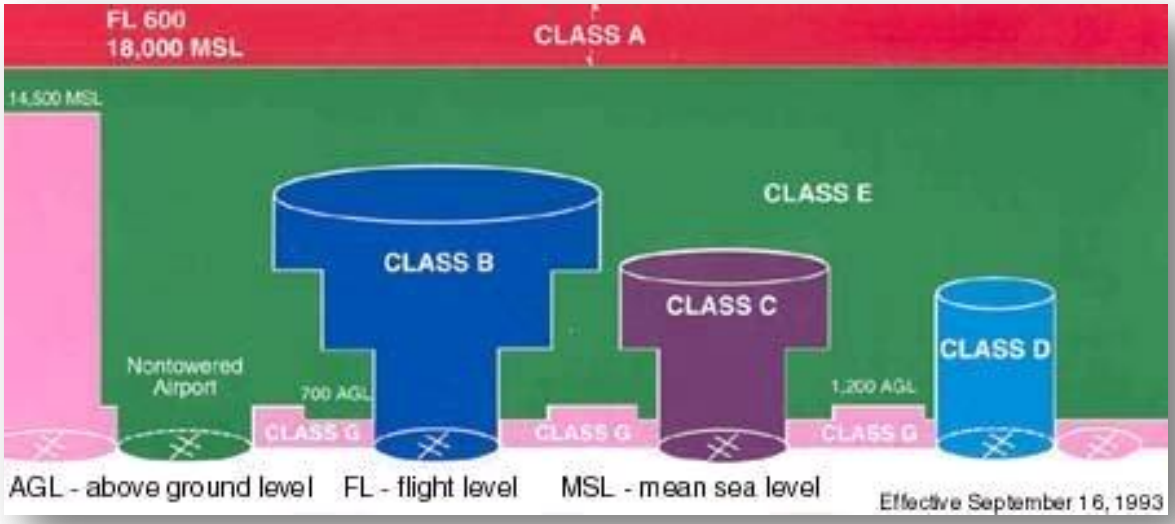

Figure 1: Airspace Classes [5]

2.4 Radio

The most relied upon technology in an aircraft for situational awareness is voice communication over radio. This allows the pilot to be informed about weather and restrictions in the area as well as have situational awareness of surrounding aircraft and hazards. In controlled airspaces of classes A, B, C and D, two-way radio equipment and communication is required with ATC for clearances and sequencing. Class E requires two-way communication for IFR flights.

Radio communication is established over frequencies based on the current position of the aircraft. Different ATC personnel operate these frequencies for ground control, tower (takeoff/landing clearances), en route, and approach/departure. As the pilot passes between the

boundaries for each of these control spaces, the pilot is instructed to contact the next controller on the published frequency.

2.5 Radars

Radar uses a directional radio signal to detect objects in the sky via the round trip time of the emitted signal. Based on the time of the signal's return, the device can determine the distance of the object from the radar device in the given direction that the signal was broadcast. This information is shown on a screen for the ATC personnel for awareness and to disseminate via radio communication to the pilots for improved situational awareness and traffic control. Primary surveillance radar (PSR) uses a passive signal reflection technique to detect objects in the signal path. Secondary surveillance radar uses a three-pulse interrogation to illicit responses from transponders onboard aircraft. The first and third pulses are a timing sequence that is used to interpret which mode is being interrogated. A second pulse is for side lobe suppression. It is transmitted from an omnidirectional antenna to eliminate leaked interrogation signals (called side lobes) by comparing intensity of the first, second, and third signals. This technique reduces the ground clutter, or reflections, present in the primary surveillance radar by affirming a vehicle of interest corresponds to the detected object on the radar. Radars are a mechanical technology, and rotate at a fixed rate that limits the rate of interrogation in a given direction.

2.6 Transponder

What began as Identification Friendly or Foe (IFF) equipment in World War II is now an integral part of air traffic control. Under the technological branding of air traffic control radio beacon system (ATCRBS), and after many technological advances since its inception, the aviation transponder is available in a few levels of intricacy. Beacon transponder technologies began as military protocols but have become the basic surveillance technology in commercial and general aviation. There are currently encrypted versions of transponder modes available for use by the military. Through an interrogation-reply format, a ground station can determine the surrounding aircraft to provide verification to the radar system. There are three basic modes of transponders. Interrogation from the radar system is on the 1030 MHz frequency, while the reply is on 1090 MHz [6].

There are three transponder modes, modes A, C, and S, in use by commercial and general aviation. Beginning with mode A, a four-digit code is broadcast from the aircraft when

interrogated by a radar system. The four digit number is assigned by the air traffic controller upon established communication and dialed in by the pilot in the cockpit. This allows the controller to confirm the location of the aircraft in communication on the radar. Mode C is very much the same as mode A, with the same interrogation schema and broadcast methods. The differentiating factor is that mode C broadcasts the current barometric altitude of the aircraft, allowing ATC to locate the aircraft on the third dimension. Mode C transponders are backwards-compatible and can also respond to mode A interrogations, and the ground station interrogates on both modes.

The interrogation signal from the ground station is in the form of a three-pulse output. The time between the second and third pulses determines the mode (A or C). A shorter delay of 8 µs is a mode A interrogation and a pulse delay of 21 µs is a Mode C interrogation as shown in Figure 2. Mode S transponders continue to listen after the mode A or C interrogation for an address specification, as shown in Figure 3. Mode A and C transponders ignore this extra information and respond with the corresponding information requested.

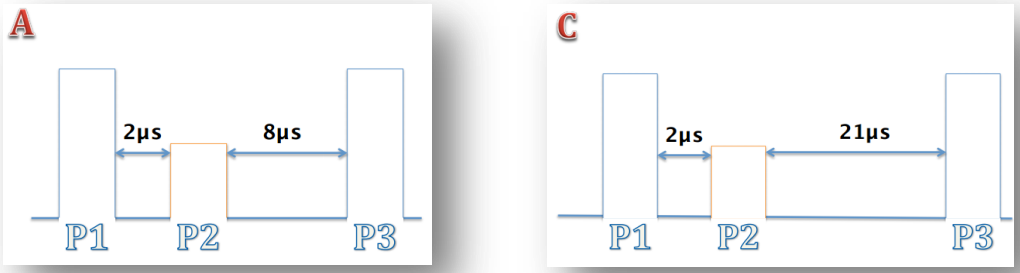

Figure 2: Mode A (Left) and Mode C (Right) Interrogation

Developed by MIT Lincoln Laboratory in the 1970s, Mode S, or select-mode, addressed the congestion by Mode A and C replies, and provides the ability to overlay a collision avoidance system on top of the transponder technology [7]. Mode S is the interrogation mode that removes the mass reply from all aircraft in the area of the interrogator by allowing probing of a single aircraft based on its distinct address. The specific identifier of the transponder that the interrogator desires a response from is appended after the interrogation mode type in the interrogation preamble.

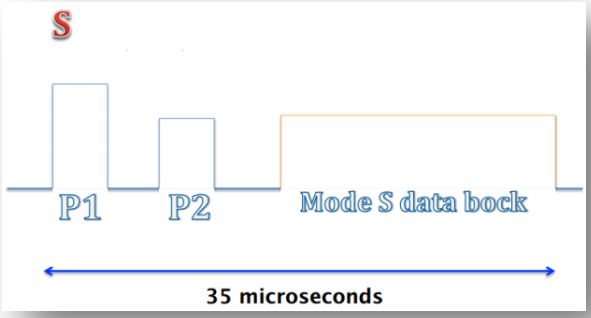

Figure 3: Mode S Interrogation

2.7 Collision Avoidance

Using some of the same transponders as the secondary surveillance radar (SSR) system, a traffic collision avoidance system (TCAS) provides rudimentary collision avoidance. TCAS uses the mode S transponder to interrogate for nearby aircraft with modes C and S. Mode A transponders do not interact with this system.

The proximity of an intruding aircraft is determined via round trip time and information received from the intruder. Via a two ellipsoid-zone setup around an equipped aircraft, the TCAS system can negotiate with other equipped aircraft for maximum separation by issuing traffic and resolution advisories (TAs and RAs). The TA is an alert of nearby aircraft with instruction to visually maintain distance from the intruder. The RA is issued in imminent danger of collision, and is often issued in the form of vertical course augmentations, sending one aircraft to a climb and the opposing aircraft to descend.

As shown in Figure 4, an intruder's proximity to the aircraft determines the level of the aforementioned alerts. Note that the larger region covers approximately 20 nautical miles ahead of the aircraft, giving under a 30 second warning for most commercial aircraft travelling over 400 knots.

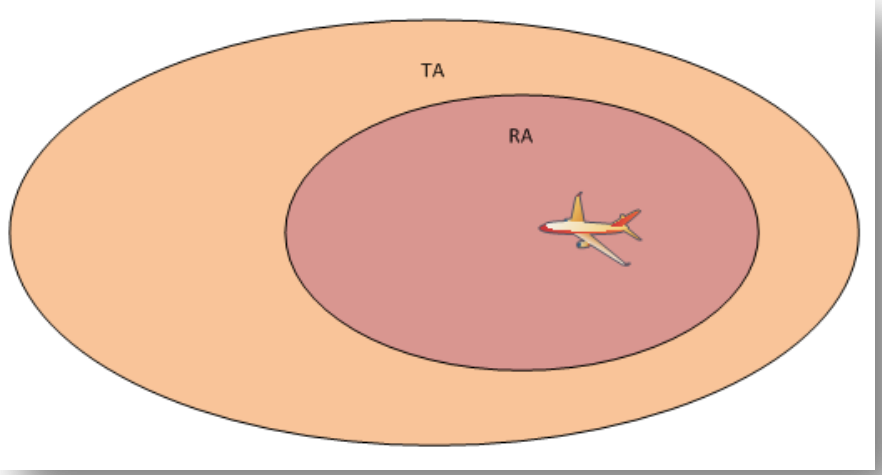

Figure 4: TCAS Advisory Regions

Through some evolution of procedures, it is now required that pilots follow the issued RAs over the instruction of the ATC. One incident leading to this resolution was a mid-air collision in which the pilot of a commercial aircraft followed ATC directives over the RA issued and took down a smaller general aviation aircraft. The ATC, while situationally aware from the SSR in a given area, is not informed of RAs issued in the cockpit.

The currently deployed TCAS technology is TCAS II version 7.0/7.1. While the revisions of TCAS II are mere updates to patch issues found with the technology, subsequent versions of TCAS have been developed, but abandoned. Most notably, TCAS III was developed to address some shortcomings of TCAS II, but abandoned. TCAS III shows the feasibility of detecting a relative horizontal position between two conflicting aircraft using the existing antennas from the TCAS unit. The technology can determine horizontal separation via timing difference, and thus differences in bearing.

CHAPTER 3 NEXT GENERATION AIR TRANSPORTATION SYSTEM AND AUTOMATIC DEPENDENT SURVEILLANCE-BROADCAST

3.1 Overview

The next generation air transportation (NextGen) system is the solution proposed by the FAA to address issues with efficiency, safety, and congestion. A major part of the NextGen system is the surveillance method of ADS-B.

3.2 Next Generation Air Transportation System

The NextGen system, developed incrementally for the past decade, is a suite of new technologies aimed to bring more information into the tower and cockpit for better control and greater safety. With the majority of the physical systems now in place, many software and algorithmic technologies can be implemented. Most notably, these technologies enable moving towards an ever more automated flight experience, opening the doors for unmanned aerial vehicles (UAV) to be commonplace in the NAS. The system has been deployed in some areas of the United States, as well as in parts of Europe and Australia.

A major motivator for the FAA's push for the NextGen upgrades is infrastructure cost savings, which will be realized when the majority of systems such as SSR are taken offline. While these projected dates have not been set, the notion of removing SSR poses many security-related questions, which are discussed further in Chapter 4.

Included in the NextGen system are equipment upgrades for both ground infrastructure and onboard avionics. Ground stations, spaced at intervals much greater than the existing radar systems, are deployed to concatenate and broadcast the announcements of aircraft in the area and provide services to be later realized in the *ADS-B In* technologies. *ADS-B In* is the data stream for in-cockpit information.

A hindrance for the general aviation community has been the upgrade costs for the new avionics required. As the NextGen technology progresses, the cost for this equipment will drop to what the FAA hopes to be a reasonable level for aircraft owners after the initial product development cost has been compensated.

3.3 Automatic Dependent Surveillance-Broadcast

Automatic Dependent Surveillance-Broadcast (ADS-B) is a technology built upon the transponder concept of the identify friend or foe (IFF) developed during World War II. The ADS-B technology is simply a protocol which can sit on a number of different carrier frequencies, but the two chosen for implementation are the 1090 MHz (of previous transponder technologies) and the 978 MHz bands. Shown in Figure 5 are the separation of frequencies and protocols on the ADS-B system. The 1090 MHz frequency, previously and currently used for Mode S transponders, is the main channel for ADS-B, but to address some congestion problems, the FAA is also supporting the use of 978 MHz Universal Access Transceiver (UAT). Ground stations have the capability to rebroadcast the messages on each frequency over the other frequency.

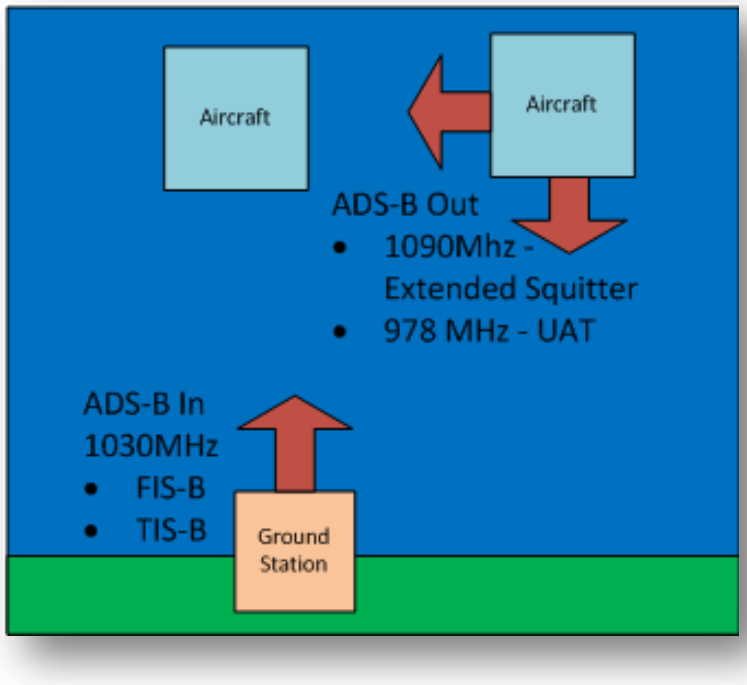

Figure 5: ADS-B Data Flow

The final decision by the FAA is for the 1090 MHz Extended Squitter (ES) to be used by commercial and high-powered aircraft, while the 978 MHz UAT channel will be used for general aviation [8]. It is the responsibility of the ADS-B ground stations to translate between the two frequencies, and provide cross-platform data flow and awareness.

The main ADS-B system is based on the 1090 MHz, with an extended squitter message type support. The ES message contains the pertinent information for ADS-B: position, velocity, and time. ADS-B is structured much like the Mode S transponder of the ATCRBS with the major difference that ADS-B is automatically broadcast once every second. The message is restricted to 144/272 bits on UAT and 56/112 bits on 1090ES [9]. The protocol operates at an update frequency of close to five seconds faster than ground-based radar systems. This allows for the system to absorb dropped messages and still operate with better performance.

The aforementioned ADS-B, while generally coined ADS-B, is known by *ADS-B Out* in FAA documentation. *ADS-B In* is the non-required, still evolving, largest benefit of the NextGen system to a pilot. *ADS-B In* provides free data services to the pilot displayed in the cockpit.

ADS-B In sits on the 1030 MHz frequency and currently comprises *ADS-B Out*, FIS-B and TIS-B, two information services broadcast from ground stations. Flight Information Service-Broadcast (FIS-B) consists of routine weather reports and forecasts such as METARs and TAFs as well as airspace restrictions and notices such as NOTAMs. Traffic Information Service-Broadcast (TIS-B) consists of detected traffic in the area via *ADS-B Out* signaling and secondary surveillance radar (where it still exists) for visualization in the cockpit for situational awareness [10].

In the FAA's final ruling in 2010, all aircraft within airspace classes A, B, and C, as well as class E above 2500 feet above ground level (AGL) or 10,000 feet mean seal level (MSL), are required to be equipped with an *ADS-B Out* device, with a 1090 MHz ES device required above 18000 MSL [8]. These requirements can be seen in Figure 6. The technology is also being applied while aircraft are grounded at large airports to prevent collisions and runway conflicts.

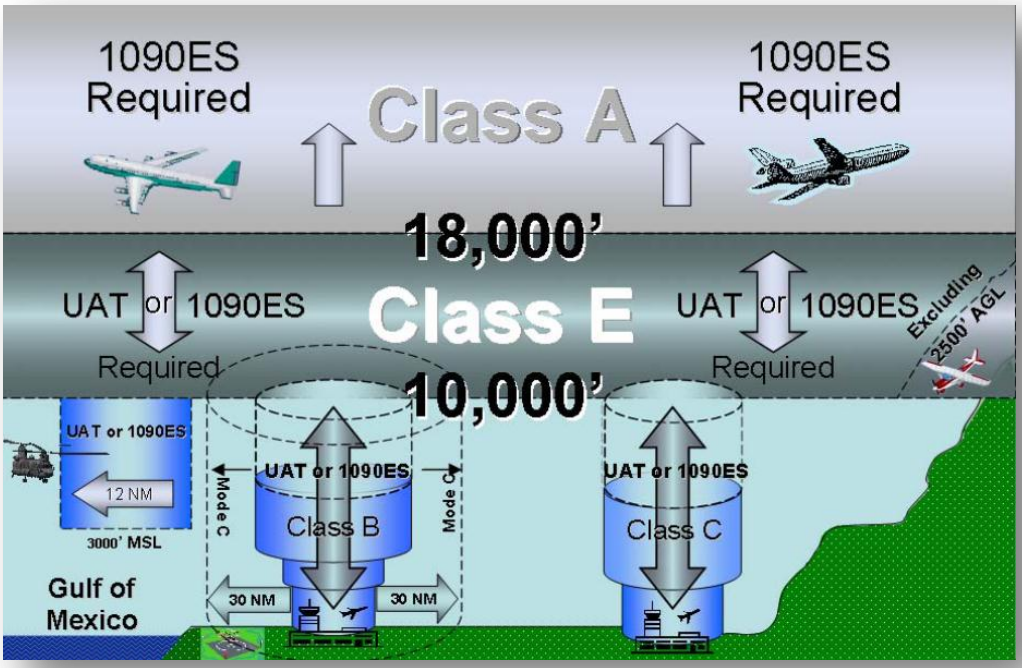

Figure 6: ADS-B Platform Requirements [11]

CHAPTER 4 SECURITY AND ATTACK VECTORS

4.1 Overview

As aviation becomes a bigger part of the infrastructure of the country and global economy, it becomes a more enticing target of attack. While attacks can target many components of the system, this thesis will focus on those targeting the ADS-B, as it is the basis for the success of much of the NextGen system. The FAA's stance on the security of ADS-B is that there are no concerns with the technology, but the agency maintains the results of a security study in a classified report. The FAA has indicated that a successful attack on the system would be too difficult [12].

Security problems with ADS-B include the lack of authentication of the aircraft in the skies. Currently, the air traffic controller will establish a verbal handshake with each pilot in the area, determining who they are, what kind of aircraft (and thus the aircraft's capabilities), where they are (radar based location with a broadcast altitude location) and their intentions (where they are heading). While many of these attributes, by themselves, provide a small risk to the system, the combination of this information can provide an easily compromised system in which money, time, or even lives can be lost. The transition from the current ground-based surveillance system to a satellite-based ADS-B system should be done with the utmost care and consideration of the opportunities such a system would provide for malicious adversaries.

4.2 Automatic Dependent Surveillance-Broadcast

ADS-B Out is a simple broadcast protocol emanating from all aircraft on unencrypted channels. Regardless of anyone listening, an aircraft persistently announces its position and identity every second. Anyone with a receiver or transmitter, in any location within the radio transmission distance from an aircraft or ground station, can listen, or even participate in the system. Since the system is unauthenticated and uses plain text, it heavily relies on the honesty of the participants. The system can therefore be compromised in a number of ways and for a variety of intentions.

It has been shown that commercially available hardware, e.g. FPGAs and microcontrollers, can be utilized to build ADS-B receivers. There are a number of commercial websites that offer ADS-B receiver kits for a considerably lower price than that of the commercially available and

FAA mandated avionics. While these devices only cover the information gathering attacks, a functional transmitter can be built with commercially available hardware, and can even be built more powerful than commercial ADS-B products. Programming the transmitter and receiver can be accomplished with published ADS-B documentation and some observation of the current overhead implementation. A recent open-source system known as GNU Radio has been carving the way for a software-defined radio, allowing modulation and signal processing to be accomplished in software, without the need for specifically constructed hardware for a given transmission type. GNU Radio is yet another tool at the hands of an adversary to the ADS-B system.

With ADS-B, the data broadcast contains who and exactly where they are located at any given period of time. Many systems that will be migrated depend on the ADS-B information received onboard an aircraft. Collision avoidance systems such as TCAS or even a pilot flying VFR will depend on this information for situational awareness. A single aircraft, either maliciously or via equipment failure, failing to report a correct position can result in a collision in a system such as ADS-B. Equipment failure will not be considered here, as it is up to the manufacturer to either provide fault tolerance of the equipment or to provide warning to the pilot in the instance of a loss of a loopback signal or mismatched calculation of groundspeed and GPS position change. Attacks can be classified by a few characteristics: intent, target, system, and location.

There are many motives that could drive an adversary to attack the NextGen system through the use of ADS-B. The more obvious reason would be to maliciously cause harm to people or property. Using ADS-B alone, a terrorist could route one or many aircraft to desired locations, create chaos in the sky, and even cause mid-air collisions or fatalities on the ground. The terrorist might want to harm a particular person onboard the aircraft, use the aircraft as a guided missile, or just create havoc for extortion.

The target of the attack can be one or multiple aircraft, the ground stations, or even a person or structure on the ground. As demonstrated on 9/11, a strategically downed aircraft can cost many lives. With some reconnaissance, an attacker can determine a particular aircraft overhead, possibly targeting a passenger, company or just an aircraft known to be fully fueled. A multi-aircraft attack is much more complicated, but can be accomplished with a well programed

controller and artificial intelligence algorithms. If the ground station is attacked, all aircraft in the area will be affected.

In aviation, there are many intertwined and interdependent systems that provide services such as navigation, congestion control, and weather notifications, some being more critical than others. Attacking the weather broadcast would be an inconvenience, but the information could be disseminated via other channels such as verbally over the radio or through internet access from satellite or cellular towers. Systems often do not maintain the same level of criticality over time due to weather and time of day. For instance, failed radio communication over VFR traffic can be mitigated by use of a directed light beam from the tower. This procedure would not necessarily work during IFR conditions (poor visibility, precipitation or fog), and thus a radio jamming attack during these conditions would be more critical. An attack on *ADS-B Out* would be most successful during IFR conditions, as there would be no visual confirmation of the surrounding aircraft in the sky. An attack on *ADS-B In* would be a simpler multi-aircraft attack, as it would not require the game techniques of predicting a pilot's reactions, but would require enough power to override the ground station. The authentication of aircraft in the skies to eliminate the ability to compromise navigation and air traffic control in poor conditions is critical.

The location of an attack adds another level of complexity. If an attack is launched in a controlled airspace around an airport, likely, there is a ground radar system that can verify that there is an attack in progress. This information can be announced to surrounding aircraft and the problem can be mitigated. However, in more rural areas, these already scarce radar systems are slated to be removed from service after being replaced by ADS-B ground stations. Without the verification of radar, the rural arena is ideal for having specific targets. Given the direct path flight and increased efficiency rewards of the NextGen system, avoiding VOR waypoints, usually located in congested areas, would be possible, increasing the likelihood of a given target being present outside of radar surveillance. Another consideration with the location of attacks is a single or distributed location of the attacker. A distributed attack would be more effective, having a greater coverage area.

With a simply constructed transmitter and trivially programed controller, the terrorist can project false aircraft into the sky, visible to aircraft and ground stations in the area. In the interest of the pilot and ATC employees, it would be natural, following the precautionary principle, to

avoid anything that appears on the ADS-B user interface. This scenario develops the idea of puppet control of the targeted aircraft and the notion that the attacker can indirectly steer the aircraft to a location of choice. Anything currently operating with mode S or depending on ADS-B is vulnerable to this attack.

An example scenario for target-based attacks involves knowing the specific party of the target aircraft. If it is an attack on a specific person or a company, an attacker can obtain the specific ID from the aircraft via FAA databases, or in the instance of military or special persons aircraft, a random ID that does not correspond to the database properly could indicate that the aircraft is of importance, identifying the target for the attacker. With this information, the attacker has the power to do many things, from simple tracking and corporate espionage to maliciously grounding the aircraft [13].

In the instance of a wide-scale attack, the attacker may have a goal of causing mass chaos with multiple targets. This attack could be launched by flooding the skies with ghost aircraft, similar to a denial of service attack in a computer network. As discussed earlier, ADS-B will enable reduced separation between aircraft, dependent on the speed and class of airspace the flights are traversing. With a system that can be compromised in the discussed vectors, the reduced separation will further reduce the margin of acceptable error to prevent collisions.

4.3 Global Positioning System Spoofing and Jamming

While not the main focus of this thesis, it is imperative that an attack on the fundamental protocol, that gives ADS-B its satellite-based positioning, be considered. The global navigation satellite system (GNSS), also known as global positioning system (GPS), is based on a network of satellite-based transmitters and a localized receiver. The receiver uses timing differences of the satellites in view to determine its location in space for translation to latitude, longitude and altitude.

There are two angles of attack on the GNSS. One is a standard denial of service (DOS) attack. By overpowering the weak signals from the satellites with a more powerful transmitter near the target, the victim relying on the positioning system will no longer have a position fix. This attack can be focused on an individual, or a group, based on the resources and power available to the attacker.

A second attack, aimed directly at one aircraft, has been demonstrated on UAVs both in research at the University of Texas at Austin and, allegedly, in the real world on a United States Central Intelligence Agency drone in 2011 over Iran [14]. This attack involves the adversary synchronizing with the GNSS while slowly overpowering the specific satellites the target receiver is polling. By overpowering the infrastructure, the attacker can slowly augment the reported location of the vehicle such that it is somewhere other than where it is expected to be. On automated vehicles, an incorrect location can initiate procedures, such as landing, or corrective actions that can be leveraged to cause the aircraft to land unintentionally. On a manned flight, a pilot could become aware of this course variation via a visual discrepancy or by the current, but dated, VOR (VHF omnidirectional range) or ILS (instrument landing system) technologies expected during navigational waypoints and during landing.

CHAPTER 5 AUTHENTICATION

5.1 Overview

Authentication is the method of ensuring that a message or assertion from an entity did, in fact, come from the entity that it claims to have come from. As indicated in Chapter 4, there are several major attack vectors within the ADS-B system. This chapter will discuss authentication methods for the critical players in the system. The major scenarios considered are:

- Authenticity of an individual aircraft with *ADS-B Out* to the ground station
- Authenticity of an individual aircraft with *ADS-B Out* to peer aircraft
- Authenticity of *ADS-B In* ground station broadcasts of FIS-B

5.2 *ADS-B Out* with Ground Station (Downlink)

ADS-B Out is the pinnacle of the satellite-based surveillance in the NextGen system. With the attack potentials discussed in Chapter 4, *ADS-B Out* must be secured in a manner that does not inhibit the data transfer to an aircraft's peers and local ground stations. When considering the ground station, they do not operate in a resource-constrained environment in terms of power, computation, or physical space. Currently, it has PSR and SSR to verify an aircraft's claimed location in an ADS-B message. As stated previously, the FAA has decided it would like to see the cost-benefit analysis of retiring parts or all of these radar systems. With the absence of SSR, the ADS-B information can continue to be verified by a corresponding match on the PSR. However, due to the removal of the rotating antenna, the ability to verify that the ADS-B report came from the direction reported is lost. While this removes some granularity of authentication, an attacker would still be unlikely to convince a ground station that he is in the sky when he remains transparent to the radar. Without knowledge of what is represented on the radar, the attacker is only able to authenticate a fake object (e.g., a weather balloon) with the ground station for the period of time the physical object remains synchronized with the position in the broadcasted messages. This form of attack is therefore less efficient than an attacker physically flying an aircraft.

In the instance that an attacker floods the sky with ghost aircraft in a manner similar to that of a Denial of Service (DoS) attack, aircraft in the surrounding area would receive these

messages and accept the data at face value. While not all aircraft are required to have *ADS-B In* at the time of the FAA's mandate for *ADS-B Out*, there exists the ability for a ground station to send the verified traffic in the area over an *ADS-B In*/TIS-B message for comparison by an aircraft. In a controlled airspace with a ground station, in the currently defined ADS-B system, there is no procedure that defines how the system on-board an aircraft will complete this verification. It should be known that by adding an on-board comparison mechanism and trusting the *ADS-B In* information with the proposed authentication in Section 5.4 over its own received data within the well-defined map area, an aircraft can quickly dismiss any ghost aircraft during an attack.

A single ghost infiltration, if executed at a critical time and place, could cripple the area until the attack can be stopped or verifiably ignored. Consider the case of a ghost aircraft injected in the glide slope of landing aircraft, and how it would parallel a snowstorm's effect on the larger NAS landscape with cancellations, delays, and redirects. The severity of the attack depends not on the part of the system attacked, but on the ingenuity of the attacker and their ability to adapt.

With *ADS-B Out* in an airspace controlled by a ground station, it is proposed that the PSR system remain in place for verification from aircraft to the ground station. Currently, in the United States, PSR covers the majority of the nation, with over 510 government owned radar systems. The airspace coverage at 5000 feet AGL can be seen in Figure 7 [15]. This existing coverage provides a sufficient system by which to verify aircraft in the NAS. SSR is severely less abundant; making the push for the NextGen system allows a verified surveillance system to match or exceed that of the current SSR/PSR system.

Figure 7: PSR Coverage at 5000 Feet AGL [15]

With PSR coverage over the majority of the NAS, an *ADS-B Out* system verified by PSR can sufficiently ensure that many of the previously discussed attacks do not have the ability to subvert into a controlling position in the system. Ideally, full coverage of PSA can circumvent the peer-to-peer authentication problem further discussed in Section 5.3. One downside to this approach is that PSA is expensive, and more unlikely now than ever to cover remote areas, i.e. over the Gulf of Mexico. ADS-B ground stations, however, have already been installed on oil drilling rigs for surveillance coverage not previously realized. An authentication method without using PSA would require directional antennas for the ability to determine the direction from which the *ADS-B Out* message is coming with respect to the ADS-B ground station. Likely seen as a costly addition to the ADS-B ground stations, directional antennas in strategic locations can provide location verification in areas below radar envelopes, such as in glideslopes around runways, that are otherwise uncovered areas in the PSA verification scheme.

5.3 *ADS-B Out* with Peers (Cross-link)

A vulnerable part of the ADS-B system exists when an aircraft is outside of an airspace monitored by a radar and ground station. These areas, while not abundant, exist outside of the NAS, above oceans and large bodies of water, and in geographically difficult areas such as the mountains in Alaska. In addition to these incompatible PSA areas, the FAA is pushing for the removal of all radar systems to fully realize the economies of ADS-B. The benefit of having the

radar verification is lost for detecting ghost aircraft by an attacker. For the NextGen system to operate safely and efficiently, it is imperative that an aircraft be able to trust the information it is receiving about its surroundings. Of the three areas for authentication in ADS-B, peer-to-peer authentication is the most difficult problem to solve.

While in the veil of a controlled airspace with ground station, the peer-to-peer location reporting can be compared to the periodic *ADS-B In* updates, including the PSR verified traffic (a combination between downlink and uplink authentication). The real challenge presents itself outside of ground station surveillance. How does an aircraft ensure that a peer aircraft is precisely at the location it is reporting?

TCAS technologies present a method of rough location detection that is already on-board all commercial aircraft. TCAS III, now abandoned for continued use of TCAS II, introduces the ability to also determine horizontal separation with no interaction to the ground. In [16], the authors state that the technology is only viable if an estimated separation between the aircraft is available. This information is available with the use of ADS-B, trusting the peer aircraft's position just long enough to calculate the horizontal separation. The horizontal separation, with the initial TCAS distance, can be used, in turn, to verify the location of the aircraft relative to the other.

In TCAS III, the horizontal separation is calculated via a simple formula with the bearing rate, magnitude of the relative velocities between the two aircraft, and the range between the two aircraft [16]. The value needed from ADS-B is the bearing rate between the two aircraft. As the formula is calculated over time, the estimated horizontal distance should match with the messages received by the unauthenticated aircraft. It is recommended that this technology be reexamined for use with ADS-B for air-to-air authentication.

5.4 *ADS-B In* (Uplink)

The severity of an *ADS-B In* attack outweighs that of *ADS-B Out* in the sense that with one broadcast, an attacker can compromise the entire broadcast area by inserting as many aircraft into the airspace map as necessary and place that information in the cockpit of every aircraft overhead with a single broadcast.

In [17], a symmetric encryption scheme is discussed for authenticity of the ground station's *ADS-B In* broadcasts. The author considered the necessity of maintaining an authentic

data link and the ease with which this need could be met by securing it using cryptography, but missed the inconvenience and flaw in a shared key system. The proposed data link authenticated encryption algorithm requires that a pilot obtain and enter a key into the system for encrypting and decrypting functions. This key, likely given in person to anyone with a pilot's license or over the clear radio channel, can be easily obtained by an attacker to be used to compromise the system.

With minimal computing power, a ground station can be verified by implementing a public key infrastructure (PKI). An asymmetric encryption scheme would allow each ground station to maintain a public and private key. The private key is held secret, such that only the ground station, the authenticated source, can encrypt data. The public key, much like the published radio frequency for ATC communications, can be loaded into the ADS-B unit permanently. This public key allows anyone to read the information provided by the ground station, knowing that the message, indeed, came from that source.

One major issue with the above approach is the required computing power to encrypt and decrypt and the overhead to transmit an encrypted message. These concerns can be mitigated by using a digital signature. By taking a hash over the entire message, and signing the hash by encryption with a private key, the message payload can remain in plain text for quick dissemination by all peers, but can be authenticated at leisure by the receiving party with minimal transmission overhead [18]. A public key cryptographic signature method such as Digital Signature Algorithm (DSA) can be appended to the ADS-B message. This algorithm generates a couple constrained random numbers and uses them with the hash of the message to create a two-value signature. When the message is received by the aircraft, it calculates, based on one of the signature values, hash of the message and public key, a value to compare against the second signature value. The hash provides message integrity, while the encryption of the hash provides the authenticity. A properly decoded signature allows an aircraft to authenticate a ground station.

A concern with encrypting ADS-B messages is the low bandwidth available for the system. Cyclic redundancy check bits and forward error correction are built into the protocol to provide higher availability, but without message collision mitigation, the system inherently relies on the rebroadcast of messages in poor reception, and thus error correction is not nearly as

pertinent as simple error detection. The check bits can therefore be replaced by the digital signature just discussed to achieve the same results.

Currently, *ADS-B Out* is packaged in the same manner as the *ADS-B In* message format, with 24 bits of cyclic redundancy check data appended to the end of the message. Using these 24 bits for a signature rather than simple error detection can provide a solution to the unauthenticated message issue. While current digital signature schemes provide a low overhead for larger packet sizes on less congested network media, the overhead exceeds the restrictions placed on the current ADS-B protocol. Without either a modification to the protocol or an adaptation of the digital signature algorithm, an authentication method would not be possible.

To achieve a signature scheme within the 24-bit parameter of the ADS-B protocol, a reduced DSA scheme can be created or the data field must be reduced in length. The algorithm, having a two integer signature, a four integer public key and one integer private keys, can be reverse engineered to determine the lengths of the keys based on a 24-bit signature. While this approach will not provide the same assurance that DSA provides for messages without such tight data restrictions, it will provide more authenticity to the *ADS-B In* transmitted data of the status quo.

With a restriction of 24 bits placed on the two signature values of r and s, the lengths can be assigned as 6 and 14 bits, respectively. Adjusting the length of the signature changes the length of the keys from a length of 159 bits down to a 15-bit random prime number. This also reduces the domain of possible signature values, and thus increases the chance for hash collisions. Collisions occur when the signature matches for two distinct messages, and give an attacker the ability to determine the private key.

Another alternative to a reduced DSA digital signature on the *ADS-B In* protocol is an algorithm called the Schnorr signature [19]. With an unconstrained length of random number seed, the algorithm can produce more concise keys and signatures. It produces a shorter signature of a length of 40 bits and provides a low probability of collisions. This signature would require the length of the data fields in the message to be shortened by 16 bits, reducing it to 72 bits. This signature, with a shorter key length than modern standard signature algorithms, could likely be cracked by brute force. Especially with mobile technologies and current general-purpose graphic processing units (GPGPU) that have become increasingly beneficial in cryptanalysis.

The Schnorr signature is more robust than the original DSA, having not used the SHA-1 hash function, which had a known attack as a result of predictable collisions. The current standard of DSA now uses a longer, significantly changed cryptographic hash function SHA-2, which gives DSA the edge over Schnorr [20]. The Schnorr signature algorithm is an older algorithm and because of its age, it is inherently less secure than more recent advances in cryptographic signatures such as the recent standards of RSA and DSA as well as elliptical curve DSA (ECDSA). Without a signature algorithm that is secure, the purpose of signing ADS-B messages is lost with the trust of each message being from a potential attacker.

With a protocol that is consistently and persistently broadcasting information, there is the ability to sign a message and distribute the signature over multiple messages. The receiver would then buffer messages and their respective signature piece in order to compile the pieces back into the signature to verify the first message. By distributing the signature, a more secure algorithm and longer key size can be utilized for the signature generation. Signature schemes such as DSA or ECDSA can then be used.

Two options arise with this schema: sign the first message of the block of messages it takes to send the full signature and sign the number of concatenated messages necessary to fit the full signature pieces needed for authentication. In the first option, the messages that are in between signed messages - the messages carrying only the signature piece from the signed message - are unauthenticated, and can be compromised by overriding the ground station signal strength during the data block only, leaving the signature piece untouched. The second option requires that enough messages for signature transmission are known prior to the first message being signed. This could introduce a delay on the messages.

Overall, there are existing methods to provide a secure signature algorithm to a publically broadcasted message, but due to the constraints of the ADS-B message format and length, these signatures have a larger overhead than the protocol can support. For a truly secure communication channel, the message length should provide the necessary bits to provide a secure signature, either by shortening the data, adding to the length of the message overhead, creating a less secure algorithm and allowing the public signatures to be updated at an interval such that hash collisions are reduced or span a modern digital signature across multiple messages. In each of these methods to secure *ADS-B In*, however, a weakness is introduced into the signature.

CHAPTER 6 DESIGN, RESULTS, AND FUTURE WORK

6.1 Considerations

For all parts of the ADS-B protocol, bandwidth is a large hindering factor for conventional methods of authentication and verification. Any proposed solutions must ensure that the added authentication information does not reduce the payload available to the ATC system.

Simulation can be useful in this research for verification that authentication methods do not interfere with the current operation of the system. But without current equipment and detailed data sheets of the implemented infrastructure for the ADS-B equipment, any simulations would be invalid.

6.2 Results

The taxonomy of attacks in the NextGen ADS-B system has been thoroughly examined and areas of concern indicated. For each of the three main concern areas with authentication in the national airspace, a feasible technology or algorithm has been suggested. The downlink from an aircraft can be verified in software using the existing PSA infrastructure; otherwise, new hardware would be required to provide directional support. The crosslink will require a new system in place to determine relative locations of incoming signals for verification. Leveraging some technologies of the abandoned development of some collision avoidance systems could provide a trusted area around the aircraft. The uplink can be authenticated by using digital signatures on a public key system.

6.3 Future Work

The FAA's NextGen development plan does not clearly state the agency's intentions for maintaining equipment rendered obsolete by ADS-B technologies. Further investigation into which older technologies are essential for backup and verification of the ATC is essential for FAA cost savings during the transition to the NextGen system.

As the technology and equipment for ADS-B continue to be designed, it is unclear what levels of computation will exceed the specifications for which the equipment was originally designed. Using a hybrid simulation technique, authentication methods can be verified with minimal hardware and software models. Hybrid simulation works by pulling real-time data from

a small subset of hardware and using computer models to scale up the simulation. With access to the newest revisions of hardware for the ADS-B protocols, the timing of the system can be verified.

The field of cryptography is constantly evolving. Responding to the increase in adversarial computer power used to crack older and less complicated algorithms, the life of any one given algorithm is short-lived. Even algorithms like RSA and DSA have changed since their releases, incorporating, for instance, a hash function less likely to produce a collision. Emphasis should thus be placed on discovering a secure signature algorithm that meets the special requirements of the *ADS-B In* protocol.

CHAPTER 7 CONCLUSIONS

In conclusion, the NextGen system brings many promising technologies to the table in aviation. In the push for decreasing the reliance on humans-in-the-loop and more efficiency in increasingly crowded skies, the required authenticity of the underlying technology of the NextGen system is grossly overlooked.

As enumerated in the attack taxonomies presented, the ADS-B system must be augmented to provide the trust that is needed to ensure the safety and integrity of the NextGen system.

By maintaining the PSR system already deployed around the United States, verification of announced aircraft location can be completely achieved with few or no resources beyond those required under the current SSR system. With the research presented for TCAS III, the technology holds promise for supplementing that of ADS-B. The ability to authenticate the location of another aircraft in the sky without the use of an on-board radar system is a tremendous step. The implementation of a public key digital signature to any messages that originate from ground stations will prevent attackers from impersonating a ground station.

REFERENCES

[1] National Air Traffic Controllers Association, “ATC Air Traffic Controller,” 2011.

[2] A. Altman and T. Sharples, “Air Traffic Controller Sounds Alarm,” *Time*, 2008.

[3] Federal Aviation Administration, “NextGen Implementation Plan,” 2012.

[4] Federal Aviation Administration, “Aeronautical Information Manual,” 2012.

[5] Department of Defense, “Unmanned Aircraft System Sense and Avoid Technology Development Roadmap,” 2010.

[6] E. Chang, R. Hu, D. Lai, R. Li, Q. Scott, and T. Tyan, “The Story of Mode S,” 2000.

[7] T. Bailey, “All About Mode S Transponders,” *Avionics News*, no. April, pp. 44–49, 2005.

[8] Federal Aviation Administration, “14 CFR Part 91 Automatic Dependent Surveillance-Broadcast Out Performance Requirements To Support Air Traffic Control Service; Final Rule,” *Federal Register*, vol. 75, no. 103, pp. 1–37, 2010.

[9] E. Valovage, “Enhanced ADS-B Research,” *Aerospace and Electronic Systems Magazine*, 2007.

[10] Federal Aviation Administration, “New Technology - ADS-B, TIS-B, and FIS-B,” *Air Traffic bulletin*, no. 2005–3, 2005.

[11] Federal Aviation Administration, “Automatic Dependent Surveillance-Broadcast Operations,” *Advisory Circular*, 2012.

[12] L. E. Dixon, “FAA Faces Significant Risks In Implementing the Automatic Dependent Surveillance-Broadcast Program and Realizing Benefits,” 2010.

[13] D. L. McCallie, “Exploring Potential ADS-B Vulnerabilities in the FAA’s NextGen Air Transportation System,” Air Force Institute of Technology, 2011.

[14] D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, “Drone Hack: Spoofing Attack Demonstration on a Civilian Unmanned Aerial Vehicle,” *GPS World*, pp. 30–33, 2012.

[15] M. E. Weber, J. Y. N. Cho, J. S. Herd, J. M. Flavin, W. E. Benner, and G. S. Torok, “The Next-Generation Multimission U.S. Surveillance Radar Network,” *Bulletin of the American Meteorological Society*, vol. 88, no. 11, pp. 1739–1751, Nov. 2007.

[16] D. W. Burgess, S. I. Altman, and M. L. Wood, “TCAS: Maneuvering Aircraft in the Horizontal Plane,” *Lincoln Laboratory Journal*, vol. 7, no. 2, pp. 295–312, 1994.

[17] T. Chen, “An Authenticated Encryption Scheme for Automatic Dependent Surveillance-Broadcast Data Link,” in *Cross Strait Quad-Regional Science and Wireless Technology (CSQRWC)*, 2012, pp. 127–131.

[18] M. Bishop, *Computer Security: Art and Science*. Prentice Hall, 2003.

[19] A. Menezes, P. Van Oorschot, and S. Vanstone, *Handbook of applied cryptography*. CRC Press, 1996.

[20] Z. Cao and O. Markowitch, “Security Difference between DSA and Schnorr’s Signature,” *2009 International Conference on Networks Security, Wireless Communications and Trusted Computing*, pp. 201–204, Apr. 2009.