

Security in Next Generation Air Traffic Communication Networks



Martin Strohmeier
Kellogg College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2016

Acknowledgements

Personal

I thank my parents for the unwavering support throughout my whole life and for giving me the chance to write this thesis.

I thank Anna for always being there for me during this thesis and for providing the much-needed counterbalance to work.

I thank my friends all around the world simply for always being awesome. You make life so worthwhile.

I thank the OpenSky team Matthias, Markus and Vincent for creating such an incredible open research network, and Rui for providing his knowledge and enthusiasm on all things air traffic control.

I thank my workgroup and so many other people in the Robert Hooke Building for making most workdays both interesting and enjoyable.

I thank Andrew Martin, Andrew Simpson, and Niki Trigoni for all their helpful feedback and advice throughout the duration of my DPhil.

I also thank my examiners Andrew Simpson and Chris Johnson for the inspiring discussion during the viva and all their comments that helped to improve this thesis.

Most of all, I thank my supervisor Ivan, who made the whole process of my DPhil at Oxford not only one of immense learning but also one of great joy.

Institutional

I thank the EPSRC, the Department of Computer Science at Oxford, and Kellogg College for their generous funding which enabled me to pursue this work.

Abstract

A multitude of wireless technologies are used by air traffic communication systems during different flight phases. From a conceptual perspective, all of them are insecure as security was never part of their design and the evolution of wireless security in aviation did not keep up with the state of the art.

Recent contributions from academic and hacking communities have exploited this inherent vulnerability and demonstrated attacks on some of these technologies. However, these inputs revealed that a large discrepancy between the security perspective and the point of view of the aviation community exists.

In this thesis, we aim to bridge this gap and combine wireless security knowledge with the perspective of aviation professionals to improve the safety of air traffic communication networks. To achieve this, we develop a comprehensive new threat model and analyse potential vulnerabilities, attacks, and countermeasures. Since not all of the required aviation knowledge is codified in academic publications, we examine the relevant aviation standards and also survey 242 international aviation experts. Besides extracting their domain knowledge, we analyse the awareness of the aviation community concerning the security of their wireless systems and collect expert opinions on the potential impact of concrete attack scenarios using insecure technologies.

Based on our analysis, we propose countermeasures to secure air traffic communication that work transparently alongside existing technologies. We discuss, implement, and evaluate three different approaches based on physical and data link layer information obtained from live aircraft. We show that our countermeasures are able to defend against the injection of false data into air traffic control systems and can significantly and immediately improve the security of air traffic communication networks under the existing real-world constraints.

Finally, we analyse the privacy consequences of open air traffic control protocols. We examine sensitive aircraft movements to detect large-scale events in the real world and illustrate the futility of current attempts to maintain privacy for aircraft owners.

Contents

List of Figures	xi
List of Tables	xiii
List of Abbreviations	xv
1 Introduction	1
1.1 Motivation	1
1.2 Contributions of our Research	5
1.3 Scope of this Thesis	8
1.4 Ethical Considerations	9
1.5 Outline	10
2 Building a New Threat Model for Air Traffic Communication	13
2.1 The Traditional Electronic Warfare Model	13
2.2 The New Threat Model	14
2.3 Taxonomy of Threat Agents in Civil Aviation	17
2.4 Summary	20
3 Wireless Communication Technologies in Aviation	23
3.1 Air Traffic Control	26
3.1.1 Voice (VHF)	26
3.1.2 Controller Pilot Data Link Communications (CPDLC)	26
3.1.3 Primary Surveillance Radar (PSR)	28
3.1.4 Secondary Surveillance Radar (SSR)	29
3.1.5 Automatic Dependent Surveillance – Broadcast (ADS–B)	29
3.1.6 Multilateration (MLAT)	32
3.2 Information Services	33
3.2.1 Aircraft Communications Addressing and Reporting System (ACARS)	33
3.2.2 Traffic Alert and Collision Avoidance System (TCAS)	34
3.2.3 Flight Information System – Broadcast (FIS–B)	35

3.2.4	Traffic Information System – Broadcast (TIS-B)	35
3.3	Potential Future Technologies	35
4	Examining Current Vulnerabilities in Aviation Technologies	37
4.1	Wireless Attack Vectors	37
4.2	Attacks on Air Traffic Control Technologies	40
4.3	Attacks on Information Systems	48
4.4	Communication in Military ATC	50
4.5	Summary of Vulnerabilities and Attacks	51
4.6	Mitigating Factors	52
4.7	Reasons for the Current State of Aviation Security	55
5	Raising Awareness of Cybersecurity in Aviation	59
5.1	Survey Design	60
5.2	Survey Limitations	61
5.3	Survey Results	62
5.4	Assessment of Concrete Scenarios	67
5.5	Summary & Outlook	70
6	Developing a Taxonomy of Air Traffic Communication Security Approaches	73
6.1	Secure Broadcast Authentication	74
6.1.1	Non-Cryptographic Schemes on the Physical Layer	75
6.1.2	Public Key Cryptography	77
6.1.3	Retroactive Key Publication	80
6.2	Secure Location Verification	82
6.2.1	Multilateration	82
6.2.2	Distance Bounding	84
6.2.3	Kalman Filtering and Intent Verification	86
6.2.4	Group Verification	87
6.2.5	Plausibility Checks	88
6.3	Summary	90
7	Transparent Security for Air Traffic Communication	93
7.1	Modeling False-Data Injection Attackers	94
7.2	The OpenSky Network	95
7.3	Data Link Layer Fingerprinting	97
7.3.1	Feature Engineering	98
7.3.2	Experimental Design	100
7.3.3	Evaluation	102

7.3.4	Detection of Anomalies	104
7.3.5	Discussion	106
7.4	Physical Layer Intrusion Detection	107
7.4.1	Attacker Model	107
7.4.2	Feature Selection	108
7.4.3	Combined Anomaly Detection	111
7.4.4	Experimental Design	111
7.4.5	Evaluation	112
7.4.6	Discussion	114
7.5	Lightweight Aircraft Location Verification	115
7.5.1	Attacker Model	115
7.5.2	Considerations about Aircraft Localization	117
7.5.3	Designing Lightweight Aircraft Location Verification	122
7.5.4	Experimental Design	127
7.5.5	Evaluation	130
7.5.6	Discussion	136
7.6	Comparison of Transparent ATC Security	137
8	New Privacy Challenges in Wireless Air Traffic Communication	141
8.1	Introduction	141
8.2	Data Sources for Aircraft Information	143
8.3	Passive Tracking of Aircraft	144
8.3.1	Tracking using Commercial Web Services	144
8.3.2	Tracking using OpenSky	145
8.3.3	Discussion	148
8.4	Information Collection using the Data Link Layer	150
8.4.1	Transponder Identification	150
8.4.2	Privacy Implications	151
8.4.3	Mitigation	153
8.5	Large-scale Event Detection	154
8.5.1	Experimental Design	154
8.5.2	Feature Selection	155
8.5.3	Evaluation	156
8.5.4	Discussion	159
8.6	Summary	160
9	Summary & Future Work	161
9.1	Summary of Results	161
9.2	Future Work	162
9.3	Final Conclusions	164

Appendices

A Large-Scale Event Detection Data 2014-2015	169
---	------------

References	173
-------------------	------------

List of Figures

1.1	Google Scholar results over time for two searches related to air traffic communication security. Note that the graph does not account for the increase in indexing and digitization occurring over the examined timespan.	4
3.1	An overview of the wireless technologies used in air traffic communication, between ground stations, aircraft and satellites.	24
3.2	Overview of the ADS-B system architecture.	30
3.3	ADS-B protocol hierarchy.	31
3.4	1090 ES Data Link.	32
4.1	Denial of service due to an overload of aircraft on a controller's screen.	41
4.2	Result of an ATC injection attack.	46
4.3	DF19 data format	51
5.1	Survey assessment of the flight safety impact, the likelihood of being attack targets and the trustworthiness against manipulation of each discussed protocol.	63
5.2	Survey assessment of the safety impact of aviation technologies. . .	64
5.3	Survey assessment of the security capabilities of aviation technologies.	65
6.1	Taxonomy of air traffic communication security approaches.	74
6.2	Illustration of Uncoordinated Frequency Hopping	77
6.3	TESLA's utilization of one-way chains.	81
6.4	Basic multilateration architecture.	84
6.5	Principle of distance bounding protocols.	85
6.6	Practical application of intent verification in ATC.	87
6.7	Illustration of the group concept.	88
7.1	OpenSky sensor coverage in Europe (June 2016).	96
7.2	Illustration of five different transponder types.	99
7.3	Schematic showing two slots as used by the transponder implementations, determined by measured inter-arrival times.	100
7.4	Visualization flight trajectories used for our data analysis.	101

7.5	Representative illustration of different message types.	103
7.6	RSS samples of a flight's two separate antennas.	110
7.7	Visualization of the flight trajectories used for anomaly detection. .	111
7.8	Example of a 2D-Parzen classifier used for anomaly detection with 200 collected flight samples.	113
7.9	Complete classifier comparison with 5-fold cross validation.	114
7.10	Graphical overview of four distinct attacker types considered for aircraft location verification.	116
7.11	Illustration of geometric dilution of precision.	119
7.12	Map of the practical reception ranges of an 8 sensor ADS-B system.	121
7.13	An example illustrating the calculation of expected TDOAs.	123
7.14	Location estimation with 3-NN in an adversarial setting.	127
7.15	Optimal choice of neighbours k for different square sizes.	132
7.16	Median location error depending on the level of noise N affecting the measurements.	134
7.17	Median location error depending on the geometric dilution of precision affecting the measurements.	135
8.1	Coverage of two sensors in the Davos, Switzerland area in 2016. . .	154
8.2	Illustration of three time-series features in the Eastern Swiss surveil- lance area of OpenSky around the time period of Davos 2014. . . .	157
8.3	Illustration of two time-series features in the Eastern Swiss surveil- lance area of OpenSky around the time period of Davos 2015 and 2016.	157

List of Tables

2.1	Overview of threat agents in air traffic communication.	18
3.1	Short-handles and full names of aviation communication technologies, systematized into applications.	25
3.2	Detailed characteristics of air traffic control protocols.	27
3.3	Detailed characteristics of information services protocols.	33
4.1	Overview of attacks and security requirements for wireless aviation communication systems.	52
5.1	Occupations of survey respondents.	62
5.2	Answer distribution for nine hypothetical safety-critical scenarios. .	68
6.1	Overview of capabilities of various security approaches against feasible wireless attack primitives on ATC protocols.	90
6.2	Overview of feasibility attributes of various security approaches for air traffic protocols.	91
7.1	Statistics about the OpenSky dataset utilized for data link finger- printing.	101
7.2	Feature combinations of different transponder implementations. . .	102
7.3	Distribution of different transponder types in our dataset.	103
7.4	Time stability of the analysed transponder features over different days.	104
7.5	Evaluation of attacker classes 1-3 using inter-arrival time-based fingerprinting.	106
7.6	Effectiveness of the examined anomaly detection approaches.	112
7.7	Statistics on OpenSky dataset used for aircraft location verification.	122
7.8	Averaged horizontal distances from the four attackers' positions to their claimed aircraft positions during the time that flight data is injected.	129
7.9	Results of the location verification approach.	130
7.10	Horizontal errors in different grid square sizes using k-NN vs. MLAT.	131
7.11	Average horizontal errors using k-NN.	133

7.12	Mean distances between estimates and claimed location injected by an attacker & mean distances to actual horizontal location of an attacker.	136
7.13	Comparison of the proposed transparent ATC security approaches. .	139
8.1	Breakdown of aircraft seen by OpenSky and their privacy restrictions on public tracking websites.	146
8.2	Breakdown of identifiable types of aircraft with and without flight history, and with full blocks on tracking websites used for comparison.	147
8.3	Manufacturers of the various transponder behavior classes.	151
8.4	Aircraft movement and meta data features for the World Economic Forum 2016.	158
A.1	Aircraft movement and meta data features for the World Economic Forum 2014.	170
A.2	Aircraft movement and meta data features for the World Economic Forum 2015.	171

List of Abbreviations

A2A	Aircraft-to-Aircraft
ACARS	Aircraft Communications Addressing and Reporting System
ACF	Autocorrelation Coefficient
ADS-B	Automatic Dependent Surveillance - Broadcast
AM	Amplitude Modulation
AMS	ACARS Message Security
ANSP	Air Navigation Service Provider
AoA	Angle of Arrival
ASDI	Aircraft Situation Display to Industry
ATC	Air Traffic Control
ATM	Air Traffic Management
CA	Certificate Authority
CNS	Communication, Navigation, Surveillance
COTS	Commercial Off-the-shelf
CPDLC	Controller–Pilot Data Link Communications
CRC	Cyclic Redundancy Check
DF	Downlink Format
DOP	Dilution of Precision
DoS	Denial of Service
DSSS	Direct Sequence Spread Spectrum
ECDSA	Elliptic Curve Digital Signature Algorithm
ESCAT	Emergency Security Control of Air Traffic
FAA	Federal Aviation Administration
FHSS	Frequency Hopping Spread Spectrum
FIS-B	Flight Information System-Broadcast

FP	False Positive
GA	General Aviation
GDOP	Geometric Dilution of Precision
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HFDL	High Frequency Data Link
ICAO	International Civil Aviation Organization
IDS	Intrusion Detection System
IFF	Identification, Friend or Foe
IFR	Instrument Flight Rules
k-NN	k-nearest Neighbour
LoS	Line of Sight
MAC	Message Authentication Code
METAR	Meteorological Terminal Aviation Routine Weather Report
MLAT	Multilateration
NIST	National Institute of Standards and Technology
OSINT	Open Source Intelligence
PKI	Public Key Infrastructure
PPM	Pulse Position Modulation
PSR	Primary Surveillance Radar
PUF	Physically Unclonable Functions
RA	Resolution Advisory
RF	Radio Frequency
RFID	Radio-frequency Identification
RSS	Received Signal Strength
RTCA	Radio Technical Commission for Aeronautics
SDR	Software-defined Radio
SSR	Secondary Surveillance Radar
STANAG	Standardization Agreement
TA	Traffic Advisory
TCAS	Traffic Alert and Collision Avoidance System

TDoA	Time Difference of Arrival
TESLA	Timed Efficient Stream Loss-Tolerant Authentication
TIS-B	Traffic Information System - Broadcast
ToA	Time of Arrival
UAT	Universal Access Transceiver
UAV	Unmanned Aerial Vehicle
UFH	Uncoordinated Frequency Hopping
VANET	Vehicular Ad Hoc Network
VDL	VHF Data Link
VDOP	Vertical Dilution of Precision
VFR	Visual Flight Rules
VHF	Very High Frequency
WEF	World Economic Forum

*Sehn wir doch das Große aller Zeiten
Auf den Brettern, die die Welt bedeuten,
Sinnvoll still an uns vorübergehn.*

*Yet we see the great of ev'ry age
Pass before us on the world's wide stage
Thoughtfully and calmly in review.*

— Friedrich Schiller's *An die Freunde*

1

Introduction

Contents

1.1	Motivation	1
1.2	Contributions of our Research	5
1.3	Scope of this Thesis	8
1.4	Ethical Considerations	9
1.5	Outline	10

1.1 Motivation

Air traffic control (ATC) is the backbone of what is arguably the key means of personal transport in the modern world. As the traffic load continues to grow dramatically, ATC has to manage ever more aircraft. Large European airports, such as London Heathrow or Frankfurt/Main, experience spikes of more than 1,500 daily take-offs and landings, and industry forecasts predict that world-wide flight movements will double between 2015 and 2034 [1]. Additionally, as Unmanned Aerial Vehicles (UAV) enter the civil airspace, they must learn to co-exist with manned aircraft and existing air traffic control systems. While forecasts project a steady 5% annual increase of global manned flight traffic over the next 20 years, UAV are projected to outgrow traditional air traffic by several orders of magnitude:

In 2035, 250,000 UAV are expected to be operating in the US alone, compared to a mere 45,000 passenger aircraft around the globe [2]. However, as this paradigm shift progresses, many technical and policy issues are yet to be solved to ensure the safe control of both manned and unmanned aircraft. In this dissertation, we will address one of the most urgent ones: the fact that current wireless air traffic communication technologies are inherently insecure.

Historically, air traffic management (ATM) and its associated wireless communications technologies are rooted in the military. Most of the improvements in communication, navigation and surveillance (CNS) technologies are direct results of wartime developments [3]. For instance, surveillance radar systems and navigation functions which were developed originally for the armed forces were later adopted for civilian aviation. This change of purpose and application also shifted the threat models affecting these wireless technologies considerably. Where the military can often also rely on secrecy, security through obscurity, and superior proprietary technologies to prevail in an arms race, the requirements in a civil setting of worldwide collaboration are different. In this environment, a pure security by design approach, such as the protection of critical wireless communication through standard cryptographic countermeasures, would be highly preferable. Unfortunately, in the slow-changing industry of aviation, such a radical switch of technologies is not currently on the horizon.

The civil aviation community emphasizes safety and has a sound and steadily improving safety record. Security, however, is not safety, and requires a different approach. While we encountered many helpful and interested people and institutions in aviation during our investigations over the past four years, the prevalent feeling too often still seems to be: “Why is security needed? Is air traffic communication not safe currently?”. Indeed, historically, few incidents had been (publicly) recorded where communication technologies were maliciously exploited to successfully cause distress to aircraft. Consequently, even recently-developed aviation technologies, which make the shift from traditional radar to modern digital communication

networks, do not include security by design in their specifications; instead, the systems rely almost exclusively on redundancy and secured supply chains.

However, with the widespread availability of cheap and powerful tools such as software-defined radios (SDR), the aviation community has lost the considerable technical advantage that protected its communication over the past decades. This disruption is illustrated by the recent proliferation of reports about potential cyber attacks on wireless ATC technologies. High-profile incidents, such as the case of hijacked emergency signals [4] or alleged military exercises causing aircraft to vanish from European radar screens [5], created a lot of speculation in the media about the potential impact of insecure technologies on the safety of air traffic [6, 7]. Naturally, such speculations are widespread now, simply because it has become clear that attacks on the wireless communication systems of aviation are potentially feasible.

In this thesis, we seek to analyze the disruptive effect of SDRs on aviation, examine the extent of both feasibility and impact of potential wireless attacks, and aim to develop countermeasures appropriate to the wireless ecosystem found in aviation.

History of Security Research in Air Traffic Communications

To understand the currently poor state of wireless security in aviation, it is helpful to take a look at the historical context in terms of research and technological development. For a long time, wireless security was not a concern for aviation as the technological advantage rendered attacks highly unlikely. Consequently, little academic research on the topic was conducted (see Fig. 1.1).

In contrast, securing wireless protocols has been a long-standing and popular research topic for the security community. Indeed, many security issues around widely deployed technologies such as WiFi are now considered solved thanks to the application of cryptography, despite common failures of concrete implementations.

However, as we have learned throughout our research and will outline in this thesis, the security of real-world air traffic communication systems is not a similarly well-defined problem. Thus, even though wireless security research offers many

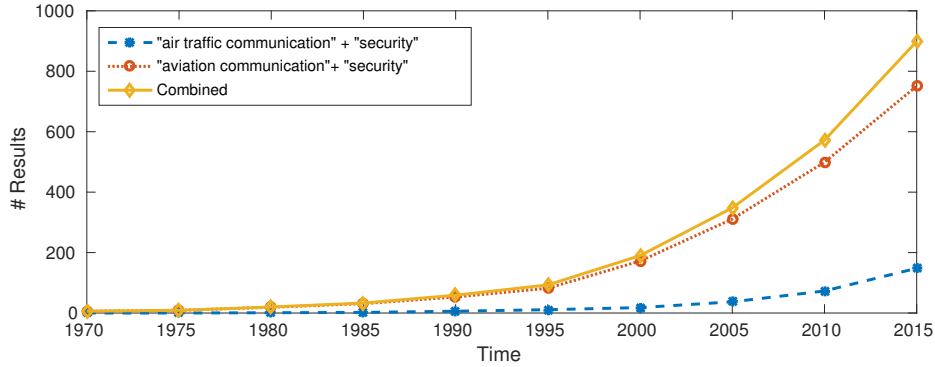


Figure 1.1: Google Scholar results over time for two searches related to air traffic communication security. Note that the graph does not account for the increase in indexing and digitization occurring over the examined timespan.

mature solutions, which hold under the strongest threat models, they are not easily transferable to the aviation context.

Indeed, over the last five years, parts of the security community recognized this issue. With the roll-out of new “Next Generation” air traffic control technologies, academic researchers and hackers took a closer look at the new protocols.

The current discussion was sparked by a series of articles and presentations on the Automatic Dependent Surveillance - Broadcast (ADS-B) protocol dating back to 2009. Righter Kunkel addressed the problem in two talks at the DEFCON hacker conference, as did Sampigethaya et al. [8] Two years later, McCallie et al. [9] from the Air Force Institute of Technology also pointed out that the soon-to-be-deployed next generation protocols remained fundamentally vulnerable. This concern was validated through emerging proof-of-concept attacks by the hacker community in the following year [10, 11]. Academic security researchers followed up with more in-depth analyses of the ADS-B protocol [12, 13]. Significant research has since been conducted, both on the security of newly-developed air traffic control protocols, and on the future of the “e-enabled” aircraft with its numerous electronic systems in general [8, 14].

These revelations generated many headlines in the mainstream press [15–18], leading to accusations of overblown media reporting from members of the aviation community. Some distrust the possibility of the cited hacks of aircraft IT systems in the real world [19, 20], while others generally doubt the impact of attacks on

wireless air traffic communication systems in practice due to the widely deployed checks and balances common to aviation [21].

We believe that these instances show that many who understand wireless security do not have the necessary aviation expertise. Likewise, many stakeholders in aviation know the processes and procedures but do not realize the severity of modern cybersecurity issues. However, things have been moving in the right direction over the past few years, as indicated, among other things, by the number of research articles published on this topic. For example, large aviation conferences have for the first time included cybersecurity tracks in their programme, significantly raising the profile in the community.

Based on these insights, we formulated the following research goals for this thesis:

- **Analyse the vulnerabilities present in current and next generation air traffic communication technologies.**
- **Understand the security knowledge and awareness of the aviation community.**
- **Understand the domain specific problems of bringing security to aviation.**
- **Finally, develop security solutions that are both practical and acceptable for the aviation community.**

1.2 Contributions of our Research

This section explains the contributions made by our research by outlining the impact our published work, which provides the foundation of this thesis.

- Relying on the insight that cheap and available technologies void the technological advantage that the aviation community enjoyed over the 20th century, we developed a new threat model for wireless air traffic communication in our paper **Assessing the Impact of Aviation Security on Cyber Power**

published in *8th International Conference on Cyber Conflict (CYCON), 2016* [22]. We elaborate on new threat agents to be taken into consideration to protect critical infrastructures (Chapter 2).

- Chapters 3 and 4 are based in large parts on our paper **On Perception and Reality in Wireless Air Traffic Communication Security** to appear in a forthcoming issue of *IEEE Transactions on Intelligent Transportation Systems* [23]. Created in collaboration with an air traffic controller, this work serves as an introductory read to the field of wireless security in air traffic communication networks and examines the consequences of using authentication-less protocols. It clearly identifies the vulnerabilities in existing and next generation technologies and relates these findings to the slow-changing aviation field, which has not kept up with the technological advancements.
- For the same work, we have conducted a large-scale study with almost 250 aviation experts, which is reported in Chapter 5. From our casual conversations with members of the aviation community, we had previously learned that there is a problematic ignorance surrounding wireless security in aviation, which we quantify with this survey. We consider it a further contribution that awareness of these problems has been increasing as a result of our engagement.
- Chapter 6 provides a comparative evaluation of potential solutions taken from the academic literature. This work is based on the article **On the Security of the Automatic Dependent Surveillance - Broadcast Protocol** published in *IEEE Communications Surveys & Tutorials* [24]. It looks at existing solutions from the wireless security community and examines their feasibility in the aviation context. We conclude that no currently existing countermeasure is both practical and effective.
- Consequently, we argue that the concrete research problem is to find transparent countermeasures that do not require changes in current aviation systems but instead use existing inputs to improve the security of the system quickly

and effectively. We have developed and implemented three such solutions (presented in Chapter 7):

1. **Passive Data Link Layer Fingerprinting of Aircraft Transponders** published in *1st ACM Workshop on Cyber-Physical Systems Security & Privacy (CPS-SPC), 2015* [25].

2. **Intrusion Detection for Airborne Communication using PHY-Layer Information** as published in *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA), 2015* [26].

3. **Lightweight Location Verification in Air Traffic Surveillance Networks** in *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security (CPSS), 2015* [27]. This paper won the best paper award and has been extended for journal submission after the initial submission of the present thesis [28].

- Beside security, we discuss privacy issues in the context of our new threat model for aviation. As laid out in [22] and our paper **OpenSky: A Swiss Army Knife for Air Traffic Security Research** published at the *34th Digital Avionics Systems Conference (DASC), 2015* [29], having neither authentication nor confidentiality can lead to severe privacy concerns. Again, SDR-powered advances in wireless technology and the broad availability of flight data require awareness and action from aviation circles. We discuss the consequences of this development in Chapter 8. Our paper at DASC also won the best paper award (out of more than 170 accepted papers), reflecting the appreciation of the avionics community.
- Finally, we have built the OpenSky Network (in conjunction with TU Kaiserslautern, Germany and armasuisse, Switzerland) to conduct our research and enable and facilitate other researchers to do the same. As described in [29] and in **Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research** published at the *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), 2014* [30], it is a

collaborative sensor network that collects and stores air traffic communication data on a large scale. While OpenSky itself is not an integral topic of this dissertation, it provides for most of the underlying data collection used in our work. We strongly believe that our project can help bring security in aviation into the 21st century and also support other valuable research. OpenSky is available at <http://www.opensky-network.org>.

We believe that private and public networks that collect air traffic communication have disrupted the landscape of security and privacy in aviation and will continue to do so for the immediate future. We argue in this thesis that the growing ubiquity of SDRs and cheap sensor networks such as OpenSky forms part of both the problems and the solutions for security. In terms of privacy, we realize these technologies at the same time open up new challenges, discussed in Chapter 8, which the aviation community at large will have to address quickly. Traditionally, aviation is a very slow-moving area, similar to many critical infrastructure industries, making it difficult to deal with significant disruptions. However, anecdotal evidence seems to indicate that things may be moving in the right direction. Our work seeks to help with that development by providing new insights and lessons for this current and potentially similar future disruptions.

1.3 Scope of this Thesis

This thesis deals with the wireless technologies that provide the basis for the communication between aircraft and ground. We focus on currently deployed protocols and thus exclude technologies that are still in development and thus neither finalized yet nor allow practical analysis (e.g., AeroMACS). We further exclude navigation aids such as GPS, which have shown to be vulnerable in their own right [31]. Lastly, we do not consider the large and growing area of malware-related security vulnerabilities and attacks on the physical infrastructure, i.e. the soft- and hardware on endpoints that power aviation. Instead, we focus exclusively on the wireless channel as a separate attack vector with which the ATC system can be compromised.

1.4 Ethical Considerations

We acknowledge the potentially sensitive nature of analysing air traffic communication protocols and increasing the awareness of easily exploitable security flaws. ATC is considered a critical public infrastructure and protects the lives of billions of people every year. However, while in this thesis we scrutinize the wireless communication system used in aviation as a whole, some of these technologies have already been widely analysed by both hacker communities and academic researchers recently.

That these protocols are inherently vulnerable has been discussed for years and even concrete exploits have been demonstrated in hard- and software. As such, we by and large only describe the possibilities that we believe are already available to the different threat agents as described in Chapter 2. With this thesis, we ultimately aim to increase the awareness of this fact in aviation circles, as we strongly believe that this particular awareness (in contrast to unsubstantiated media narratives) is what it takes to secure the aviation communication systems of the future. Indeed, we feel that not only the academic research into this topic has accelerated over the past few years as indicated in Fig. 1.1, but also the perception within aviation circles has changed for the better.

While we consider the overall state of aviation communications security a serious one, redundancy and existing processes still mostly protect current airspaces as we explain in Chapter 4. Nonetheless, we do not make available any of our penetration testing tools used to emulate real aircraft communication using SDRs. While they are not difficult to write for a determined attacker – and indeed briefly before the submission of this thesis a ADS-B signal generator has first been made publicly available on GitHub – we have no interest in increasing the ease of exploitation any further.

We further follow a responsible disclosure process, working with ATC institutions during all phases of our research. We notified these institutions of our results and plan to work closely with them in the future.

1.5 Outline

This dissertation is structured as follows:

- **Chapter 2** discusses the impact of recent technological developments on aviation security. Based on this, we build a new threat model for air traffic communication, which is assumed throughout the later chapters.
- **Chapter 3** provides the necessary background to the wireless technologies and systems used in air traffic communication. We explain their usage within the aviation system, detail their technical features, and discuss their safety impact.
- **Chapter 4** reviews the existing literature on air traffic security, both concerning possible attacks on aviation systems and research into countermeasures. It also examines the reasons for the currently poor state of wireless security in aviation.
- **Chapter 5** looks at crucial human factors and reports from a survey of 242 international aviation experts to examine their awareness of the security issues present in aviation.
- **Chapter 6** evaluates and compares the relevant research avenues from the wider field of wireless security and discusses their applicability to the field of air traffic communication.
- **Chapter 7** proposes an attack detection system for air traffic communication. We discuss, implement, and evaluate three different approaches based on physical and data link layer data obtained from live aircraft. We show that our attack detection approaches are able to defend against the injection of false aircraft data and can significantly and immediately improve the security of air traffic communication systems under existing real-world constraints.

- **Chapter 8** analyses the privacy consequences of aircraft tracking carried out by merely passive observers using commercial services or private receiver installations. We look at sensitive aircraft movements and blocking policies of public trackers and illustrate the futility of current methods to maintain traditional privacy for aircraft owners. We further show that it is possible to use aircraft meta data to detect unusual real-world events and validate our approach using the well known World Economic Forum.
- **Chapter 9** finally summarizes the results, discusses the future work that is required to secure modern air traffic communications, and concludes this thesis.

The only man I know who behaves sensibly is my tailor; he takes my measurements anew each time he sees me. The rest go on with their old measurements and expect me to fit them.

— George Bernard Shaw’s *Man and Superman*

2

Building a New Threat Model for Air Traffic Communication

Contents

2.1	The Traditional Electronic Warfare Model	13
2.2	The New Threat Model	14
2.3	Taxonomy of Threat Agents in Civil Aviation	17
2.4	Summary	20

In this chapter, we define a new threat model for wireless systems in aviation. We distinguish between the traditional adversarial electronic warfare model and the recently emerged modern threat model. The modern threat model is based on a) the widespread distribution of accessible software-defined radios and b) the ongoing move from analogue to digital communication systems (as pointed out in, e.g., [32]). We consider purely *passive* observers as well as *active* adversaries with the capability to eavesdrop, modify, and inject data on the communications channel.

2.1 The Traditional Electronic Warfare Model

The traditional threat model has been implicitly and explicitly used in aviation since the introduction of radio communication and radar in civil air traffic control in the first half of the 20th century. Surveillance radar, navigation, and communication

systems originated from military applications and were later integrated into the civil aviation airspace [33]. We characterize the threat model used as comparatively naïve, reflecting the general state of computer security considerations within industrial and infrastructure systems during this period. In short, the model makes the following main assumptions about adversaries and technologies on which today’s aviation communications security is still based:

- **Inferior technological capabilities:** Active adversarial capabilities were ascribed only to military and nation-state attackers with the ability to conduct electronic warfare [34].
- **Inferior financial capabilities:** Similarly, it is further assumed that electronic devices capable of distorting radar are financially out of reach for all but the most capable attackers.
- **Requirement of inside knowledge:** An impactful attacker needs to be (or have contact to) an *insider* to be able to obtain the knowledge of communication systems and general aviation conduct necessary for an attack.
- **Use of analogue communication:** Typical attacks on analogue communications are easier to detect for the *user*. For example, somebody hijacking the voice channel or causing a denial of service at a PSR will typically be detected immediately.

2.2 The New Threat Model

With the technological advancements of the late 1990s and 2000s, the aviation threat model changed drastically, as the assumptions about adversaries and their capabilities ceased to hold. In the 1990s, software-defined radios were first practically adapted for military and closed commercial use [35]. Later, open-source projects such as GNURadio [36], released in 2001, and finally the availability of cheap commercial off-the-shelf (COTS) software-defined radio transceivers spread adversarial capabilities to a large and expanding group of people. They enabled a broad community

with basic technological understanding to receive and process, craft and transmit arbitrary wireless signals – including those used in aviation. Contrary to the pre-SDR era, hardware need not be purpose-built any more (requiring considerable technical and financial resources) but can simply be programmed and re-programmed on the fly with the necessary code and knowledge easily shared via the Internet.

We make the following assumptions for an adversary model that is adequate for wireless communications security in modern aviation and which we use throughout this work:

- **Increased digitization and automation:** There is a general trend in aviation towards transmitting sensitive data (such as flight clearances) using unauthenticated digital communication networks. While attacks on analogue technologies such as VHF have been included in the traditional threat model, new *digital attacks* are emerging which are easy to execute, potentially devastating, and difficult to detect on the data link level for increasingly *automated systems* and their users [37].
- **Increased technological capabilities:** With the widespread availability of cheap SDR technology, it is reasonable to assume that a large amount of people are capable of conducting wireless attacks on aviation systems. The financial barrier is all but gone with SDR receivers available from as little as \$10 while capable senders cost less than \$300 with a strong downwards trajectory. In conjunction with downloadable software, this development enables a new class of *unsophisticated* attackers.
- **Easy availability of aviation knowledge:** Attackers today can easily gain the necessary knowledge about processes and conventions in aviation communications. Syntax and semantics of wireless protocols can be obtained by *outsiders* through openly accessible means, such as specification protocols, forums, plane-spotting websites, and finally by capturing and examining real-world communication data.

SDR Impact on Aviation

In summary, the fact that wireless attack capabilities are shifting from military adversaries to script kiddies, hobbyists, white hat hackers, cyber crime organisations, and terrorists increases the likelihood of attacks manifold. In conjunction with the move towards unsecured digital networks, and increased deployment of homogeneous COTS hard- and software, the new threat model goes beyond denial of service through traditional jamming and requires us to rethink and address wireless security in aviation.

The advent of SDR technology has provided a surge of accessible applications for radio communications in general. The former assumption that access to the frequencies used by important communication technologies is hard has been voided. Modulations of virtually all radio applications are well known and made available freely through the SDR community. Thus, the ability to eavesdrop and manipulate any wireless communication channel is available to any interested observer without the requirement of significant resources or specialist knowledge. Examples of such possibilities are the trivial access to mobile phone networks, satellite signals, television channels, or wireless sensor networks.

One of the most active and enthusiastic SDR communities is concerned with aviation communication and flight tracking. Using, for example, the popular RTL-SDRs, a \$10 USB stick re-purposed as software-defined radio receiver, a plane-spotter can choose between several different software options to receive virtually all air traffic communication protocols in use today. Countless enthusiasts and volunteers around the world use this hard- and software to power a multitude of services such as <http://flightradar24.com>, <http://opensky-network.org>, or <http://adsbexchange.com>, where an ever-increasing number of flight movements can be followed live and without delay. Data from flight trackers has been involved regularly in investigations following flight incidents such as the Germanwings crash [38] or the two Malaysian aircraft lost over the Ukraine [39] and the Indian Ocean [40] in 2014, illustrating the impact of the changing communications landscape on aviation.

2.3 Taxonomy of Threat Agents in Civil Aviation

Based on the insights from the previous sections, we develop a new threat agent model for wireless attacks in the aviation context. We analyse possible attackers based mainly on a) their resources, b) their subject-matter expertise and c) their motivation. Table 2.1 presents the threat agents applicable to wireless communications in aviation, sorted by their approximate capabilities, who we discuss in detail in the following. Our taxonomy is very loosely inspired by the relevant definitions of the US National Institute of Standards and Technology (NIST) [41], but adapted for the unique context of the cyber-physical aviation system. While our approach to threat agents in aviation is novel, we believe that tying it into the existing NIST framework leads to easier application in practice. Our taxonomy provides new insights into the specific technological capabilities of different classes of threat agents and how these are utilized to achieve their respective goals, even in light of potential countermeasures.

While out of scope for this dissertation, similar active agents could be part of a threat model for the protection of the hard- and software supply chain for ATC systems, which is a general risk vector in safety critical systems [42]. As recently reported by the United States Government Accountability Office [43], there have been several incidents where malicious software has been found on computer systems related to ATC. Thus, securing the computers and networks used in ATC is clearly imperative to safely control the airspace. However, even when a system itself is never compromised – if its inputs over the wireless channel can be tampered with, all security guarantees are void.

Furthermore, we do not consider insider threats as a separate agent in our model but rather an orthogonal threat, i.e., they permeate all five agents, passive and active. For wireless communication attacks the concept of an insider is less well-defined compared to conventional threat models that pertain to wired networks and computer systems. Through the inherent broadcast quality of the wireless medium an outsider in principal has the same physical access as an insider, both for sending and receiving data. While an insider might still enjoy privileged

Threat Agent	Capabilities	Hardware/Cost	Motivation
Passive Observers	Eavesdropping, use of websites & mobile apps.	Internet access, \$10 SDR receiver stick.	Information collection / Financial or personal interest
Script Kiddies / Hobbyists	Eavesdropping, replay attacks, denial of service.	COTS SDR transmitter, \$300-\$2,000.	Any noticeable impact / Thrill and recognition
Cyber Terrorism	Resources for specific high-impact operations, though usually on a limited scale.	Directional antennas, small UAVs with SDR transmitters, \$5-10,000.	Political or religious motivation / Massive disruption and casualties
Cyber Crime	Resources for large-scale operations with sophisticated transponders.	As with cyber terrorism but potentially on a larger, more targeted scale.	Maximizing impact / Financial gains using e.g. blackmail or valuable information
Nation States	Anything computationally and physically possible.	Military-grade radio equipment, capability for electronic warfare.	Weapons / Targeting specific, potentially military objects

Table 2.1: Overview of threat agents in air traffic communication. The four active attackers are sorted by their approximate capabilities.

information such as the exact position of wireless hardware, or deeper knowledge of the data processing functions in the typically proprietary systems, we subsume these capabilities under the existing threat agents.

Passive Observers

Passive observers are interested persons who exploit the open nature of air traffic communication protocols to glean information. This class of threat agents does not actively interfere with air traffic communication but instead uses public and private websites and mobile applications, which display air traffic and its communications in real time, to gather information about private or secret air traffic movements. Alternatively, they can employ cheap SDR receivers to gather their own undistorted picture of all air traffic in their vicinity, in real time or stored for historic analysis. The information collected by such merely passive observers can be exploited in multiple ways, ranging from privacy concerns to detection of military operations,

which are discussed in detail in Section 8.3. Apart from this, passive observation often forms the basis of attacks executed by active threat agents.

Script Kiddies and Hobbyists

Script kiddies and hobbyists are the lowest active threat in our model based on their abilities considering both hardware and knowledge. Their aim is to exploit well-known security holes with existing, easy-to-use attacks with typically low sophistication. Their motivation is regularly not rational; instead any identifiable impact is sought for thrill and recognition. We assume a typical attack to be the following: using a programmable transponder, they listen in on legitimate radio communication, modify the call sign and/or information such as position and velocity, and play it back. The objective of the attacker is to have their signals either show up as a new aircraft with an unexpected call sign or as an existing aircraft causing conflicting information. We assume that the attacker is on the ground and sends with the standard parameters of their transponder.

Hobbyists are typically interested in plane-spotting and more familiar with the norms and protocols in modern ATC, either due to personal interest in aviation or because it relates to their job. Further to this, they are more knowledgeable about radio communication and the basic characteristics of the wireless channel. They have access to SDRs and are able to operate it with matching software frameworks such as GNU Radio. Their attack is similar to the script kiddies' but it is not detected by naïve plausibility checks on the content and data link level.

Cyber Terrorism

Attacks on cyber-physical systems powering critical infrastructures such as aviation are a natural target for terrorists and politically motivated attacks. Terrorists seek to threaten national security, cause mass casualties, weaken the economy, and damage public morale and confidence [41]. By exploiting vulnerabilities in wireless aviation communications, terrorist groups, who traditionally hijack or crash planes using physical force, could mount attacks on planes from the ground within safe

distances. While there may currently be simpler options with higher threat levels (such as the use of UAV or malware) to achieve the same purpose, we cannot entirely discount this threat in the future.

Cyber Crime

The cyber crime attacker class seeks to attack systems for monetary gain. Equipped with sufficient subject-matter knowledge, software-defined radios, and potentially even small unmanned aerial vehicles, they are able to inject new messages and/or modify existing ones in such ways that they are not flagged by current detection systems. Cyber crime attackers are typically interested in causing maximum damage and exerting credible threats, as a pre-requisite for, e.g., blackmail or to take advantage of captured inside knowledge [41]. Consequently, they are interested in exploiting any potential and effective way to attack ATC and aircraft systems.

Nation States

With sufficient knowledge and near-unlimited resources, it is possible to bypass plausibility checks and redundancy-based defences. While it becomes increasingly difficult to deceive multiple ATC systems at the same time, it remains possible. However, we argue that attacks on PSR remain achievable only by a nation state actor and thus part of the electronic warfare threat model traditionally outside the scope of securing civil aviation. In this case, new protocols with authentication through cryptographic means may help further, although this is unlikely to happen in the foreseeable future due to the reasons outlined in Section 4.7.

2.4 Summary

This chapter contrasted the old electronic warfare-based threat model that aviation has dealt with for many decades with the new threat model facilitated by cheap COTS SDRs. The key takeaway is that the latter enabled new classes of threat agents below the nation state which have to be considered in the future. Naturally, it is also conceivable that different types of actors work together, for example when

cyber criminal networks act as an agent of state-based policy similar to the attacks on Georgia several years ago[44]. In this case, the distinctions between criminals, terrorists and nation state actors can blend together and become blurry.

Nevertheless, we will refer to the specific threat agents where appropriate in this work, in particular when considering the effectiveness of the security mechanisms in Chapter 7 with regards to the different capabilities.

*On parle toujours mal quand on n'a rien à dire.
One always speaks badly when one has nothing to
say.*

— Voltaire [45]

3

Wireless Communication Technologies in Aviation

Contents

3.1	Air Traffic Control	26
3.1.1	Voice (VHF)	26
3.1.2	Controller Pilot Data Link Communications (CPDLC) .	26
3.1.3	Primary Surveillance Radar (PSR)	28
3.1.4	Secondary Surveillance Radar (SSR)	29
3.1.5	Automatic Dependent Surveillance – Broadcast (ADS-B)	29
3.1.6	Multilateration (MLAT)	32
3.2	Information Services	33
3.2.1	Aircraft Communications Addressing and Reporting Sys- tem (ACARS)	33
3.2.2	Traffic Alert and Collision Avoidance System (TCAS) .	34
3.2.3	Flight Information System – Broadcast (FIS-B)	35
3.2.4	Traffic Information System – Broadcast (TIS-B)	35
3.3	Potential Future Technologies	35

Throughout this dissertation, we focus on the whole picture of aviation as found under Instrument Flight Rules (IFR), where navigation depends on electronic signals. IFR usually apply to large commercial aircraft, even though they are open to any aircraft with the necessary equipment, including general aviation (i.e., civil aviation that is not a scheduled air service or air transport for remuneration or hire).

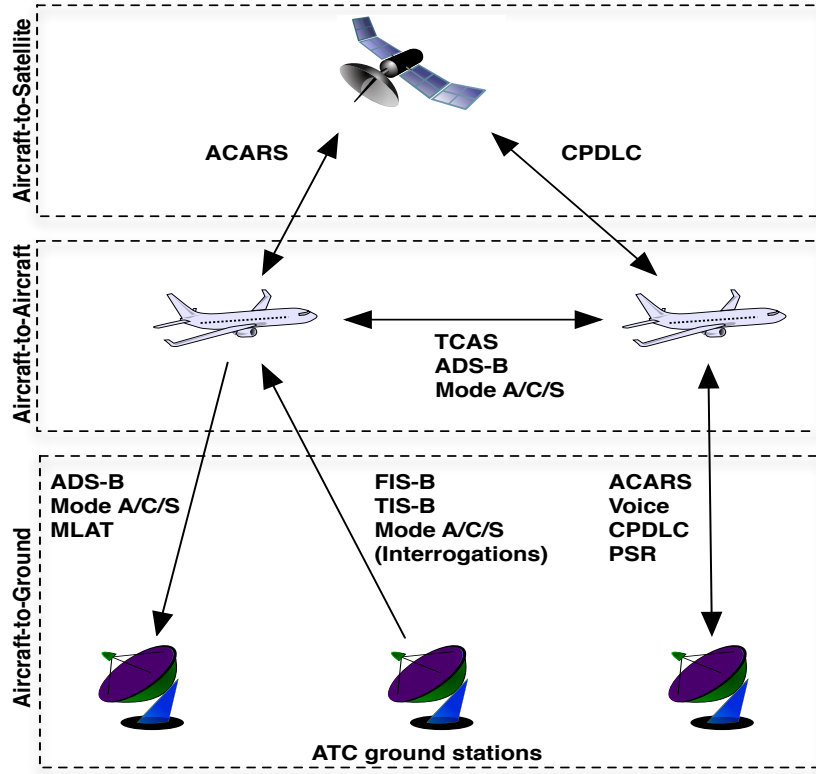


Figure 3.1: An overview of the wireless technologies used in air traffic communication, between ground stations, aircraft and satellites. The arrows indicate the direction of the communication for each protocol.

However, many of our findings also apply to flying under Visual Flight Rules (VFR). Under European Organisation for the Safety of Air Navigation (EUROCONTROL) and Federal Aviation Administration (FAA) regulations, aircraft under VFR need to equip fewer communication systems and enjoy more freedom in their choice compared to commercial aircraft, but also have fewer options in case of their failure.

Fig. 3.1 provides a comprehensive, high-level picture of currently employed wireless communication technologies in commercial aviation, focusing on the interactions between the technologies and their utilization during different flight phases. To aid the reader’s understanding throughout the dissertation, we have collected the most important acronyms in Table 3.1. We generally use broad definitions of *communications* and *protocols*, which can include analogue technologies as well as message-based protocols transmitting digital data.

In order to focus on the systems view, we have divided all technologies into

Abb.	Full Name
Air Traffic Control	
VHF	Voice (Very High Frequency)
PSR	Primary Surveillance Radar
SSR	Secondary Surveillance Radar (Mode A/C/S)
ADS-B	Automatic Dependent Surveillance-Broadcast
CPDLC	Controller–Pilot Data Link Communications
MLAT	Multilateration
Information Services	
ACARS	Aircraft Communications Addressing and Reporting System
TCAS	Traffic Alert and Collision Avoidance System
FIS-B	Flight Information System-Broadcast
TIS-B	Traffic Information System-Broadcast

Table 3.1: Short-handles and full names of aviation communication technologies, systematized into applications.

two categories according to their application: air traffic control and information services. ATC protocols are used to enable communication between controllers and pilots or their aircraft. They include VHF, PSR, SSR, ADS-B, MLAT and CPDLC, which are used during all flight phases, usually on line-of-sight frequencies, although the use of satellite communication is possible (e.g., CPDLC). Information services offer a more general platform for the exchange of data such as weather and traffic information: we discuss ACARS, TCAS, FIS-B, and TIS-B.

While there is considerable fragmentation among aviation systems all over the world, we aim to be as comprehensive and generally applicable as possible. Technologies and procedures related to military systems, both secret in nature and exclusive to a country’s air force, are out of the scope of this dissertation. We further appreciate that even some of the same systems differ across regions, and due to space limitations, we cannot address every exception. However, our findings are broadly applicable, as the underlying technologies and principles are the same.

In the following, we briefly introduce the wireless technologies used in aviation and their impact on the system as a whole. As the application of the technologies determines the consequences and severity of security breaches, we divide the technologies accordingly into two categories. *Air traffic control* comprises technologies which support air traffic services. This includes communication links between controllers and pilots, and technologies for monitoring air traffic. *Information*

services are technologies which provide information to pilots to improve their situational awareness (e.g., weather or traffic information).

In general, all technologies in the same category fulfil the same greater purpose and can be thought as backup of each other, hence the need to look at the system as a whole. However, some of them do offer irreplaceable functions not covered by any other technology, and their loss will lead to a degradation of expected service levels, depending on the equipment of airspace and aircraft.

3.1 Air Traffic Control

ATC protocols enable communication between controllers and pilots or their aircraft. They establish information about the aircraft's position and intent, and thus ensure the safety of the airspace. Table 3.2 details their technical characteristics.

3.1.1 Voice (VHF)

Voice communication [46] is the primary means of communication between ATC and the aircraft. It is used to transmit all ATC instructions (clearances) to the aircraft, which are acknowledged by the pilot, as well as pilots' reports and requests to ATC. Flight information services, weather reports, and airport information broadcasts can also be provided by voice communication. It is further used for operational communication between the airline operator and the aircraft, as far as the aircraft is in range of the operator's transmitter. Voice communication is conducted by analogue radio on VHF and HF (outside VHF range, e.g., over oceans) [47].

3.1.2 Controller Pilot Data Link Communications (CPDLC)

CPDLC is a message-based service offering an alternative to voice communication between ATC and pilot. ATC can use CPDLC via a terminal to send clearances or requests. The pilot can send requests and reports by selecting predefined phrases (e.g., REQUEST, WHEN CAN WE) or by using free text. CPDLC has great advantages over VHF: the number of acoustic misunderstandings is reduced, messages are saved for accountability, and it is easier, more efficient and safer to

	Voice	PSR	Mode A/C/S	ADS-B	CPDLC
Use	Communication ATC-Cockpit	Non-cooperative aircraft detection and positioning	Cooperative aircraft detection, positioning and data exchange	Broadcast aircraft data relevant for ATC and collision avoidance	Communication ATC-Cockpit
Type	Selective & Broadcast	Broadcast	Interrogation	Broadcast	Selective
Sender	Aircraft & Ground	Ground	Aircraft	Aircraft	Aircraft & Ground
Receiver	Aircraft & Ground	Original Sender	Aircraft & Ground	Aircraft & Ground	Aircraft & Ground
Frequency	3.4-23.35, 117.975-143.975, 225-400 MHz	1-2 & 2-4 GHz band	1030 & 1090 MHz	978 & 1090 MHz	VDL2: 136.975 MHz
Data Rate	Not applicable	Not applicable	1 Mbps (Mode S)	1 Mbps	30 kbps
Contents	Clearances, pilot requests, any other information	Pulses	A: squawk, C: altitude, S: similar to ADS-B but no position	ID, call sign, position, altitude, velocity, intent, type, etc.	Clearances, requests, weather, any other relevant information
Link Layer	Radio (amplitude modulation)	Pulse position modulation	Mode A/C/S	UAT / Mode S 1090ES	VDL / HFDDL / satcom
Data Source	Pilot & Controller	Radar	Aircraft	Aircraft	Several
Signal	Analogue	Analogue	Digital	Digital	VDL2+: digital
Adoption	In use	In use	In use	Parts of the world, in adoption	Parts of the world, in adoption
References	[46, 47]	[48]	[49]	[50–52]	[53]

Table 3.2: Detailed characteristics of air traffic control protocols.

transmit and receive long messages such as flight plan changes during flight. For example, VHF depends on the controller to catch a wrongly understood flight level instruction while he/she is also busy with several other aircraft. With CPDLC such mistakes can be eliminated and preliminary studies show that the communication demands on the pilot can be reduced by as much as 84% [54].¹

Some busy airports already employ CPDLC for automated clearance delivery and start up approval; in some European airspaces it is also used in-flight for minor tasks. Currently, CPDLC uses VHF Data Link Version 2 (VDL) [53] as its data

¹As unintended consequences may occur when reducing voice usage (the situational awareness of other aviation users on the channel may suffer for example, as they cannot overhear potentially safety-related conversations), further research is required to assess the concrete impact and feasibility of this move.

link. Coverage is provided by ground stations and satellites to ensure availability where required, even in oceanic regions. It has been successfully used for more than a decade in airspaces not covered by VHF as it offers easier communication (via satellite) compared to HF with its signal propagation difficulties. Without CPDLC, time-critical ATC clearances or pilot requests very often cannot reach the destination in due time, which forces pilots to deviate from ATC clearance without permission (for example to avoid bad weather), resulting in a safety problem. Indeed, the transit times of HF Data Link (HFDL) messages do not meet the requirements for some aircraft separation standards. Such issues are eliminated with CPDLC. As the technology is not yet mandatory, many short- and mid-range aircraft are not CPDLC equipped. Thus even in CPDLC-enabled airspaces, VHF remains the primary communication channel.

3.1.3 Primary Surveillance Radar (PSR)

PSR is the acronym for non-cooperative aircraft localization systems using radar. In aviation, these usually consist of a rotating antenna radiating a pulse position-modulated and highly directional electromagnetic beam on a low GHz band [48]. The pulses are reflected by targets and subsequently the bearing and round trip time are measured to get the target's position. PSR is *independent* of an aircraft's equipment and does not require the aircraft's cooperation, but it does depend on the reflecting area (surface material and size, distance and orientation of the aircraft in space). Due to this, and the fact that the signal has to travel two-way, very high radiation power is required (several MW). As the received information is carried by analogue signals, the system has to deal with numerous disturbing echoes caused by terrain, obstacles, weather, flocks of birds, or even cars on elevated roads. This makes complex signal processing necessary to extract the desired information.

In military airspace surveillance PSR is strictly required as it is crucial to detect uncooperative aircraft with intentionally non-working transponders. In civil ATC, however, PSR is used merely to detect aircraft with rare transponder failures and not as standard backup. Neither identification nor altitude are provided by PSR;

the tracker software system uses PSR solely to verify and improve the quality of targets obtained by other sensors.

3.1.4 Secondary Surveillance Radar (SSR)

The transponder modes A, C, and S (short: Mode A/C/S) form part of the Secondary Surveillance Radar [49]. This *cooperative* technology provides more target information on ATC radar screens compared to PSR, which only offers an unidentified target position without further supporting data. SSR ground stations broadcast interrogations of aircraft transponders, which reply with the desired information. SSR uses digital messages with different frequencies and modulations for the interrogation (1030 MHz) and the reply (1090 MHz).

The reply is also used to locate the aircraft's position using the antenna's bearing and the message round trip time. In this digital process, a radiation power of about 1 kW is sufficient, much lower compared to PSR.

The older Modes A and C (which report ID and altitude, respectively) are being substituted by Mode S, which supports selective interrogations of single aircraft instead of broadcast requests to all aircraft in range. This feature is supposed to relieve the saturated 1090 MHz reply channel, currently suffering from severe message loss (as discussed in [24]). Mode S also offers a worldwide unique transponder ID and more message formats with information on, for example, aircraft intent or autopilot modes. Note that it does not transmit the aircraft's position, however, which must be obtained by separate means (e.g., PSR, MLAT or ADS-B).

3.1.5 Automatic Dependent Surveillance – Broadcast (ADS-B)

The FAA and EUROCONTROL named ADS-B [50] as the satellite-based successor of both PSR and SSR. At its introduction, ADS-B presented a completely new paradigm for air traffic control. Contrary to before, every ADS-B participant retrieves their own position and velocity by using an on-board GNSS receiver, rendering it a *dependent* technology (see Fig. 3.2 for an illustration).

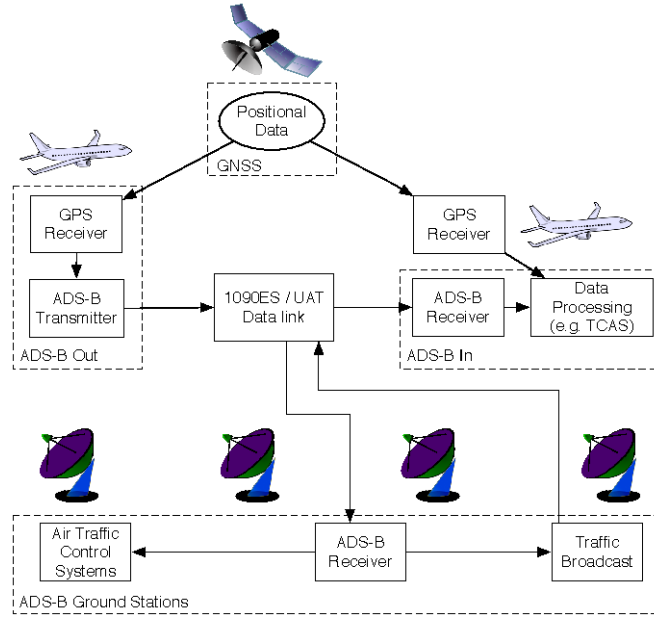


Figure 3.2: Overview of the ADS-B system architecture. Aircraft receive positional data that is transmitted via the ADS-B Out subsystem over the 1090ES or the UAT data link. It is then received and processed by ground stations and by other aircraft via the ADS-B In subsystem.

The ADS-B data link is then used by aircraft to automatically and continually broadcast their own ID, position and velocity as well as further information such as intent or urgency codes. These broadcasts happen twice a second in case of position and velocity, and once every 5 s for identification. This subsystem, ADS-B Out, is mandated for use by 2020 in US and European airspaces, promising to improve on location accuracy and decrease system costs by replacing existing radar systems [55]. Mandates for the installation of a receiving subsystem ADS-B In on commercial aircraft have not been set yet, but such systems are available and data is used by flight information services today (see Section 3.2). ADS-B’s importance in the future and its close links to current SSR systems warrant a more detailed look in this section.

Two competing ADS-B data link standards exist: Universal Access Transceiver (UAT) and 1090 MHz Extended Squitter (1090ES). UAT has been created specifically for use with aviation services such as ADS-B, utilizing the 978MHz frequency with a bandwidth of 1Mbps. Since UAT requires fitting new hardware, as opposed to 1090ES, it is currently only used for general aviation in EUROCONTROL and

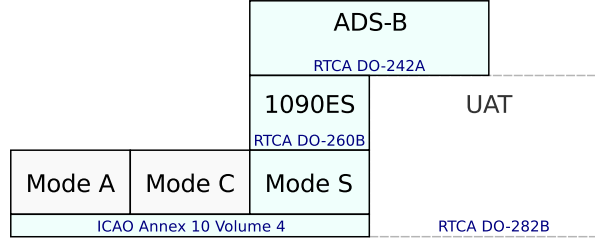


Figure 3.3: ADS-B protocol hierarchy [13]. The 1090 MHz Extended Squitter is based on the traditional Mode S system and provides the data link for ADS-B in commercial aviation. UAT is a new development but currently only mandated for general aviation in the US.

FAA-mandated airspaces. Scheduled airliners, on the other hand, employ SSR Mode S with Extended Squitter, a combination of ADS-B and traditional Mode S known as 1090ES (see Fig. 3.3). In other words, the ADS-B function has been integrated into traditional Mode S transponders.

In this work, we focus on the commercially used 1090ES data link. The complete overview over the ADS-B protocol can be found in the specification documents [56–58] while various other works give succinct, higher level descriptions of the protocol (e.g. [9, 12, 13]).

As the name suggests, the 1090ES data link predominantly uses the same 1090MHz frequency that Mode S uses for communications by aircraft. Fig. 3.4 provides a graphical view of a 1090ES transmission, which starts off with a preamble of two synchronization pulses. The data block is then transmitted by utilizing pulse position modulation (PPM). With every time slot being $1\mu\text{s}$ long, a bit is indicated by sending a $0.5\mu\text{s}$ pulse either in the first half of the slot (1-bit) or in the second half (0-bit). It is important to note that PPM is very sensitive to reflected signals and multipath dispersion, a fact that can play a major role in security and protocol considerations.²

There are two different possible message lengths specified in Mode S and its dependent protocols, 56 bit and 112 bit [56]. ADS-B exclusively uses the longer format. The downlink format field DF (alternatively UF for uplink messages) assigns the type of the message. 1090ES uses a multi-purpose format as shown in Fig. 3.4.

²See [59] for more information on PPM and multipath.

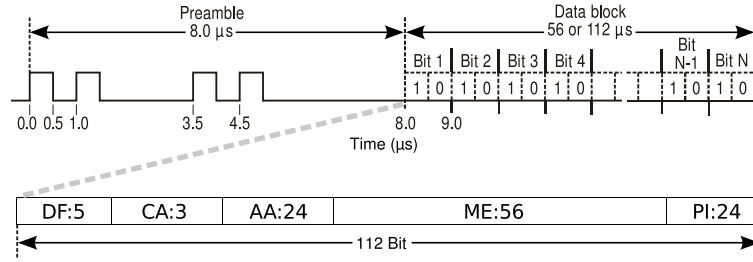


Figure 3.4: 1090 ES Data Link [13].

When set to 17, it indicates that the message is an extended squitter, enabling the transmission of 56 arbitrary bits in the ME field. The CA field indicates information about the capabilities of the employed transponder, while the 24 bit AA field carries the unique International Civil Aviation Organization (ICAO) aircraft address which enables aircraft identification. Finally, the PI-field provides a 24 bit CRC to detect and correct possible transmission errors. It is possible for recipients to correct up to 5 bit errors in 1090ES messages using a fixed generator polynomial of degree 24.

Overall, ADS-B exemplifies the move to cooperative data communication networks in the next ATC protocol generation. While as of now, its impact is still limited compared to VHF and SSR, its deployment is already considerable with more than 70% of all aircraft supporting the protocol [60]. Its crucial position in the next generation of ATC protocols informs our choice to make it the centre of our investigations in Chapter 7 and 8 along with Mode S and multilateration.

3.1.6 Multilateration (MLAT)

Multilateration, or hyperbolic positioning, has been successfully employed for decades in military and civil applications, not limited to navigation. It differs from other ATC aids as it is not a separate protocol but exploits the time differences of arrival of signals received from aircraft independently through other protocols such as SSR or ADS-B [61]. Using the reception times of four or more receivers, it is a purely geometric task to find the origin of the signal and thus the position of the sender/aircraft. Consequently, MLAT is a dependent surveillance technology as it requires other cooperative technologies to function. We will elaborate more on multilateration in Section 6.2.1.

	ACARS	TCAS	FIS-B	TIS-B
Use	Dispatch, operations, engineering, maintenance...	Collision avoidance	Flight information	Traffic information
Type	Broadcast	Interrogation	Broadcast	Broadcast
Sender	Aircraft & Ground	Aircraft	Ground Radar	Ground Radar
Receiver	Aircraft & Ground	Aircraft & Ground	Aircraft	Aircraft
Frequency	129.125-136.900 MHz	1030 & 1090 MHz	978 MHz	978 & 1090 MHz
Data Rate	2400 bps	1 Mbps	1 Mbps	1 Mbps
Contents	Position, weather, fuel & engine information, delays, maintenance reports...	Altitude, relative position (derived from round-trip time), transponder status	Weather text & graphics, notices to airmen, terminal information	Non-ADS-B equipped aircraft
Link Layer	Several	Mode S & 1090ES	UAT	UAT & 1090ES
Data Source	Various	ADS-B & Mode S	FIS-B Provider	Radar station
Signal	Digital	Digital	Digital	Digital
Adoption	In use	In use	Parts of the US	Parts of the US
References	[62]	[63]	[64]	[50]

Table 3.3: Detailed characteristics of information services protocols.

3.2 Information Services

Information services are air traffic systems that provide a more general platform for the exchange of information, from traffic and weather information to free text. These protocols use a variety of sources and supply the backbone for a wide array of use cases. Table 3.3 provides the technical details of the discussed technologies. Exploiting them can lead to a range of potential issues, from mere privacy problems to serious disasters where collision avoidance is concerned. Contrary to ATC technologies, they are not typically set up as redundant, although they can theoretically handle some of the same functions.

3.2.1 Aircraft Communications Addressing and Reporting System (ACARS)

ACARS [62] is a digital data link system developed in the 1970s for general communication between aircraft and ground stations. ACARS messages are used for ATC, flight information and alerting, and also by airlines to communicate with their aircraft. It is used in all flight phases, for services as varied as dispatch, operations,

engineering, catering, and customer service. ACARS transmits: safety-critical data such as aircraft weight, fuel, engine data, and weather reports; privacy-related information about passengers or catering requests; and information critical to business operations such as gate assignments, crew schedules, and flight plan updates.

ACARS offers five data links, depending on the aircraft's equipment: VHF, Inmarsat satcom, Iridium satellite, VDL Version 2, and High Frequency Data Link. The messages are character-oriented and only accept valid ASCII symbols [3].

3.2.2 Traffic Alert and Collision Avoidance System (TCAS)

TCAS [63] is an airborne system for collision avoidance independent of ground-based ATC. The current version TCAS II uses the available information (i.e., identity, altitude) from ATC protocols such as Mode C and S to provide a traffic surveillance display of all equipped aircraft in the proximity [3]. It determines the relative velocity and distance of nearby transponder-equipped aircraft through interrogation. When a broadcast Mode S message is received, the transmitted ID is added to a list of aircraft that is then interrogated at about 1 Hz. With the reply, distance and altitude of the interrogated aircraft are determined. ADS-B messages will be incorporated into TCAS in the future. Currently, the state of the art consists of so-called hybrid surveillance systems, which use ADS-B information to reduce the interrogation rates of TCAS systems. This is achieved by identifying aircraft deemed at a safe distance and not interrogating them until they come closer. Full use of ADS-B messages would make the interrogation step unnecessary. Based on the obtained relative velocities and positions, potential threats are identified and presented to the pilot as a Traffic Advisory (TA). When proximity thresholds are violated, TCAS issues a Resolution Advisory (RA) and proposes an avoidance manoeuvre to eliminate the threat (in the latter case TCAS can be classified as ATC protocol, too). Advisories are also broadcast for the attention of air traffic controllers.

3.2.3 Flight Information System – Broadcast (FIS–B)

FIS-B [64] is a general flight information service that requires aircraft to be equipped with ADS-B In. It uses the Universal Access Transceiver [51, 52] data link on 978 MHz which offers more flexibility through larger ADS-B messages. It is in use in parts of the US, with a wider adoption possible in the future. FIS-B provides data about airspace restrictions or meteorological advisories. The data is supplied by the FAA for general aviation below 24,000 ft [65].

3.2.4 Traffic Information System – Broadcast (TIS–B)

TIS-B [50] is another ground-based traffic information service used in the US that broadcasts additional data about aircraft that are not equipped with ADS-B transponders. TIS-B is used for increased situational awareness and collision avoidance. The system uses the same frequencies as ADS-B and the same message format and provides users with a full surveillance picture as seen by ground radar, i.e., the broadcast data can be compiled from all available ATC sources such as PSR, SSR, ADS-B, or MLAT.

3.3 Potential Future Technologies

Apart from the technologies introduced in this chapter and considered throughout this dissertation, the international aviation authorities ICAO, EUROCONTROL, and the FAA have started planning for further upgrades of the current communications systems and are seeking to develop new data links. Specifically, L-band Digital Aeronautical Communications System (L-DACS) and Aeronautical Mobile Airport Communications System (AeroMACS) are supposed to replace the current VHF system. Since these systems can provide much higher data throughput comparing to the existing data links, some of the applications currently provided by other technologies could also utilize these new technologies one day. Thankfully, L-DACS and AeroMACS have begun to at least consider the issue

of wireless security and some corresponding designs are already included by the specifications or will be in the future.

Unfortunately, L-DACS is still in the very early specification phase and in line with typical technological cycles in aviation will not be deployed before the 2030s [66]. Furthermore, since its specification is not finished – many parts are up in the air and there are still competing proposals – they could strongly benefit from an immediately increased awareness about security concerns in aviation, which we aim to provide with our work.

AeroMACS takes the form of a profile of IEEE 802.16-2009 [67], known as WiMAX. It intends to provide a surface data link for use at the airport, allowing ATC, airlines, and airports to communicate with the aircraft [68]. It has line-of-sight range of up to 3 km per cell and uses commodity radios to communicate. While the current standards include cryptography, making it a serious step forward, AeroMACS will not solve the security problems currently found in aviation. Besides the prevalent issue of long deployment time frames (the beginning of deployment is not projected before the middle of the next decade), many security questions such as the protection of management frames are still undecided [69]. Most importantly, AeroMACS will only be able to replace current data links on the ground and in the immediate vicinity of an airport, leaving the vast amount of air traffic communication unprotected.

AeroMACS is further along in the development cycle compared to L-DACS, with test deployments going on at some airports around the world. However, at the time of writing, many of the necessary avionics standards and specifications were still in the planning phase [68].

Considering these facts, we exclude both L-DACS and AeroMACS from the scope of this dissertation; in the next chapter we only examine the security of all currently deployed technologies. However, we strongly believe they should see input from the security community as soon as possible to avoid the costly deployment of insecure protocols.

Es hört doch jeder nur, was er versteht.

Everyone hears only what he understands.

— Johann Wolfgang von Goethe [70]

4

Examining Current Vulnerabilities in Aviation Technologies

Contents

4.1	Wireless Attack Vectors	37
4.2	Attacks on Air Traffic Control Technologies	40
4.3	Attacks on Information Systems	48
4.4	Communication in Military ATC	50
4.5	Summary of Vulnerabilities and Attacks	51
4.6	Mitigating Factors	52
4.7	Reasons for the Current State of Aviation Security . .	55

In this chapter, we analyse the known security flaws in current wireless air traffic communication technologies. We use existing literature and publicly available standard documents as sources for our analysis and introduce new potential attacks on these unsecured technologies where appropriate. Lastly, we examine the reasons for the current state of aviation security.

4.1 Wireless Attack Vectors

Contrary to wired networks, there are no practical obstacles, such as buildings or security guards, for an attacker trying to access a wireless network, making access control mechanisms very challenging. This section outlines the classes of

vulnerabilities inherently stemming from the broadcast nature of radio frequency (RF) communication when RF is used without appropriate security measures. Concretely, these are eavesdropping, jamming, message injection, message deletion, and message modification. Attacks exploiting all of these vulnerabilities have become practical and accessible to a wide range of people relatively recently with the proliferation of SDRs.

Eavesdropping

The most straightforward attack is the act of listening in on unsecured broadcast transmissions. As protocols send unsecured messages over an inherently broadcast medium, the possibility of eavesdropping is not surprising and well known. Many non-adversarial services use this obvious privacy concern, e.g., to visualize air-traffic on the Internet.¹ Yet, eavesdropping also serves as preparation for more sophisticated active attacks such as the injection of real, previously saved, messages (a so-called replay attack), or the tracking of private aircraft as discussed in Chapter 8.

Eavesdropping is not only difficult to prevent without confidentiality provided through encryption, but is also practically impossible to detect. A small number of countries (such as the United Kingdom) have long-standing and very broad laws against listening in on unencrypted broadcast traffic which is not intended for the recipient.² However, since it is extremely difficult to even detect an eavesdropper, the technical reality renders such legal and regulatory approaches insufficient.

Jamming

Almost as simple as eavesdropping is the jamming attack, where a single node (either a ground station or an aircraft) or an area with multiple participants is

¹Prominent examples are <http://flightradar24.com> and <http://flightaware.com> among many others.

²Section 48 of the Wireless Telegraphy Act of 2006 states that (1) “A person commits an offence if, otherwise than under the authority of a designated person— (a) he uses wireless telegraphy apparatus with intent to obtain information as to the contents, sender or addressee of a message (whether sent by means of wireless telegraphy or not) of which neither he nor a person on whose behalf he is acting is an intended recipient, or (b) he discloses information as to the contents, sender or addressee of such a message.”

effectively disabled from sending/receiving messages by an adversary sending with sufficiently high power on the frequency of the attacked technology (e.g., 1090 MHz for SSR/ADS-B and dependent protocols). It has also been proven that it is feasible to conduct reactive jamming in real time, targeting only packets which are already in the air as assessed in [71]. While jamming is a problem that is difficult to solve in all of wireless communication, the impact is severe in aviation due to the system's inherent wide open spaces which are impossible to control as well as the importance and criticality of the transmitted data. Besides digital communications systems such as ADS-B or SSR, primary radar may also be a target for jamming attacks. However, due to rotating antennas and a higher transmission power, typical PSRs are more difficult to jam than receivers for modern ATC protocols, especially for non-military grade attackers. The jamming of ATC frequencies, like all other active tampering with RF communication, is prohibited by law and regulations in most countries. However, while it is possible to locate an active offender, this is naturally only a necessary but not sufficient protection for the ATC system.

Message Injection

Slightly more sophisticated is the injection of non-legitimate messages into the air traffic communication system. Since no authentication measures are implemented at the data link layer, there is no hurdle for an attacker to build a transmitter that is able to produce correctly modulated and formatted messages. Schäfer et al. [13] provide a detailed example on how to conduct an attack on ADS-B with limited knowledge and very cheap and simple technological means which have been easily and widely available for some time. As another direct consequence of missing authentication schemes, a node can deny having broadcasted any data (legitimate or not) and/or claim to have received conflicting data, making any kind of accountability impossible.

Message Deletion

Legitimate messages can be physically “deleted” from the wireless medium by utilizing destructive or constructive interference. Destructive interference means transmitting the inverse of the signal broadcast by a legitimate sender. Due to superposition, the resulting signal should be erased or at least highly attenuated. In practice, however, this approach has very precise and complex timing requirements, making it extremely challenging.

Constructive interference on the other hand does not require synchronization but simply causes a large enough number of bit errors to destroy the message. Since, for example, Mode S extended squitters’ CRC can correct a maximum of 5 bit errors per message, if a message exceeds this threshold, the receiver will drop it as corrupted. While the message is effectively destroyed, the receiver might at least be able to verify that it has been sent, depending on the implementation and the circumstances. In any case, message deletion is more subtle than complete jamming of the frequency.

Message Modification

Modifying messages on the physical layer during transmission is typically done using two different approaches: overshadowing and bit-flipping. Overshadowing means that the attacker sends a high-powered signal to replace part or all of the target message. With bit-flipping, the attacker superimposes the signal, converting any number of bits from 1 to 0 (or the other way around). In both cases arbitrary data can be injected without the knowledge of any of the participants. This effect can also be achieved by combining message deletion and injection, but physical layer message modification can in some cases be regarded as even more sinister than the injection of a completely new message, since the manipulated message was originally legitimate. The feasibility of such message manipulation has been shown in [72] and [73].

4.2 Attacks on Air Traffic Control Technologies

ATC is the backbone of safe aviation, and its technologies are imperative for the smooth, safe and efficient functioning of the world-wide aviation system.




Figure 4.1: Denial of service due to an overload of aircraft on a controller's screen [13].

Manipulating data used by ATC can have severe consequences, such as loss of situational awareness and denial of service (DoS). In this section, we will discuss potential attacks on ATC systems based on wireless attack vectors.

Fig. 4.1 provides an illustration of what such an attack could look like on the part of the controller. While some consequences of attacks on ATC seem obvious, we investigate the outcomes further in Chapter 5.

VHF

While analogue voice communication via VHF is a highly safety-critical technology, it does not have an impact on the radar screen as seen by a controller. Successful VHF communication depends on the correct understanding of the message by the communication partners, meaning a high quality signal must be ensured. While it is not possible to attack VHF in the same sense as a digital protocol, simultaneous use of the frequency also leads to a partial or full DoS in practice, and/or to confusion for pilots and controllers when attackers pose as legitimate users [74].

Despite the fact that VHF employs amplitude modulation (AM), which allows reception of multiple channels on the same frequency, it is difficult to maintain service with an attacker dedicated to disturbing the intended communication. Authentication procedures for VHF are available for military flights only, in case the pilot insists; as they are time- and capacity-consuming, they are not applied for civil flights. When CPDLC is not available, as in most regions currently, or the aircraft is not equipped with it, there is no backup protocol for VHF. Consequently, losing this main communication layer within highly-populated airspaces results in a severe threat.

There are many reported incidents with spoofed voice communication in the aviation literature and on the Internet [74], and recent works also discuss the urgent need to improve the security of VHF [74, 75]. While even knowledgeable intruders could be detected through changes in signal or voice levels, this is obviously not a secure solution. Furthermore, attackers could effectively disable VHF (e.g., by jamming) and make aircraft rely on backup systems using unauthenticated digital data links (e.g., CPDLC), where manipulation is even harder to detect, and rendering the redundancy void.

CPDLC

Compared with VHF, relying on unauthenticated data links such as the one provided by CPDLC becomes a larger problem as attack detection is typically more difficult in digital networks compared to the analogue voice technology.

Attacks on CPDLC's availability are less critical currently, as CPDLC is still used as a secondary communication layer. However, protocol attacks such as message manipulation or injection are severe when undetected, as clearances and other flight safety-related information are transmitted using CPDLC. As there is no authentication, it is trivial to eavesdrop on or spoof clearances and execute replay and message alteration attacks as discussed further in [76].

Impersonation of aircraft is easily possible: to login to the responsible Area Control Centre, the pilot simply puts the correct Location Indicator (e.g., SBAO

for ATLANTICO, the control centre in Recife, Brasil) into the terminal. After a handshake, the user is successfully logged in. Likewise, an attacker can claim the identity of an ATC unit and send instructions to an aircraft causing the pilot to perform unnecessary, dangerous manoeuvres or increasing ATC workload through additional inquiries. Considering the plans of many aviation authorities to shift more responsibilities towards CPDLC in the future, CPDLC should be high on the list of protocols to be secured.

PSR

As PSR systems use a signal-based detection approach, they are not subject to protocol attacks such as message injection. However, jamming on any of the operational frequencies is possible [77], although, due to high power requirements, this remains in the realm of military electronic warfare. Normally, missing PSR information (caused by jamming) does not impact controllers as the main target information (position, identification, altitude, intent) is provided separately. While military PSR can offer security measures such as frequency hopping or modulation schemes, these are unavailable in civil aviation. Similar to other sensor systems, PSR may also be vulnerable to attacks on its timebase (e.g., GPS).

Despite these potential vulnerabilities, PSR can be considered relatively secure compared to other technologies. Thus, it is safety-relevant that PSR belongs to the oldest ATC technologies and is phased out in favour of modern data communication protocols using more accurate satellite systems. This could potentially reduce long-term security as reliance on unauthenticated, digital, and dependent ATC technologies increases.

SSR

With the publication of Mode S implementations for SDRs on the Internet (e.g., *dump1090*³), a somewhat knowledgeable attacker can exercise full control over the communication channel, i.e., by modifying, jamming, or injecting Mode A/C/S

³The currently most advanced fork of *dump1090* can be found at <https://github.com/MalcolmRobb/dump1090>.

messages into ATC systems, he can create a fully distorted picture of the airspace as seen by ATC.

Every Mode S message carries an identifier which can be replaced with an arbitrary one. Using a known and trusted aircraft identifier may, for example, reduce the likelihood for detection compared to an unknown and unexpected object on the radar. Mode S also offers special emergency codes selected by the pilot (7500 for hijacking, 7600 for lost communications, and 7700 for an emergency) and injecting these can cause immediate ATC inquiries. While this happens occasionally through transponder failures and wrong settings [4] and thus procedures to deal with such an occurrence are in place, this feature can actively be used to create confusion at a busy ground station.

As Mode A/C/S is the only source for all necessary information displayed on ATC radar screens, manipulation or jamming is a severe threat since no equivalent backup exists. Mode S messages and its data link are also used by other ATC systems, which consequently inherit its vulnerabilities as we will discuss in the relevant sections.

A vulnerability that is specific to Mode S is the amplification attack. Exploiting the interrogative nature of Mode S, an attacker can cause large-scale interference on the 1090 MHz channel without sending on the target frequency (but on the 1030 MHz interrogation frequency instead). Interrogations are limited to a maximum of 250/s now [78], but these restrictions are placed on the interrogators, not on the Mode S transponders in aircraft. As Mode S messages are unauthenticated, a malicious sender can easily circumvent these measures and use non-selective interrogations to amplify her sending power and frequency. By changing her own identifier code, the attacker can make all receiving aircraft answer the interrogations, increasing the range and capability of the interference attack manifold. While the aircraft continue to send useful information, the interference level in even moderately busy airspaces would quickly cause a partial DoS as important data gets lost.

As a concrete example, the FAA uses a stochastic probability of 0.25^4 to calculate the number of replies to a given Mode S interrogation broadcast (“all-call”) [78].

⁴Accounting for interference on the 1030 MHz channel and the time when a transponder is already busy.

Thus, an attacker can create an amplified response of $x * y * 0.25$ messages per second, where x is the number of receiving aircraft and y the number of all-calls per second. Based on data from our own Mode S receivers, 200 aircraft are easily in range within a normal airspace. Thus, an attacker can use only 100 messages to create $200 * 100 * 0.25 = 5000$ additional Mode S messages, adding to the significant existing interference experienced by all ATC receivers.⁵ If done with consideration, such an attacker is also more difficult to detect from the ground. In addition, a recent incident showed that some transponders are susceptible to over-interrogation. Sending many interrogation calls can lead to overheating and result in a complete loss of the target from ATC displays due to a full DoS of the transponders. Investigations revealed that transponders transmit at rates beyond their requirements and design limits, worsening amplification attacks [80].

ADS-B

Securing ADS-B communication was still not a very high priority when it was specified to be the new standard in civilian secondary surveillance during the early 1990s. Neither the official standards of the Radio Technical Commission for Aeronautics (RTCA) [56–58] nor other requirements documents [81, 82] mention security in this context. However, the problems in ADS-B have been well known for a long time, mostly because they are relatively obvious to the interested researcher. On the Internet, warnings are traceable back to as early as 1999.⁶ In 2012, weak ADS-B security finally got some broad attention in the mainstream press [6, 7, 15–17] due to presentations at hacking conferences [10]. On the academic side, Costin and Francillon [12] as well as Schäfer et al. [13] analysed ADS-B security, too, focusing on the ease of exploiting ADS-B with cheap hard- and software.

As the commercially used ADS-B data link 1090ES is based on unauthenticated Mode S, it suffers from the same passive and active attacks described in the previous section. For example, it is possible to selectively jam all ADS-B messages of a single aircraft, which would make it vanish from the ADS-B channel. This feat is

⁵The signal load is approaching 100% in some scenarios [79].

⁶For example, <http://www.airsport-corp.com/adsb2.htm>.

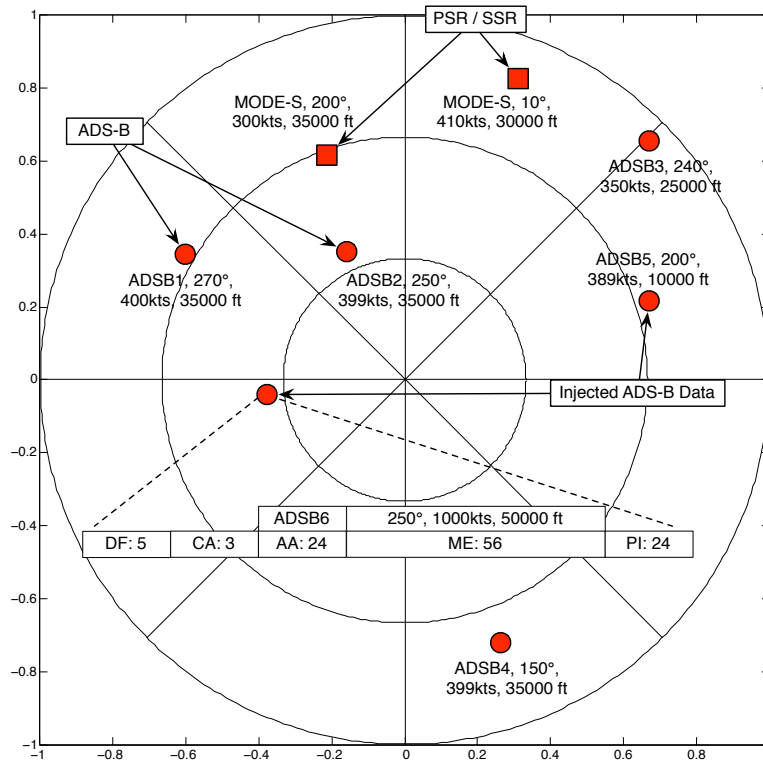


Figure 4.2: Result of an ATC injection attack. The radar screen shows legitimate ADS-B equipped aircraft and aircraft detected by PSR/SSR surveillance alongside aircraft injected by an adversary. On the screen, the injected aircraft (ADSB5, ADSB6) are indistinguishable from real ones.

much more easily accomplished with ADS-B's regular broadcasts compared to the Mode S interrogation system using directional antennas and much more frequent, irregular, and bursty interrogations.

Furthermore, as ADS-B additionally broadcasts the position of aircraft, this opens some new attack vectors which only require standard off-the-shelf hardware to execute as demonstrated in [12, 13]. Trivially injected ADS-B messages claiming to be non-existing aircraft are impossible to tell apart from authentic ones on the link layer. Without security added at the application layer, they could end up on a controller's screen as shown in Fig. 4.2. Other attacks virtually modify the trajectory of an aircraft by selectively jamming an aircraft's messages and replacing them with modified data. This causes discrepancies between the real position and the one received by ATC [13]. This is a worrying prospect, as ADS-B is set to be the main ATC protocol in the long term, with the FAA considering elimination

of Mode A/C/S transponders at some point in the future.⁷

MLAT

In theory, MLAT as a technology does not rely on the *contents* of the received messages but (similar to PSR) works purely on the signal level. This provides a strong theoretical security advantage because it does not suffer from compromised message integrity. Even if the contents of, e.g., an ADS-B message are wrong, the location of the sender can still be identified. Thus, MLAT offers additional security based on physical layer properties (here the propagation speed of electromagnetic waves) which are difficult to cheat.

However, in practical ATC implementations this assumption fails. Typical MLAT systems heavily rely on fusing the location data obtained from the signals with SSR message contents to display identification and altitude of the targets, leaving the system as a whole as vulnerable as Mode A/C/S or ADS-B.

Independent of this problem, a well-coordinated and synchronized attacker could still manipulate a message's time of arrival at the distributed receivers of an MLAT system such that using these signals for location estimation would result in a position of the attacker's choice [83]. This is shown in [31] for the similar case of spoofing a group of distributed GPS receivers. The authors find that even though more receivers severely restrict the possible attacker placement, attacks are generally feasible but harder and thus less likely.

Despite these drawbacks, MLAT can still offer improved security over the sole use of SSR/ADS-B and is thus a favoured SSR backup solution in aviation [84] and the academic community [85]. Unfortunately, it is very expensive to deploy (in part due to its susceptibility to multipath propagation) and thus not a preferred option in all environments [27].

⁷See <https://www.faa.gov/nextgen/programs/adsb/faq/>.

4.3 Attacks on Information Systems

Information services provide information other than those used for air traffic control. Attacks on these technologies can affect the situational awareness of pilots about their surroundings, for example impacting their knowledge of weather or other aircraft nearby. As there is such a wide variety of use cases, not all potential exploits can be discussed here but we will focus on some illustrative examples.

ACARS

ACARS security issues have long been considered, shown for example by a 2001 military study: “The military is uncomfortable with the ease at which eavesdropping on ACARS can be achieved” [86]. Today, ACARS eavesdropping has become much more widely accessible as SDR-based decoders are available on the Internet,⁸ although the satellite data link is physically more difficult to attack than the VHF-based data link.

To counter this, the ACARS message security (AMS) standard was developed [87]. It provides end-to-end encryption using ECDSA with SHA256 for digital signatures and offers message authentication codes with HMAC-SHA256 of a default length of 32 bits. AMS currently enjoys very little adoption as only few airlines (e.g., Lufthansa [88]) consider securing ACARS transmissions. Others (e.g., Ryanair [89]) forego ACARS completely and use airport-based mobile phone technologies.

Furthermore, airlines use their own semantics for data packets transmitted by ACARS, providing some security by obscurity.⁹ Due to this wide variety in implementations and applications supported by ACARS, discussing all potential attack vectors is not possible here. Existing examples in the literature, however, include the potential exploitation of soft- and hardware using the interface offered by ACARS [11] or the issuing of wrong ATC instructions [86]. On top of this, one can imagine a serious impact on business intelligence and personal privacy

⁸acarsd, for example, <http://acarsd.org>.

⁹A US military presentation [86] considers binary ACARS messages less vulnerable as they are not human-readable.

when passenger lists, crew information, or engine data are transmitted in clear text via ACARS, as studied in [90].

TCAS

Due to its crucial function, TCAS has some of the highest safety-related consequences, particularly for commercial aviation. As TCAS is based on data and message formats of Mode A/C/S (and ADS-B, which it will integrate in newer versions), it suffers from the unauthenticated nature of these protocols as described above. The potential attack vectors of TCAS differ, however, as the main targets are aircraft, not ground stations. Attacking aircraft at cruising altitude from the ground requires a strong transceiver, making an attack from closer range or within the aircraft more likely.

One concern is an attacker who falsifies the data used by TCAS to be aware of the surveillance picture around an aircraft. To do this, answers to Mode S interrogations by TCAS are spoofed using wrong information and message timings. The attacked TCAS system will classify such “ghost aircraft” as a threat and initiate an RA to which the pilot needs to respond. As even real advisories can lead to serious incidents,¹⁰ a loss of situational awareness caused by multiple fake RAs is very possible.

Another attack focuses on the RA messages themselves. Issuing fake advisories to ATC ground stations using Mode S RA reports is an easy way to cause a partial loss of situational awareness and control. Since controllers are prohibited from interfering with RAs, effective control of the air traffic is strongly inhibited. TCAS provides the biggest contrast of all protocols between the relative ease of such attacks, and the potentially severe impact of attacking a system in charge of collision avoidance.

FIS-B

FIS-B is based on the unauthenticated ADS-B data link UAT. It is thus trivial to manipulate or replay the broadcast messages sent out by ground station service

¹⁰See, e.g., http://www.skybrary.aero/index.php/Misinterpretation_of_TCAS_RA_Aural_Annunciation_Messages.

providers [91] and change the information available to a pilot. The payload encoding of FIS-B is available at <http://fpr.tc.faa.gov>, requiring a non-verified registration. Decoders of weather data sent over FIS-B such as METAR (Meteorological Aviation Reports) are widely available on the Internet.¹¹ Thus, it is not difficult to send out forged broadcasts of weather reports or severe weather forecast alerts, even raster scan pictures, by simply following the standard specifications.

TIS-B

TIS-B uses both available ADS-B data links, UAT and 1090 ES. Again, as both are unauthenticated, it is trivial to manipulate or replay the broadcast messages sent out by ground station service providers with the same means as explained above. Forged TIS-B messages broadcast to airborne targets can advertise non-existing aircraft or manipulate information (e.g., position) about aircraft without their own beacon transponder [65]. Both FIS-B and TIS-B are currently of limited safety impact outside areas busy with general aviation in the US, giving their vulnerabilities a lower priority.

4.4 Communication in Military ATC

There is undoubtedly a much stronger need and motivation to implement stringent ATC security in a military context. Though it is not within the scope of this dissertation, anything in practical use or development by airforces or navies could naturally be of interest for civil security solutions as well. There are various standards developed by the US and NATO military, the most relevant among them are the cryptographically secured Mode 4 and Mode 5 as defined in the NATO Standardization Agreement (STANAG) 4193. Mode 4, which employs a 3-pulse reply to a challenge, has been in use for decades and according to the forecasts of the NATO Minimum Military Requirements is to be superseded by Mode 5 by 2020 (full operational capability) [92].

¹¹There are some providers using proprietary data links, e.g., Weather Services International or Honeywell, the latter offers encryption.

Format #	DF					
19	1 0011	AF : 3	Military Application : 104			Military Extended Squitter

Figure 4.3: DF19 data format [95]

While the legacy Mode 4 indeed only allows aircraft to respond to challenges, Mode 5 adopts the ADS-B broadcast capability. Thus, participants can announce their presence without a prior query, which is very useful for identification, friend or foe (IFF) [93]. On the security side, Mode 5 uses proprietary hardware and encryption algorithms with a black key concept.¹² The signal modulation is done via spread spectrum and operation requires a platform identification number (PIN) [94]. Mode 5 hardware is equipped with a unique identifier that also informs about national origin. It offers two different levels: Level 1 is the interrogation response mode, providing time, position and identification based on both GPS and other instruments. Level 2 is the broadcast mode and entirely based on GPS. There is little available detail on the security mechanisms including the applied cryptography of Modes 4 and 5 as this information is classified.

The ADS-B specification itself also mentions the message types/downlink formats Military Extended Squitter (DF19, see Fig. 4.3) as well as Military Use Only (DF22) without detailing them further, although it is known for example that DF19 makes ample use of bursts instead of regular beacon messages [95]. Despite incomplete or unavailable information on the performance of Mode 5 compared to civil technologies, it is safe to say that cost, scalability and ease of use of the known aspects of the system are prohibitive to widespread use in commercial ATC. Spread spectrum techniques and cryptography could, however, be a part of a future security approach in wireless air traffic communication and will be discussed in Chapter 6.

4.5 Summary of Vulnerabilities and Attacks

Table 4.1 lists the currently available literature on the analysed attacks on the aviation communication technologies as discussed in this chapter. It also provides

¹²Black keys are safe to transmit since they are encrypted with an encryption key. Red keys on the other hand are unencrypted and classified as highly sensitive.

Technology	Attacks	Requirements	
		Confidentiality	Integrity
VHF	[74]	-	X
PSR	[34, 77, 96, 97]	-	X
SSR	[80, 98]	-	X
ADS-B / TIS-B / FIS-B	[9, 11–13, 20, 99]	-	X
CPDLC	[32]	X	X
MLAT	[83]	-	X
ACARS	[11, 20, 86, 100]	X	X
TCAS	[13, 101]	-	X

Table 4.1: Overview of attacks and security requirements for wireless aviation communication systems.

an overview of their security requirements in terms of confidentiality and integrity. We omit *availability* here as it is a key requirement for all systems in aviation, especially during critical flight phases. Attacks on availability of wireless technologies in particular through jamming are notoriously difficult to handle; we assess the potential flight safety impact on the next chapter.

Likewise, we assume all technologies to require *integrity* for a safe aviation environment. However, this requirement may differ in severity, in line with each technology’s impact on safety. Concerning *confidentiality*, the desire in the aviation community is to stay with open ATC systems that are available to everyone (see Section 4.7), which is generally unproblematic for commercial airlines.¹³ On the other hand, the two data links (ACARS and CPDLC) fulfil a broad range of purposes, some of which may be critical to privacy and safety. Thus, for these data links, having full confidentiality is strongly desirable, regardless of the user.

4.6 Mitigating Factors

After examining the vulnerabilities present in all wireless protocols, we want to consider some of the factors that can help mitigate the threat of attacks on the

¹³Yet, a number of privacy problems arise for private aircraft and general aviation, which are discussed in Chapter 8.

applications and services provided by these technologies. These factors can be divided into a) *rules and procedures* and b) the existence of *redundant communication systems*. Both factors do not traditionally consider security, i.e., malicious attacks on the analysed systems, but focus on maximizing the overall safety of the aircraft, regardless of the nature of the interfering elements.

Procedures

Procedures and practices for systems failures are plentiful in aviation and try to cover all imaginable high-risk cases. While they are purely aimed at non-deliberate failures, many of them (e.g., lost communications procedures, the FAA's Emergency Security Control of Air Traffic (ESCAT) plans for major crises such as 9/11, or so-called ATC Zero procedures to handle local failures of ATC centres) will have mitigating consequences for deliberate attacks as well.

Overall, human factors are of crucial importance in aviation. Many pilots have experienced incorrect instrument readings and are trained to check and double-check at all times. On the ATC side, if a controller noticed more than one aircraft using the same transponder identity, they would call the plane and provide a different one, followed by a request to the pilot to set an IDENT flag on the aircraft transponder, which will be displayed on the ATC screens.

Unfortunately, even without malicious attackers, procedural defences are not always successful. One particular example is when priorities and instructions of different protocols — both human and technical — are not well-defined or even conflicting, such as in the fatal 2002 Überlingen mid-air collision.¹⁴

On top of this, it is highly unlikely in reality that *every* single malicious communications interference is detected by humans and defended using only rules and procedures. Sophisticated attacks on communication are typically not considered in current aviation training. Existing procedures (for example, FAA regulations and the Aeronautical Information Manual [102] for general aviation in the US) cover

¹⁴The two aircraft had received different instructions from ATC and their on-board equipment. By each following a different procedure, their collision course was not resolved. See e.g., <http://goo.gl/WZFceZ> for a full analysis.

only faulty systems and always assume that available information is genuine and has not been maliciously altered. Furthermore, any new generation of pilots is trained to rely on instruments and digital systems even more,¹⁵ strongly motivating the increased need to secure the integrity of these systems in the future.

Redundancy

Related to procedure-based mitigations is the concept of redundancy. The most common type is the use of *hardware redundancy* to improve the availability of a system. This can, for example, include numerous duplicate senders or receivers within the aircraft, or completely independent systems. Similarly, any ATC ground station utilizes several receivers for their surveillance or multilateration systems, making the failure of one, or even a few, unproblematic.

Besides this, there are procedures for wireless systems failures that rely on different systems for redundancy. When a primary data source is not available, there are often other systems at the disposal of a pilot or a controller to obtain the required information. For example, an aircraft with a Mode A/C/S transponder failure may still be tracked using PSR. However, the altitude and ID, among other information, are lost, causing a significant drag on the controller's awareness and attention, while he/she still has to maintain the required separation for all assigned aircraft.

However, this type of *technological redundancy* has inherent security flaws. As the technologies were not developed for redundancy, but simply happened to become alternative options because of legacy reasons, relying on older technologies comes with a degradation of content quality. For example, separation minima may need to be increased when SSR systems fail, hence an attack can achieve a significant degradation of service quality. On top of this, the availability of systems can differ across airspaces, making general assurances difficult, especially for smaller airports.

Overall, technological redundancy is effective in many cases when it comes to preserving safety. However, it fails when there is no suspicion of malicious activity by the user(s), or when attacks are conducted on the procedures themselves

¹⁵This reliance is illustrated by some recent incidents, most notably the Air France disaster where the pilots flying trusted instruments even though they knew they were unreliable [103].

(i.e., multiple technologies are targeted simultaneously, based on knowledge of the processes followed by ATC).

4.7 Reasons for the Current State of Aviation Security

We conclude this chapter on wireless vulnerabilities with a look at the reasons for the current state of affairs. There are five identifiable historical causes that led to the lack of built-in wireless communications security within the air traffic system, leading to the security flaws discussed above. These causes are: long development & certification cycles, legacy & compatibility requirements, cost pressures, frequency overuse, and a preference for open systems within the aviation community. The ingrained inertia of these reasons also explains the difficulties in fixing the existing problem within the short term. Hence, we believe it is imperative to fully understand them before beginning to improve security in aviation, as we aim to do in the remainder of this dissertation.

Long Development & Certification Cycles The development and certification cycles for new technologies in aviation are typically up to two decades. Taking ADS-B as an example, the development of its current form started in the late 1990s [104]. The widespread rollout and mandatory use will however only be completed by 2020 in the most advanced airspaces. This slow and cautious approach reflects the safety-focused thinking within the aviation community, where a multitude of tests and certifications are required before giving a technology the green light [105]. Unfortunately, while this approach is extremely effective in reducing technical failures, it does not take into account the increased adversarial potential and shifting threat model created by the recent advancements in wireless technologies discussed above.

Legacy & Compatibility Requirements As a truly global and interconnected system, civil aviation requires technical protocols and procedures that are understood as widely as possible [106]. However, new protocols and technical advancements are not introduced in all airspaces at the same time, but depend on local authorities and available infrastructure. It follows that older technologies are kept in service not only for backup and capital investment reasons but also to offer the largest possible compatibility for air traffic control all over the world, as high levels of interoperability between countries require time and effort [107].

Cost Pressures Tying into the previous point, the aviation industry is famously competitive and under major cost pressures [108]. Changes to existing aircraft equipages are expensive, and thus unpopular, unless they provide immediate financial or operational benefits to the aircraft operators who foot the bill for the installation of new technologies. Apart from these two main drivers, fundamental technological equipment changes happen primarily through regulatory directives, which are often subject to long lead times and struggle with extensive industry lobbying. As a compromise, legacy technologies are sometimes overhauled to save costs [106]. An example for this is the ADS-B protocol which, for commercial aircraft, relies on the old Mode S technology instead of using a new data link developed from the bottom up.

Frequency Overuse As shown in [79] and [24], some of the ATC frequencies such as the 1090 MHz channel are severely congested. An ever-increasing number of aircraft share the same frequencies, and this is further exacerbated by UAVs set to enter the controlled airspace in the foreseeable future. As a consequence, existing ATC protocols suffer from severe message loss, a fact which at the same time also poses a problem for potential cryptography-based security solutions to overcome.

Preference for Open Systems There is a case for air traffic communication protocols to be open to every user, i.e., while authentication would be highly desirable, confidentiality through encryption of the content is not. Despite the

associated security and privacy problems, the ICAO plans for future protocols to be openly accessible. This approach is supposed to fulfil typical aviation requirements such as ease of communication, compatibility, and dealing with administrative differences across countries and airspaces [109]. While we acknowledge that open systems are a requirement for the effectiveness of air traffic control for the foreseeable future, it is crucial to start considering and mitigating the downsides, which are rapidly increasing due to fast technological changes.

Nicht die Wahrheit, in deren Besitz irgend ein Mensch ist, oder zu sein vermeint, sondern die aufrichtige Mühe, die er angewandt hat, hinter die Wahrheit zu kommen, macht den Wert des Menschen.

The true value of a man is not determined by his possession, supposed or real, of Truth, but rather by his sincere exertion to get to the Truth.

— Gotthold Ephraim Lessing's *Über die Wahrheit*

5

Raising Awareness of Cybersecurity in Aviation

Contents

5.1	Survey Design	60
5.2	Survey Limitations	61
5.3	Survey Results	62
5.4	Assessment of Concrete Scenarios	67
5.5	Summary & Outlook	70

After examining the current state of aviation cybersecurity in the previous chapter, we require additional domain knowledge, data on the awareness of cybersecurity in aviation, and an assessment of the potential impact of attacks. To collect this information, we conducted a survey across all aviation circles, which is the first to address these issues publicly, and we are thankful to all involved aviation authorities and air navigation service providers for their help.

The three main research questions that this survey looks to answer are: a) Which technologies are considered to have the biggest impact on safety? b) Are aviation stakeholders aware of security issues in the wireless technologies they utilize? c) If yes, are these issues considered a concern towards safety?

We analyse the answers to these questions after discussing the design of the

study, and the demographics of the respondents. We then analyse the respondents' assessment of nine concrete hypothetical attack scenarios in Section 5.4.

It has to be noted that increased awareness of insecure protocols does not mean that controllers' or pilots' behaviour in their daily work should change. We believe this is neither feasible nor a sensible course of action. However, increased awareness among *all* aviation stakeholders should provide the necessary basis for a general change in the aviation community's approach to cybersecurity issues. Without all parties on board, crucial regulatory changes are unlikely to be implemented within reasonable time frames.

5.1 Survey Design

We planned and conducted our survey with the help of private pilots and a full-time professional air traffic controller. They advised us on the appropriate question language with relation to aviation subject terms and helped us devise relevant attack scenarios for pilots and controllers, respectively. Furthermore, they provided us with the necessary aviation expertise and background at every stage during the design, implementation, and execution of this survey.

Our survey was conducted fully anonymously over the internet to protect respondents from potential repercussions when speaking about the security of ATC systems or disclosing safety problems. We used a questionnaire on SurveyMonkey, where we did not collect the respondents' IP addresses, so we could not make any inferences to their work place. We obtained ethical approval for this survey from the University of Oxford's Social Sciences & Humanities Inter-Divisional Research Ethics Committee (IDREC) under the Ref No: SSD/CUREC1A/15-033.

The recruiting was done first through a controlled dissemination (CD) via mailing lists of air navigation service providers, airlines and other aviation-related organisations. In a second phase, we attempted to recruit participants through a separate open dissemination (OD) of the questionnaire in eight closely-moderated aviation forums. Out of these, only two forums for private pilots were willing

to cooperate; the remaining 6 were wary of negative publicity or did not give a reason for declining our request.

Overall, we did not aim to survey people with deep knowledge in computer security, but to get a more realistic opinion of the aviation community as a whole.

5.2 Survey Limitations

Our approach has some natural limitation, which we outline here. We tried to mitigate any confounding factors through our design, but we acknowledge some potential limitations caused by the characteristics found in aviation technology:

- **Proprietary systems:** Typically, systems are implemented by different companies following loose standards. Even some of the widely used protocols (e.g., ACARS) have proprietary elements that are not freely available. To counteract this problem, we abstracted away from the concrete implementations. We designed the questions such that we could draw more general conclusions on the respondents' knowledge of the systems' security.
- **Fragmentation:** Likewise, there is a forest of different systems, regulations, and processes in aviation. Depending on the airspace, the availability, knowledge, and usage of the discussed protocols differ. However, we mitigated this problem by surveying experts from many countries, making sure their judgement of security in aviation technologies did not vary significantly.
- **Representativeness:** Considering distribution and potential self-selection, we do not claim that our results are necessarily representative of the aviation community. Yet, when reconciling with comments and conversations with experts, we believe in their validity.

Overall, we believe it is an important task to abstract away from single technologies and gather a more systemic picture of the awareness on wireless security in aviation as a whole. It is worth noting that a survey-based analysis is an accepted tool in aviation research. For example, the authors in [110] recently used it to analyse the safety of the FIS-B protocol.

Group	# Respondents	Share
Private Pilot	77	32.0%
Commercial Pilot	59	24.5%
Civil ATC	37	15.4%
Aviation Engineer	10	4.1%
Aviation Authority	7	2.9%
Military Pilot	5	2.1%
Military ATC	4	1.7%
Other	42	17.4%

Table 5.1: Occupations of survey respondents ($n = 242$).

5.3 Survey Results

Demographics

We had 242 completed surveys, 110 or 45.5% from the controlled dissemination and 132 or 54.5% from the open dissemination. We compared the results of both dissemination methods and found no significant differences in the respondents' evaluations apart from their professions: 55.7% of OD respondents were part of the general aviation community (GA, i.e., private pilots) and not otherwise working in aviation, compared to only 3.6% of CD respondents. We analyse the responses as a whole unless stated otherwise.

The participants' aviation experience was fairly evenly distributed, with 32% having 20 years or more, and about 22% offering an expertise of less than 5 years, 5-10 years, and 10-20 years, respectively. The top working countries were the UK (37.7%) and the US (23.3%). A further 37.3% work in Continental Europe, with 4 respondents from other countries around the world (Indonesia, Hong Kong, Canada, UAE).

As illustrated in Table 5.1, most of our respondents were private (32%) or commercial pilots (24.5%) followed by civil air traffic controllers (15.4%) and aviation engineers (4.1%). We had responses from professions as varied as Air Traffic Safety Electronics Personnel, researchers at Air Navigation Service Providers, ATC technicians, aviation software developers, or Flight Information Service Officers/Instructors amongst others. The average response time length was 31 minutes and

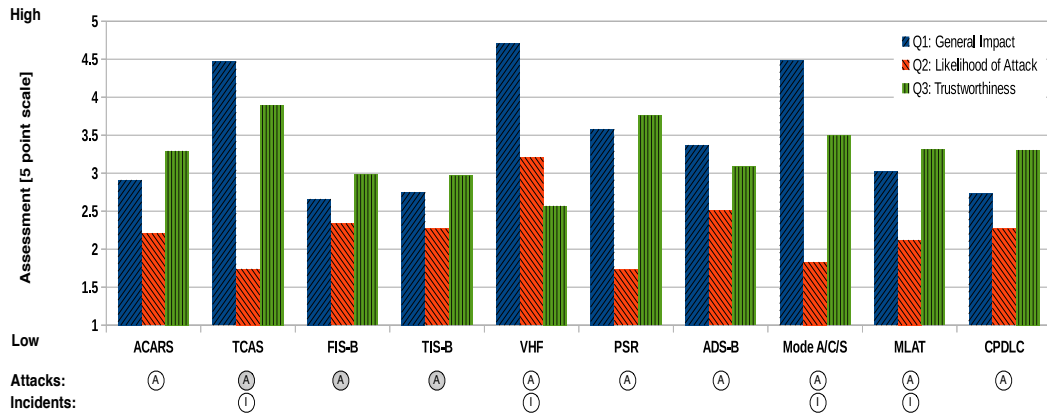


Figure 5.1: Respondents’ risk assessment of 1) the flight safety impact, 2) the likelihood of being attack targets and 3) the trustworthiness against manipulation of each protocol. The respondents ($n = 235$) answered three questions. Q1: “How would you rate the flight safety impact of each of these technologies?”, Q2: “How do you rate the likelihood that a malicious party injects false information into these technologies?”, and Q3: “How would you rate the trustworthiness of information derived from these technologies against intentional manipulation by a malicious party?”. Answers were provided on an equidistant 5-point-Likert scale. The circles below the bars show the availability of published attacks discussed in Section 4.5 and (publicly or privately) reported incidents for each protocol. Protocols with grey circles are vulnerable by extension as they depend on data from other vulnerable protocols.

10 seconds, which, along with the plethora of comments we received (more than 400), shows that the respondents took the survey seriously.

Self-assessed Knowledge and Work Environment

The respondents judged their air traffic communication knowledge above average for their field (3.76 out of a symmetric, equidistant 5-point Likert scale, where 1 is “very bad” and 5 is “very good”). The technologies that the respondents considered themselves most familiar with are VHF and Mode A/C/S. This is explained by the prevalence and importance of these technologies in current aviation processes: Mode A/C/S and VHF are also the most relied upon technologies, followed by TCAS. Very interestingly, more than 50% of the respondents answered that to some extent they already rely on ADS-B in their work, despite the protocol not being operational in most airspaces.

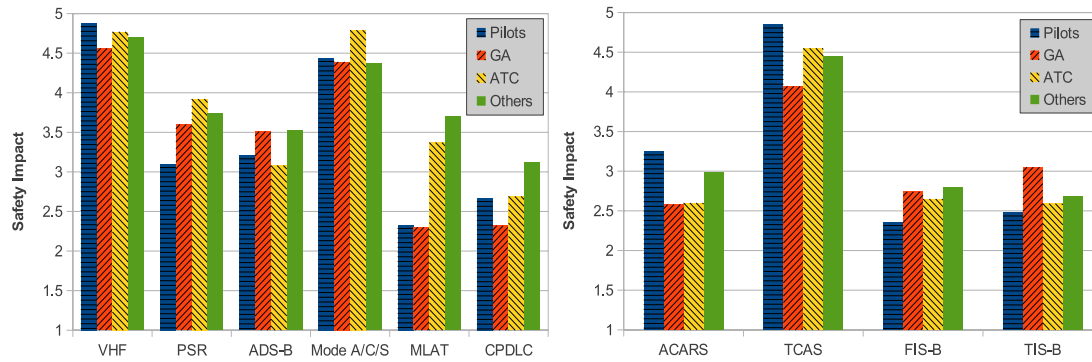


Figure 5.2: Assessment of the safety impact of aviation technologies by stakeholder group. Question: “*How would you rate the flight safety impact of each of these technologies?*” Data gathered from 43 commercial pilots, 55 private pilots, 32 controllers, 45 others.

Impact Assessment

Fig. 5.1 shows that VHF, Mode A/C/S, and TCAS are also considered the technologies with the highest safety impact, all around 4.5 out of 5 (“very high”). But even the least important technologies (CPDLC, FIS-B, TIS-B, which are not widely operational yet) still have a moderate impact according to the respondents.

Fig. 5.2 breaks the protocols down further along the lines of the different stakeholders (controllers, commercial pilots, private pilots, others). We can see that the different groups’ impact assessments do not differ significantly, with the exception of MLAT, which is rated considerably higher by controllers compared to pilots (commercial or private). The latter group does not actively use MLAT while the work of controllers partly relies upon it, depending on the configuration of their airspace.

Generally, the isolated impact of a single protocol can often be limited. A German controller represents the feelings of many (but not all): “*in case of loss there are still backup systems and cross-check possibilities so there are no really high impact ratings*”, illustrating the traditional safety approach of redundancy extended to security.

Security Assessment

In terms of security awareness, the responses showed a very unconcerned aviation community with the exception of VHF:

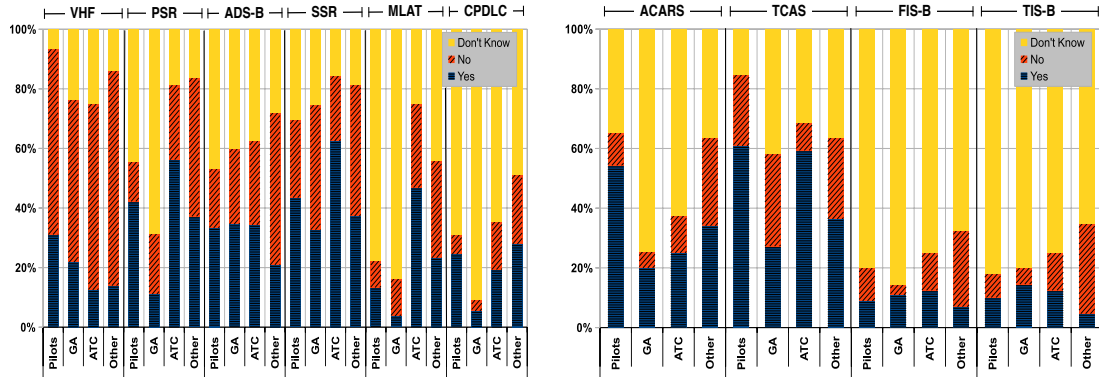


Figure 5.3: Assessment of the security capabilities of aviation technologies by stakeholder group. Question: “Which of these technologies do you believe ensure the integrity of the transmitted data and the authenticity of its origin?” Data gathered from 43 commercial pilots, 55 private pilots, 32 controllers, 45 others.

- As shown by the red (middle) bars or Q2 in Fig. 5.1, VHF is considered the least reliable by the participants when it comes to the likelihood of manipulation by an adversary; >60% of respondents say it is “very likely” or “likely” to be attacked.
- Consistently, VHF is also the least trustworthy technology (see Fig. 5.1, Q3). All other technologies were rated at least 3 points out of 5, with TCAS the most trustworthy at almost 4 out of 5 by the respondents.
- Many experts report actual experience with non-legitimate uses of the frequency, one noting “VHF is an increasingly common comms signal to be maliciously emulated by non-involved parties. Particularly on tower frequencies. Anyone can buy an aviation transceiver without licence.” Others mentioned frequent disturbances by pirate radio stations.
- Of the digital protocols, ADS-B was considered the most likely to be attacked. We speculate this might be due to the raised awareness caused by widely publicized attacks on ADS-B specifically. Yet, at 2.5 points, the responses were not highly concerned with a potential real-world incident.
- All technologies except VHF are considered relatively unlikely to be attacked, despite the prevalence of known vulnerabilities and reported incidents indicated

in Fig. 5.1 and discussed in detail in the last chapter.

Fig. 5.3 again takes a look at the question of security at the level of controllers, commercial pilots, private pilots, and others:

- Less than 10% of all respondents are aware that CPDLC is not authenticated, especially considering the plans of many aviation authorities to shift more responsibilities towards CPDLC in the future. As shown in Fig. 5.3, across all groups, a majority of those respondents who said they knew the answer, were mistaken about its security capabilities.
- A large majority of controllers (almost 60%) believe PSR offers integrity checks.
- Worryingly, more than 40% of all respondents wrongly believe the Mode A/C/S protocol offers built-in security features, including a notable 60% among controllers.
- Only about 20% of the participating pilots and controllers are aware of ADS-B's security shortcomings, despite the existing research publicised over the previous years.
- Pilots are particularly unaware of MLAT's security features (around 80% answered with "Don't know"), again reflecting the fact that it does not consciously feature in their work compared to controllers.
- Knowledge about FIS-B and TIS-B is still very limited due to their rollout only in the US airspace.

These results clearly illustrate the lack of awareness within the aviation community when it comes to the security of all utilized technologies.

5.4 Assessment of Concrete Scenarios

To further improve the understanding of insecure wireless aviation technologies, we transformed some of the previously discussed wireless attacks into concrete scenarios and evaluated the respondents' assessments of the practical impact of such realistic threats. In nine hypothetical scenarios the respondents had to rate the impact of air traffic communication technologies misreporting the current air traffic picture. Note that this could happen for different reasons, many of which are not caused by an attack. We were interested in the impact regardless of the underlying cause which would typically be abstracted away for end users of these systems. The respondents assessed these scenarios with five options: "No effect", "Minor loss of situational awareness",¹ "Major loss of situational awareness", "Full denial of aviation service" and "Don't know". Table 5.2 shows the questions and the distribution of the results.

Some participants would have liked more information as answers can depend on many specifics. Unfortunately, it is impossible to cover the large number of potential combinations of systems and situations separately. Thus, we consider the quantitative part of the questions about these scenarios as an abstract expert opinion. Consequently, we complemented all questions with a qualitative approach by inviting comments, a few of which we provide to gain a more complete picture. We focus on three systems: ATC radar, showing the air traffic picture to ground controllers (Questions 1,3,5,7), TCAS screens displaying intruders in an aircraft's immediate airspace to the pilot (Q2,4,6,8), and instrument readings in a cockpit based on navigation aids such as GPS (Q9). In terms of attack classes, we analyse message injection (Q1-2), content manipulation (Q3-4,9), selective jamming (Q5-6), and full DoS caused by jamming (Q7-8). We report the answers of 92 pilots to the scenarios relating to their expertise (Q2,4,6,8,9), and of 31 controllers to the scenarios relevant to ATC (Q1,3,5,7).

¹Situational awareness is a concept widely used in aviation despite some criticism. A historically accepted definition is "the perception of elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" [111].

What is the effect if...	Surveyed group	No effect	Minor loss of situational awareness	Major loss of situational awareness	Full denial of aviation service	Don't know
1) ...a non-existing target shows up on an air traffic control radar screen?	ATC	16.13%	54.84%	25.81%	0.00%	3.23%
2) ...a non-existing target shows up on a TCAS screen?	Pilots	10.75%	44.09%	31.18%	0.00%	13.98%
3) ...wrong label indications show up on an air traffic control radar screen (e.g., altitude, selected altitude, callsign)?	ATC	3.01%	28.31%	54.82%	6.02%	7.83%
4) ...wrong label indications show up on a TCAS screen (e.g. relative altitude)?	Pilots	3.23%	30.11%	51.61%	4.30%	10.75%
5) ...information or whole targets are selectively missing from an air traffic control radar screen?	ATC	6.67%	6.67%	83.33%	3.33%	0.00%
6) ...information or whole targets are selectively missing from a TCAS screen?	Pilots	6.52%	23.91%	48.91%	7.61%	13.04%
7) ...all data is missing from an air traffic control radar screen system (for example for 10 minutes)?	ATC	3.23%	6.45%	45.16%	45.16%	0.00%
8) ...all data is missing from a TCAS screen system (for example for 10 minutes)?	Pilots	6.52%	27.17%	45.65%	6.52%	14.13%
9) ...the position of an aircraft as shown to its pilot is incorrect?	Pilots	5.38%	22.58%	60.22%	7.53%	4.30%

Table 5.2: Answer distribution for nine hypothetical safety scenarios as answered by 92 pilots and 31 controllers, respectively.

Comparing the scenarios, we find that the impact of a single ghost aircraft on a radar screen (**Q1**) is rated as mostly “minor”. Controllers say it can cause delays during busy times due to additional work and increased separation requirements. It is generally common to experience non-existing radar targets without malicious intent, caused by transponder issues, reflections, clutter, and other reasons. The impact of a ghost aircraft appearing as an intruder on TCAS screens (**Q2**) is rated slightly higher; it is seen as likely that it would trigger unnecessary manoeuvres that could lead to further unpredictable complications. For both ground radar screens and pilots’ TCAS displays, wrong label indications are rated higher in impact compared to ghost aircraft (**Q3+4**). For both scenarios, more than 50% consistently rate it as at least a major loss of situational awareness.

Selectively missing information is also rated highly in terms of loss of situational awareness (**Q5+6**). Some comments note that it is inherently difficult or impossible for a pilot or controller to cross-check missing aircraft as they would not even know about it in many situations, and thus no procedure is triggered. A faulty TCAS would even lose all purpose and relying on it could have severe impact in bad weather and visibility conditions. Under non-selective jamming attacks on ATC (**Q7**), operations would be stopped for general aviation and commercial starts and landings would be reduced, causing a partial denial of service. All operations would go procedural and use the less efficient VHF for communication and separation (this is not hypothetical, as two major incidents in Central Europe in 2014 proved [5]). For TCAS (**Q8**), the jamming scenario is rated very similarly to selectively missing information.

For the last scenario (**Q9**), we asked about the general effect in case the instruments of an aircraft show its position to be different from the true one (e.g., in case of a GPS malfunction, whether caused deliberately or not). If an incorrect position of an aircraft is shown to its pilot, major loss of situational awareness would occur but the specific outcomes are again impossible to predict. The respondents suggest a wide range of scenarios, from additional work for the controller up to a fighter jet escort.

One very interesting comment noted that German regulations for controllers state that they should not consider outages or wrong labels as a possibility in their work but always rely on the systems (the implication being that otherwise the workload could not be handled). While experienced controllers certainly know about malfunctions such as ghost radar tracks in their work, this illustrates the importance of – and reliance on – wireless technology.

5.5 Summary & Outlook

As our investigation shows, aviation professionals at large are unaware of the state of security in aviation technology. A majority of participants falsely believe, for example, that even the most common surveillance technologies offer authentication, clearly contrasting with the security community’s knowledge. Indeed, our findings are in line with research from cognitive science, which suggests that so-called “expert blind spots” exist: i.e., having a large amount of domain-specific knowledge may prove disadvantageous on tasks such as forming remote associations among disparate concepts [112]. Hence, we conclude that raising awareness of these issues must form the foundation of all future work in aviation security.

While our results do not show very large differences in current awareness between the different stakeholders, it is clear that the desired form will take different shapes for each constituency. In other words, not all of the participants in the survey need be concerned with all aspects of wireless security. For example pilots, especially private ones, do not require in-depth knowledge of MLAT spoofing, whereas aviation engineers working for ASNPs should certainly consider this problem in their system.

Thus, pilots, controllers, regulators, and system engineers need to be educated in those areas most relevant to their work. For pilots, this would include in particular a discussion of GPS spoofing and jamming in their training, including dispelling the myth that their systems are immune to these attacks. It should also cover the vulnerability of any landing system relying on unauthenticated wireless links as well as an assessment of the trustworthiness of instructions gained over

CPDLC, which, as has been pointed out in the previous chapter, are as vulnerable as the current VHF system.

On the part of the controllers, awareness of the potential of wireless attacks needs to increase. Where such attacks are detected quickly, mitigating measures can be taken such as a fallback on procedural operations.

Systems engineers and authorities certainly bear the largest responsibility to ensure the security and thus safety of the future air traffic system, the former on the smaller operational level, the latter on the larger strategic scale. Engineers need to consider the possibility of malicious interference with the received unauthenticated wireless data in their system design. As they do not have the power to change the protocols, this could be achieved for example through data fusion or intrusion detection approaches as discussed in Chapter 7.

On the part of the authorities, it is clear that top-down regulations are crucial in aviation. They must immediately include passive and active threats in their considerations for all future ATC technologies. While it is difficult to significantly speed up the process of introducing protocols that include security by design, it is imperative that the focus on wireless security is strongly increased in research and practice.

La helicoide representa el movimiento, el paraboloide es el padre de la geometría, el hiperboloide es la luz y el tetraedro la síntesis del espacio.

The helicoid represents movement, the paraboloid is the father of geometry, the hyperboloid is light and the tetrahedron is the synthesis of space.

— Antonio Gaudí [113]

6

Developing a Taxonomy of Air Traffic Communication Security Approaches

Contents

6.1	Secure Broadcast Authentication	74
6.1.1	Non-Cryptographic Schemes on the Physical Layer	75
6.1.2	Public Key Cryptography	77
6.1.3	Retroactive Key Publication	80
6.2	Secure Location Verification	82
6.2.1	Multilateration	82
6.2.2	Distance Bounding	84
6.2.3	Kalman Filtering and Intent Verification	86
6.2.4	Group Verification	87
6.2.5	Plausibility Checks	88
6.3	Summary	90

Substantial work has been done on some parts of air traffic communication security over the last few years and several approaches have been proposed to enhance the security of SSR and ADS-B in particular. Furthermore, a large amount of research has been done in related fields such as vehicular ad hoc networks (VANETs) and wireless sensor networks where broadcast authentication and security also play an important role. While some of the discussed ideas may not be directly applicable to the unique aviation environment, they can inspire future solutions.

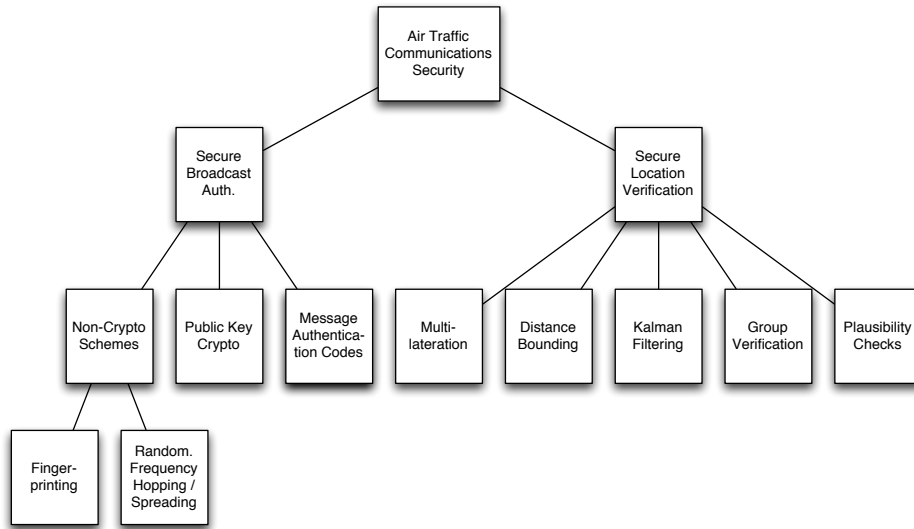


Figure 6.1: Taxonomy of air traffic communication security approaches.

As shown in the taxonomy in Fig. 6.1, we identify two distinct approaches to securing air traffic communication from the existing literature: **Secure Broadcast Authentication** and **Secure Location Verification**. Consequently, Section 6.1 examines the various schemes that apply asymmetric properties (cryptographic and non-cryptographic) to **authenticate the broadcast communication** itself while Section 6.2 reviews several different methods that seek to **verify the authenticity of the location claims** made by aircraft and other ATC participants.

6.1 Secure Broadcast Authentication

Secure Broadcast Authentication is one possible means to prevent and/or detect attacks in a wireless networks. This section will describe the various methods that have been proposed in the literature, typically for wireless sensor networks or VANETs and analyse their applicability to ATC.

Authentication of messages on a broadcast medium is harder compared to point-to-point communication. A symmetric property is only useful in point-to-point authentication where both parties trust each other. Thus, an asymmetric mechanism is inherently required so that receivers can *verify* messages but are not able to

generate authentic messages themselves [114]. For a good overview over secure broadcast communication in general, the reader is referred to [115].

The goal is to keep the open nature of ATC intact while offering a potential authentication mechanism. This could be done either globally or only selectively in cases where suspicious behaviour has been detected. Such reactive authentication could lessen the strain on the network by only requiring additional security (and thus computational and communicational overhead) at times when incidents seem more likely.

6.1.1 Non-Cryptographic Schemes on the Physical Layer

Non-cryptographic schemes such as fingerprinting comprise various methods for wireless user authentication and device identification techniques, either based on hardware or software imperfections or characteristics of the wireless channel which are hard to replicate. The goal of such schemes is to identify suspicious activity in a network. Finding a signature for legitimate beacons in a network, possibly being able to tell apart ground stations from aircraft, identifying the type of aircraft, or even individual machines, provides data useful for the development of an intrusion detection system [116]. If there are tangible differences between legitimate and non-legitimate packets on the physical layer, then machine learning techniques could be employed to develop a model for predictions of normal behaviour and also statistical thresholds beyond which an activity is considered suspicious. Even if it is only feasible to identify classes of devices instead of singular participants, this could prove to be valuable information in detecting intruders. Yet, fingerprinting does not provide sure-fire security, and various attacks and concerns have been brought forward [117].

We consider two types of fingerprinting: software-based and hardware-based.

Software-Based Fingerprinting Software-based fingerprinting techniques try to exploit distinctly different patterns or behaviour of software operating on wireless equipment. Depending on the specification of a protocol, there is a lot of leeway for

manufacturers and developers when implementing software on a given device. If there is enough entropy in information about the combination of chip sets, firmware, drivers to tell apart different wireless users, this approach can be used to verify their continuity up to a certain degree. As a downside, it seems likely that large fleets of airline operators are fitted with very similar or same hardware, making them harder or even impossible to differentiate and on the other hand easier to study and copy for a potential attacker.

Hardware-Based Fingerprinting A number of techniques have been proposed to identify devices based on unique hardware differences. Some of these differences can be used for *radiometric fingerprinting*, exploiting differences in the turn-on/off transient (see, e.g., [118]) or the modulation of a radio signal to build unique signatures.

Another unique hardware feature amongst wireless devices is *clock skew*. As no two clocks run precisely the same, this can be used to create signatures and enable identification.

Recently reviewed options for future systems include the use of so-called physically unclonable functions (PUFs), which essentially exploit specifically implemented circuits to create unique and secure signatures, thus abandoning the scope of non-cryptographic solutions.¹ Furthermore, besides requiring new hardware, this approach also necessitates an overhauled messaging protocol, including a challenge and response model [120], making it a difficult fit for current aviation technologies.

Randomized/Uncoordinated Frequency Hopping / Spreading A physical-layer scheme different from fingerprinting, Frequency Hopping Spread Spectrum (FHSS) and Direct Sequence Spread Spectrum (DSSS) are both used in wireless systems to improve protection against malicious narrow band and pulse jamming as well as eavesdropping. In their usual form they both require a pre-shared spreading code or hopping pattern between sender and receiver which makes it hard to follow or hinder the communication for anyone without access to the

¹For a good overview on PUFs, see [119].

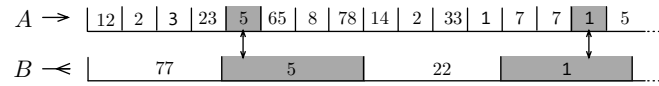


Figure 6.2: Illustration of Uncoordinated Frequency Hopping after [121].

code/pattern. This is also exploited in military communications but is not a viable option for world-wide civil and commercial ATC where such secret codes would presumably not stay secret for long.

The need for a pre-established code can be relinquished by employing random, uncoordinated versions of FHSS and DSSS. Strasser et al. [121] propose such a physical layer approach to counteract jamming in wireless broadcast scenarios. Uncoordinated Frequency Hopping (UFH) provides a viable way to broadcast initial messages without an attacker being able to jam the transmission in an efficient way. The key insight to these approaches is that, contrary to normal frequency hopping mechanisms, sender and receiver(s) only rely on the statistical chance to be on the same channel at the same time (see Fig. 6.2 for an illustration).

Uncoordinated Direct-Sequence Spread Spectrum (UDSSS) [122] and Randomized Differential DSSS [123] are techniques based on the same principle. They rely on the statistical chance that spread codes randomly chosen by sender and receiver(s) will happen to be the same every so often.

While the proposed methods can effectively defeat jamming and modification attacks, the inherently lower performance and a prolonged transmission time make them difficult to use in a system of the scale found in aviation. Nevertheless, frequency hopping or spread spectrum techniques could play a crucial role in future developments.

6.1.2 Public Key Cryptography

Cryptographic measures have been a tried and tested means to secure communication in wireless networks and must subsequently also be considered in the ATC setting. One question to examine is if the currently existing technologies can be encrypted or whether new developments are required.

Finke et al. [124] examine various encryption schemes for ADS-B, including the possibility to do the key management for symmetric encryption out of band, for example through CPDLC, as the 56 bit payload of ADS-B messages are insufficient. The authors also give an analysis of the security and practicability of asymmetric, symmetric and format preserving encryption. In their conclusion, they support a symmetric cipher using the FFX algorithm (format-preserving, Feistel-based encryption with multiple implementation variances) which can encrypt non-standard block sizes (e.g., ADS-B's 112 bit messages) with sufficient entropy. However, the difficulties concerning key management and distribution are strongly acknowledged. Most recently, Wesson et al. [125] look at the broader question of how to use encryption to secure ADS-B and conclude that the problems with symmetric cryptography are too large to overcome. They argue that a public key infrastructure (PKI) is the only feasible cryptographic approach and propose elliptic curve digital signature algorithm (ECDSA) signatures as the smallest and thus best solution. Besides the key management problem, they further analyse the interference burden on the wireless channel, showing that even without the significant additional traffic that is currently found on the 1090 MHz frequency, the decrease in operational capacity would potentially be crippling.

As mentioned before, if broadcast authentication is needed, one requires an asymmetric property, a characteristic fulfilled by public key cryptography. Samuelson and Valovage [126] report on an implementation of authentication and encryption in UAT using a public key infrastructure. Their method uses a hash to create a message authentication code (MAC) that can be used to authenticate the message and can be extended to full encryption but no further details are publicly available.

The general idea is to use a challenge/response format with an authenticator ground station, who authenticates every participant in its reach and notifies a higher authority and/or all other participants of any failed authentication. This concept requires the station to have access to a worldwide database of secure keys that is both hard to maintain globally as well as subject to possible security breaches.

Costin et al. [12] suggest a “lightweight” PKI solution which essentially amounts to a retroactive part publication of the key as discussed in Section 6.1.3: Aircraft A transmits the signature distributed over a number N of messages, so that after every N messages the surrounding participants have received A ’s signature. The recipients keep the messages until the full signature has been transmitted, at which point they can authenticate the buffered messages. The authors suggest that the PKI key distribution necessary for this scheme could be done during an aircraft’s regular check-ups.

Robinson et al. [127] analyse various different solutions to create PKI infrastructures for a general “airplane assets distribution system”. Although, the work is not discussing wireless protocols but instead focuses on the distribution of software and data on the ground, the authors identify the airline industry’s needs and requirements for a PKI infrastructure, and it seems plausible that the same system could be used to secure air traffic control data. According to their analysis, an ad hoc approach without a central authority, employing pre-loaded trust certificates, could be used as a short-term solution until a more structured, long-term public key infrastructure has been developed.

The obvious idea for a centralised key distribution would be to have aviation authorities such as the FAA act as a certificate authority (CA). But assuming the role of a CA is no easy task. Even many specialized institutions have had to report numerous security breaches over the last decades. Furthermore, if this problem is sufficiently solved, there remains the question of how aircraft from airspaces mandated by different authorities can securely communicate with each other. These challenges are somewhat analogous to the same approach in vehicular networks as discussed in [128] but arguable even worse due to the large internationalization found within aviation communication.

There are certain natural disadvantages to using an encryption solution that cannot be overcome (or only with great difficulty) as mentioned in [129]:

- Despite the encryption of data frames, management and control frames are not protected.

- It immediately and unmitigatably breaks compatibility.
- Key exchange is notoriously difficult in ad hoc networks, which are by definition without a centralized institution. They are often too dynamic, requiring constant adaptation. This would result in too much overhead in both the number and the size of messages.
- The open nature of ATC is widely seen as a feature.
- One-time signature, even using techniques such as Merkle-Winternitz, provide an overhead of 80 bytes and more, simply to sign 60 bits [114].

To conclude, while highly effective against most attacks on air traffic communication, it is difficult to build any kind of encryption scheme with currently existing aviation technologies, making new developments necessary (and desirable).

6.1.3 Retroactive Key Publication

A variation on traditional asymmetric cryptography is to have senders retroactively publish their keys which are in turn used by receivers to authenticate the broadcast messages. The key concept is simple: Any broadcasting entity produces an encrypted MAC which is sent along with every message. After a set amount of time or messages, the key to decrypt this MAC is published. All listening receivers, who have buffered the previous messages, can now decrypt the messages and ensure the continuity of the sender over time.

The TESLA (Timed Efficient Stream Loss-Tolerant Authentication) protocol [130] can provide efficient broadcast authentication on a large scale, while coping with packet loss and real-time applications. The μ TESLA broadcast authentication protocol is the adaptation of TESLA for wireless sensor networks [131].

Both TESLA and μ TESLA use one-way key chains as shown in Fig. 6.3: The broadcaster chooses a random key K_n and applies a public pseudo-random function F as often as required to acquire the keys: $K_i = F(K_{i+1})$, $0 \leq i \leq n-1$. Subsequently, every secret K_i , $i > 0$ is used for sending in the i -th interval and disclosed to the

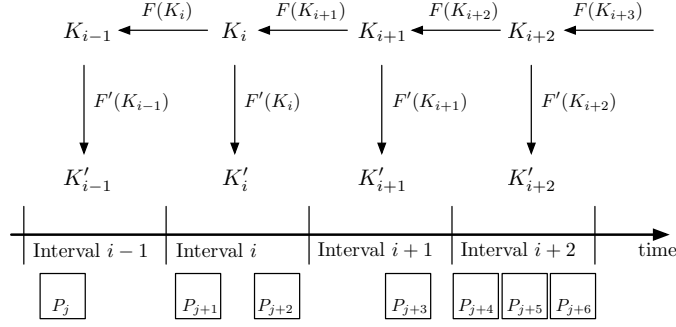


Figure 6.3: The figure illustrates TESLA's utilization of one-way chains after [130]. The first one-way function F generates the chain, following that the second one-way function F' derives the MAC keys. Time is divided into separate intervals i , all having the same length. The packets P_j are each sent during one specific interval. For every such packet, the sender computes a MAC with the key that is in accordance with that interval. E.g. P_{j+2} 's MAC is calculated based on its data and key K'_i . Disclosing the keys of previous intervals can be done either by attaching the key to sent packets or in separate messages.

public after a number of time intervals d . As every previous key can be recovered by the receiver(s) by applying the one-way function F , the receiver needs to do two things to authenticate a message: [132]

1. Authenticate the key K_i against previously received keys to ensure they are from the same key chain.
2. Ensure that the message with key K_i could only have been sent before the key has been published (requiring loose time synchronization), i.e. before interval $i + d$.

The fact that μ TESLA uses symmetric cryptography in connection with time as its asymmetric property makes it an interesting idea for adaptation to aviation since a sufficiently good time synchronization could be provided via GPS. The advantages of μ TESLA are obvious: ATC can keep its open and broadcast nature and a complex PKI infrastructure is not required to ensure a sender's continuity, although it could be added if identification and source integrity are required (e.g., for ground stations, which have good connections).

Another advantage of μ TESLA is that lost packets on the notoriously jammed 1090 MHz frequency (there is no medium access control in place) are not an

integral problem for authentication. Furthermore, the overhead for communication is significantly less than with traditional asymmetric cryptographic methods.

Haas and Yu [133] compare TESLA and ECDSA-based authentication by simulating their performance in a real-world VANET scenario (although not including certificate distribution). Regarding channel congestion and MAC layer delay, they found that the TESLA protocol with keys attached to a subsequent broadcast performed significantly better than ECDSA or a TESLA-scheme that publishes keys in separate packets. All these discussed characteristics make TESLA at least a serious alternative to more traditional PKI-structures for the aviation environment.

6.2 Secure Location Verification

Besides securing the communication — and thus the location data — of ATC protocols, there are other approaches to ensure the integrity of air traffic management. The general idea of secure location verification is to double check the authenticity of location claims made by aircraft and other participants. This is inherently different from the verification of the broadcast sources and messages. The baseline is to establish means to find the precise location of a sender, effectively offering some redundancy and thus the ability to double check any claims made. As an additional advantage any such approach creates more location data, which can be merged with SSR, ADS-B and PSR and offer a back-up system in case of failure of these primary navigation systems or even GPS. Overall, secure location verification can not protect the contents of air traffic communication and is hence less useful for data links and information services. However, it can be an incredibly helpful tool to maintain the safety of air traffic control

6.2.1 Multilateration

Multilateration, or hyperbolic positioning, is a popular form of *co-operative independent surveillance* and has been successfully employed for decades in military and civil applications. If the precise distance between four or more known locations

and an unidentified location can be established, it is a purely geometric task to find the unknown point. We can, for example, use the received aircraft communication signals which travel at the speed of light to estimate the distance. Since we do not know the absolute time a message needed to travel from an aircraft to a receiver, we have to employ the time difference of arrival (TDoA).²

Thus, multilateration requires a number of antennas in different locations that receive the same signal at different times. From the TDoA, hyperboloids can be calculated on which the aircraft's position must lie. With four or more receivers, a 3D position can be estimated by finding the intersection of the hyperboloids.

Performing multilateration by utilizing TDoA is currently the preferred solution for location verification on the ground. It is used in the field (for example by the ASDE-X system [135]) at various US airports and also being rolled out in Europe in connection with the CASCADE project.³ One major advantage of multilateration is the fact that it can utilize aircraft communication that is already in place. Thus, there are no changes required to the currently existing infrastructure in aircraft, while on the ground receiver stations and central processing stations have to be deployed (see Fig. 6.4).

While currently used mainly in comparably short distances around airports, wide area multilateration has also been a popular research topic. Compared to primary radar systems, wide area multilateration is relatively easy and cost-effective to install and use on the ground.

Despite its successful use in the field, multilateration leaves a number of open problems in terms of secure location verification and the ICAO names a few known drawbacks [136]:

1. It is susceptible to multi-path propagation.
2. A signal has to be correctly detected at comparably many receiving stations.
3. A separate link between central processing station and receivers is required.

²For a full explanation of the multilateration process in aviation see [134].

³<http://www.cascade-eu.org>

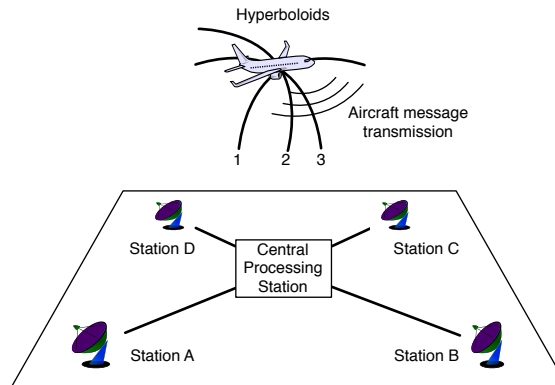


Figure 6.4: Basic multilateration architecture. Four (or more) receiver stations measure the time at which they receive the same message from an aircraft. They send this data to the central processing station which can calculate the aircraft's position from the intersecting hyperboloids resulting from the time difference of arrival between the receiver stations.

Essentially, many of the limitations stem from cost and logistic reasons as it can be difficult and expensive to deploy enough sensors and stations in very remote and inaccessible areas. Improving on these limitations could make TDoA-based systems a more widely useful approach as we will discuss in Section 7.5.

Furthermore, it is possible to attack MLAT, as pointed out in Section 4.2, although this involves both a certain cost and non-zero engineering knowledge. Thus, the difficulty of exploiting MLAT is relatively high, certainly when compared with the simplicity of spoofing the content of unsecured ATC protocols.

6.2.2 Distance Bounding

Distance bounding is another method that has been employed in wireless networks to partly localize other participants and ensure secure transactions, for example for radio-frequency identification (RFID) communication. First presented by Brands and Chaum in 1993 [137], the idea behind distance bounding is to establish a cryptographic protocol with the goal to have a prover P show to a verifier V that P is within a certain physical distance (see Fig. 6.5 for the concrete protocol). Similar to MLAT, the universally valid fact that electro-magnetic waves travel roughly at the speed of light c , but never faster, builds the foundation of all distance bounding

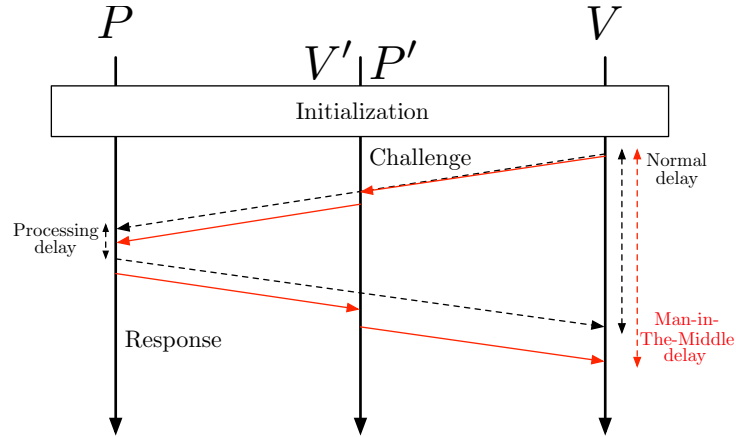


Figure 6.5: Principle of distance bounding protocols. The verifier V sends a challenge to the prover P , who, after processing, sends his response (black dashed arrows). A man in the middle (V'/P') can only increase the distance by adding further processing delays, but not decrease it (red arrows).

protocols.⁴ This enables the computation of a distance based on the time of flight between the verifier's challenge and the corresponding response by the prover.

In aviation, the determined distance could serve as an upper-bound, an additional piece of information that can subsequently be used as a means to verify and authenticate a node by checking the truth of its claims. When distance-bounding is performed by various trusted entities (such as ground stations) these can collaborate and find the actual location of the prover via trilateration.

Chiang et al. [138, 139] develop a secure multilateration scheme based on distance bounding that, under idealized assumptions, can detect false location claims with a high rate of success. By also taking into account differences in received signal strength (RSS), they can mitigate distance enlargement attacks and generic collusion attacks on the protocol. This shows how different physical layer techniques can be combined to improve theoretical security. However, the practical challenges in using such protocols in ATC are difficult to solve.

While in the literature distance bounding has been used mostly for close-up, indoor communication, it has been modelled for use in VANETs up to a distance of 225 m between prover and verifier [140, 141]. Tippenhauer and Capkun [142]

⁴This is in contrast to distance estimation via received signal strength, which can be influenced by a malicious node.

also considered the impact of moving targets on distance bounding protocols and verifiable multilateration. In their original implementation, it takes about 600 ms to perform a full localization, which at a speed of 500 km/h means that a target moves 75 m during the process. The authors propose using additional Kalman filters to smoothly keep track of the prover's location and detect malicious tampering by outsiders.

As an additional on-demand feature, distance bounding could provide crucial information about the legitimacy of nodes in areas where PSR is not present (or has been phased out due to cost reasons).

6.2.3 Kalman Filtering and Intent Verification

Kalman filters (also known under the technical term linear quadratic estimation) [143] have already seen extensive use in broader ATC applications, e.g. to filter and smoothen GPS position data in messages. Kalman filtering is used to observe noisy time series of measurements and tries to statistically optimally predict future states of the measured variables of the underlying system.

At a high level, the Kalman filtering algorithm comprises two distinct steps: a prediction step and an update step. As the procedure is recursive, it can easily be used and updated in real time, without having to save more than the last state.

Kalman filtering is a useful tool in general to predict the future values of a feature based on collected historical data. More concretely, it is used in ground systems to filter and verify the state vectors and trajectory changes reported by aircraft and to conduct plausibility checks on these data [144]. Krozel et al. [145] go on further to verify the intent of the aircraft by first defining local and global correlation functions to evaluate the correlation between aircraft motions and their intent (see Fig. 6.6). Using these methods, possible message injections could be detected as they deviate too far from the expected data.

From an attacker's point of view, Kalman filters can be tricked by a so-called frog boiling attack [146]: The adversary is jamming the correct signal, while continuously transmitting an ever-so-slightly modified position. If this is done slowly enough,

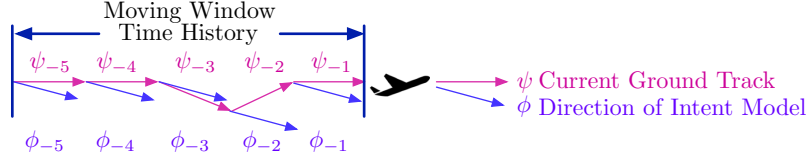


Figure 6.6: Practical application of intent verification in ATC after [145]. The example analyses the horizontal aircraft motion with a global correlation function as a moving window over local correlation functions.

the Kalman filter will see the injected data as a valid change. This attack has also been discussed recently by Chang et al. [147], who attempt to secure UAV under adversarial influence. This exposes a general weakness of Kalman filtering as the approach is based on comparatively little historical data. But it is still of great use since obviously bogus manoeuvres, velocities, or features can be detected and the complexity of an attack is greatly increased.

6.2.4 Group Verification

Group verification is another concept proposed to mitigate security and privacy concerns in ATC [85]. Illustrated in Fig. 6.7, it aims at securing the airborne SSR/TCAS communication by employing multilateration done by a group of aircraft to verify location claims of non-group members in-flight. A given authenticated group with 4 or more aircraft that have established trust can communicate with each other to utilize multilateration (based on TDoA or RSS) just as ground stations can. If a forged or otherwise inaccurate position report is detected, the sensible reaction would be to increase the circle of avoidance around nearby aeroplanes since their position cannot any more be regarded as precisely known and thus safe.

Kovell et al. [148] conducted a study about the applicability of the group concept in commercial aviation in the United States airspace. Examining the vast differences in traffic density over the US, they found that around 91% of aircraft at a given time could be part of a sufficiently large group of 4 aircraft or more.

Group verification has some downsides, too. First of all, it requires a new protocol to implement the verification and trust process. If such a protocol can be successfully implemented, there remains the central problem of how to manage

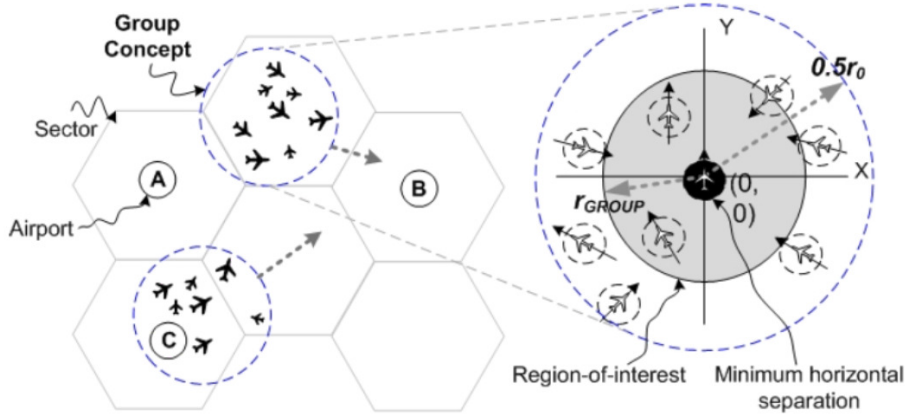


Figure 6.7: Illustration of the group concept from [85]. Four or more aircraft V are in any group G . Each group then can internally use multilateration to verify each other's location claims as well as those of outsiders in range. r_0 is the wireless communication range, $0.5r_0$ considered geographically proximate and thus acceptable for group establishment, given sufficient communication quality. To lower group overhead, the region of interest for a group can be restricted to r_{group} .

the secure authentication of members that are to be accepted into the group in the first place. It is very complicated to establish trust in new groups of aircraft and to reliably avoid malicious aircraft. Furthermore, the performance of the system in reaction to intelligent intentional jamming of some or all communication would have to be considered.

On the other hand, even without a perfectly secure solution, the group concept would raise the difficulty and engineering effort of certain airborne attacks by orders of magnitude.

6.2.5 Plausibility Checks

There are numerous relatively simple rules that can be utilized to locate potential red flags without resorting to more complex measures. No such rule is necessary nor sufficient by itself to detect an ongoing attack. But depending on the scenario and the attacker's savvy, they can indicate unusual behaviour that should be investigated either by a human or handled through further technical means. For example, it

is very plausible to outright drop a number of packets where either the data or the meta data is technically or physically impossible, illustrating a plausible attack scenario by less sophisticated threat agents such as script kiddies or hobbyists.

Red flags exist across various technologies and layers, from the physical to the application layer. Some are available to ground stations/air traffic control only, others also to aircraft-to-aircraft communication. Examples include but are not limited to [149]:

- Investigating **aircraft which suddenly appear** well within the maximum communication range of a receiver.
- Dropping **aircraft which are violating a given acceptance range threshold**, producing impossible locations.
- **Aircraft violating a given mobility grade threshold**, producing impossible minimum or maximum velocities.
- **Maximum Density Threshold**: If too many aircraft are in a given area, ATC software will typically alarm the user.
- **Map-based Verification**: Aircraft in unusual places such as no-fly areas or outside typical airways (this might possibly be better handled at the ATC software layer).
- **Flight plan-based Verification**: Flooded/attacked ground stations are able to check messages against the existing flight plan.
- Obvious **discontinuities in one of the data fields**.

Naturally, such potential red flags need to be handled with care and before any action is taken (such as dropping the packet/flight from ATC monitors).

	Injection / Modification	Eavesdropping	Jamming / Deletion
Physical Layer Authentication	+	-	-
Random DSSS/FHSS	-	+	+
Public Key Cryptography	+	+	-
μTESLA	+	-	-
Wide Area Multilateration	+	-	-
Distance Bounding	+	-	-
Kalman Filtering	+	-	-
Group Verification	+	-	-
Plausibility Checks	+	-	-

Table 6.1: Overview of capabilities of various security approaches against feasible wireless attack primitives on ATC protocols.

6.3 Summary

Table 6.1 provides a compact overview of the effectiveness of the examined solutions in combating the possible wireless attack primitives. We see that most security schemes focus on attacks of the message injection/modification class. This has two main reasons that have been mentioned throughout this survey: First of all, the open nature of ATC has been considered a desirable feature in most scenarios. So unless there is a major paradigm shift in the way air traffic control is handled currently, there is no interest in protecting against passive listeners, despite this being the first stepping stone for more sophisticated and problematic attacks. However, data links such as ACARS and CPDLC are widely used for confidential and safety-critical information and thus are in severe need of added confidentiality.

Second, passive attacks such as eavesdropping are simply much more difficult to protect against without having a full cryptographic solution. Similarly, attacks on the physical layer, such as continuously jamming the well-known frequency or the more surgical message deletion are hard to defend against, with measures on the same layer (e.g., uncoordinated spread spectrum) providing some of the only

	Difficulty	Cost	Scalability	Compatibility	References
Physical Layer Authentication	Variable	Variable	Variable	Additional hard-/software.	[116–120, 150–158]
Random DSSS/FHSS	Medium	Medium	Medium	Additional hard-/software.	[121–123]
Public Key Cryptography	High	High	Medium	Key distribution infrastructure and new technology needed.	[12, 124–129, 159–163]
μ TESLA	Medium	Medium	High	New message type and software required.	[115, 130–133, 164–166]
Wide Area Multilateration	Low	Medium	Medium	Separate system using existing technologies.	[134, 167–173]
Distance Bounding	High	Medium	Low	New messages and protocol needed.	[137–142]
Kalman Filtering	Low	Low	High	Separate software system.	[143–145, 148, 174]
Group Verification	High	Medium	Low	New messages and protocol needed.	[85, 148]
Plausibility Checks	Low	Low	High	Integration into current systems.	[149]

Table 6.2: Overview of feasibility attributes of various security approaches for air traffic protocols.

approaches to this general wireless security problem. The discussed approaches could however potentially address message insertion and tampering, either by protecting outright against it through integrity verification or by detecting anomalies in the data or its origin (e.g., Kalman filtering, multilateration).

Table 6.2 discusses the feasibility to implement each approach in the real-world aviation environment. As has been laid out, there is no single perfect solution that also fulfils the requirement of having little impact on the currently employed software and hardware used by ATC protocols. The choice is ultimately between a completely new protocol development, relatively lighter but still intrusive modifications such as new message types, or a transparent, parallel system which requires new software

and/or new hardware but does not change the existing environment.

Finally, the fusion of various ATC systems and their data (e.g., PSR, SSR, MLAT, ACARS, CPDLC) is an obvious and necessary idea for improving security. Yet, it is common knowledge in the aviation community that a major part of the business case for next generation ATC is based on reducing the reliance on analogue systems such as PSR [167]. Thus, in the longer term, it seems inefficient and not desirable to keep expensive legacy and/or new backup systems around, simply to improve the security of wireless communication technologies through redundancy.

Taking all these realities into account, we investigate cost-effective and transparent solutions aimed at fixing the existing security flaws in air traffic communication in the next chapter.

*Nothing travels faster than the speed of light, with
the possible exception of bad news, which obeys its
own special laws.*

— Douglas Adams' *Mostly Harmless*

7

Transparent Security for Air Traffic Communication

Contents

7.1	Modeling False-Data Injection Attackers	94
7.2	The OpenSky Network	95
7.3	Data Link Layer Fingerprinting	97
7.3.1	Feature Engineering	98
7.3.2	Experimental Design	100
7.3.3	Evaluation	102
7.3.4	Detection of Anomalies	104
7.3.5	Discussion	106
7.4	Physical Layer Intrusion Detection	107
7.4.1	Attacker Model	107
7.4.2	Feature Selection	108
7.4.3	Combined Anomaly Detection	111
7.4.4	Experimental Design	111
7.4.5	Evaluation	112
7.4.6	Discussion	114
7.5	Lightweight Aircraft Location Verification	115
7.5.1	Attacker Model	115
7.5.2	Considerations about Aircraft Localization	117
7.5.3	Designing Lightweight Aircraft Location Verification	122
7.5.4	Experimental Design	127
7.5.5	Evaluation	130
7.5.6	Discussion	136
7.6	Comparison of Transparent ATC Security	137

In the previous chapters we made our case for transparent, immediate security solutions that can help improve the security of air traffic communication and air traffic control in particular. In this chapter, we develop, implement, and evaluate different approaches to such transparent solutions, which could help detect message injections and anomalies in the air traffic surveillance picture. We focus on the ADS-B protocol as our case study, as it will form the basis of ATC in the next technology generation, and exemplifies the move towards digital broadcast networks. However, all examined solutions can principally be adapted to other protocols used to broadcast navigational data.

7.1 Modeling False-Data Injection Attackers

First, we describe the concrete model that an attacker uses to inject false data into an ADS-B target receiver. The injection of false data provides the basis of most of the attacks on the ADS-B system as discussed in our threat model (see Chapter 2). Executed correctly, such attacks are subtle but can have devastating effects on the system.

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

- **Replay Attack:** This attack captures real ADS-B data in the area and plays it back at a later time without modification. This is a traditional replay attack, which is trivial, considering the ADS-B protocol has no built-in authentication. Concretely, we assume that the attacker captures a given flight's ADS-B messages (positional, velocity, identification, and potentially others) and plays them back in the same order.
- **Message Injection:** This attack injects a new ghost aircraft created from scratch, by creating correctly formatted ADS-B messages according to the specified standards [57]. We also assume the attacker crafts messages with a

legitimate identifier and reasonable flight parameters (e.g., believable altitude and speed) to create an aircraft which is indistinguishable from a legitimate one for ADS-B-based radars. This forms the basis of virtual trajectory modification, virtual aircraft hijacking, and aircraft spoofing attacks [13].

For both scenarios, we adopt a non-naive attacker that has a sufficient amount of knowledge to inject valid-looking position messages. In other words, we assume these ADS-B messages are well-formed and their content is reasonable and able to withstand superficial sanity and validity checks.

7.2 The OpenSky Network

To enable experimental studies based on real data, we, in conjunction with TU Kaiserslautern, Germany and armasuisse, Switzerland, have developed the OpenSky Network,¹ a participatory sensor network for air traffic communication data, specifically ADS-B and Mode S. It provides access to 3 years of historical raw message data as well as metadata, and offers a very fast query infrastructure, ideal for large-scale research projects. As of October 2016, it has saved more than 300 billion air traffic communication messages, covers about 3,000,000 km^2 on four continents, captures more than 16,000 unique aircraft every day, and has seen over 110,000 different aircraft overall. Fig. 7.1 illustrates the system's coverage in Europe. For more detail on the creation of OpenSky, its use cases for aviation besides security, and its big data infrastructure, please refer to [29, 30, 60].

We use OpenSky's wealth of data as a basis to develop, implement, and evaluate different systems approaches to securing air traffic control protocols. We take into account that aviation is a very challenging real-world environment, where technologies are difficult to change, time horizons are measured in decades, and the cost factor is crucial to both airlines and airports. Thus, all solutions are based on OpenSky's current setup of cheaply available commercial off-the-shelf receivers, making them transparent, not costly, deployable on the ground, and reasonably

¹<http://www.opensky-network.org>

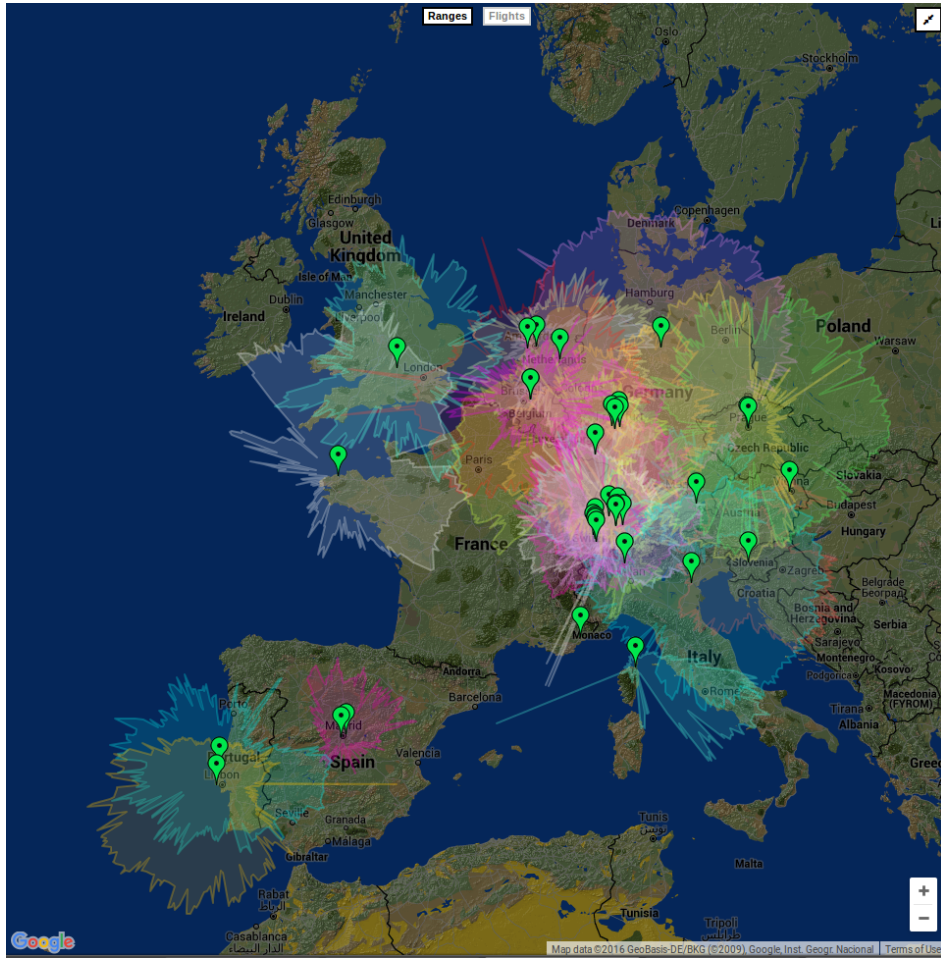


Figure 7.1: OpenSky sensor coverage in Europe (June 2016).

scalable. Where OpenSky’s current implementation of physical data collection was not enough, we used additional sensors in our lab, as indicated where appropriate.

Security Considerations

OpenSky not only forms the basis for our data collection, we also believe that crowd-sourced sensor networks will play an important part in the air traffic communication infrastructure of the future. Thus, we briefly consider some of the advantages and disadvantages, in particular concerning the security of such networks.

Besides the cost-effectiveness, one of the benefits is the robustness of the system that is inherent in its distributed nature. As COTS sensors are orders of magnitudes cheaper than traditional aviation equipment, many more of them can be deployed, safeguarding the system against the failure even of several sensors. It is further

irrelevant whether failure comes from natural causes or from a DoS attack. The massively-distributed fashion also makes attacks much more complex than attacking a mere one or two ATC receivers. Likewise, it is difficult for an adversary to obtain exact knowledge of all current sensor locations, which is required as a basis for many attacks. Mobile nodes could increase this location fuzziness even further in the future.

However, we have to consider potential new attack vectors that open up through the use of crowdsourced sensor networks. First of all, with precise, GNSS-powered sensors, an attack on the timebase (via spamming or spoofing) is principally possible. Of course, this is similar to existing radar or MLAT installations (see [175]). To defend against this, OpenSky supports non-GNSS sensors and is able to achieve time synchronisation using content of the ADS-B messages sent out by the aircraft [24]. Furthermore, as mentioned above, such an attack is complicated by requiring precise knowledge of all sensor locations. The other major attack would be insiders who sign up to OpenSky and feed false data to its database, as there are no integrity checks in ADS-B or other protocols, this is trivially possible. Defense mechanisms to such an attack include the creation of trust models which assign higher trust to known receivers and cross-check the data received for plausibility. Similarly, there are means to independently check the location claim of a sensor based on the data it sends, this could even detect larger Sybil attacks. While it is out of the scope of this dissertation, we will consider this problem in future work.

7.3 Data Link Layer Fingerprinting

Our first proposal is based on the analysis of aircraft messages to identify distinct differences between ADS-B transponder types and their implementations used in the commercial aviation market. We engineer several fingerprinting features based on transmission behaviour deduced from randomly chosen message inter-arrival times.

As is the case in many wireless networking ecosystems (see Section 6.1.1), these transponders can exhibit different behaviours on the data link level as well as the physical layer which can be utilized to distinguish incoming messages. In the

following, we identify and describe such differences on the data link layer, the selection of relevant features and the resulting classification.

7.3.1 Feature Engineering

A standard implementation of the ADS-B protocol broadcasts three types of messages in a regular manner:

- **Position messages:** The aircraft broadcasts a message with its own position on average every 0.5 seconds. A transmission mechanism is used to send the next message after a time interval randomly drawn from $[0.4; 0.6]$ seconds.
- **Velocity messages:** The aircraft broadcasts a message with its current velocity on average every 0.5 seconds. Similar to the position messages, the random message transmission interval is specified to be between 0.4 and 0.6 seconds.
- **Identification messages:** The aircraft broadcasts a message with its own ICAO 24-bit identifier on average every 5 seconds. Their transmission interval is randomly drawn between $[4.8; 5.2]$ seconds.

Through our exploratory research, we discovered a number of variations in the transmission periodicity across the aircraft data we analysed. The key insight here is that all major implementations do not use a truly random interval but instead use a number of possible *slots* placed more or less evenly throughout the specified interval. Over time, this leads to very different-looking distributions of the inter-arrival times of a given message type (see Fig. 7.2 for an illustration of some representative transponder behaviour).

The only information needed to build features based on this behaviour is a message's arrival time in the form of an absolute time stamp t_i . From this, we can calculate the inter-arrival time Δt between two subsequent messages from the same aircraft (as indicated by its transponder identification) by subtracting t_i from t_{i+1} . Indeed, while ADS-B is not encrypted, exploiting such timing and inter-arrival

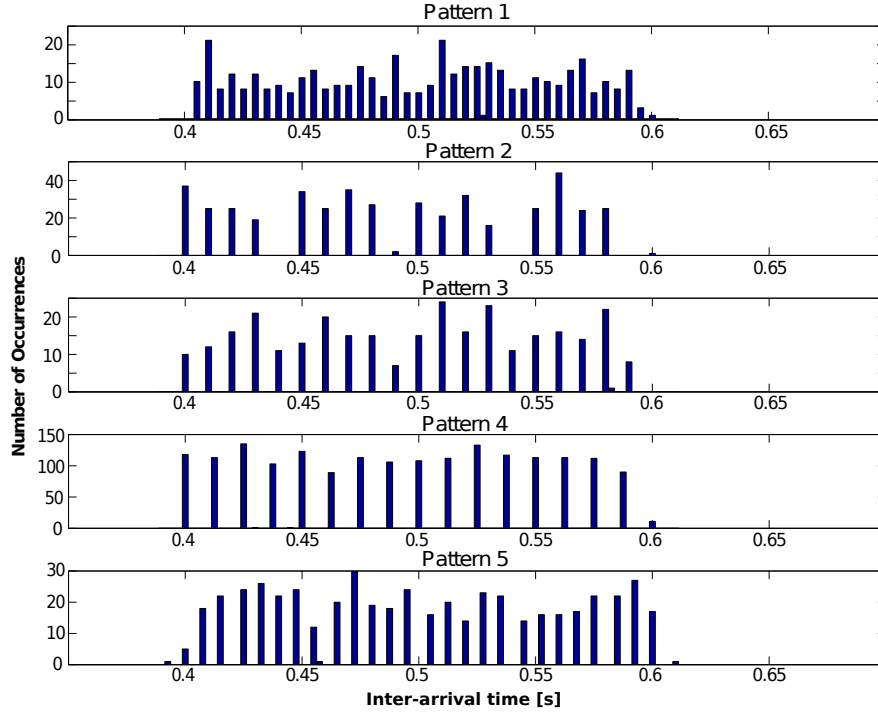


Figure 7.2: A representative illustration of five different transponder types. The graph shows the histograms of five time series of collected ADS-B position messages.

information between various message types is naturally possible even with fully encrypted messages when the same data link transmission patterns are followed.

Based on this, we develop seven distinguishing features:

- **Slot number:** The most obvious feature is the number of slots in the given interval of 0.2s for position and velocity messages and 0.4s for identification messages.
- **Slot width:** The second feature is the width of a slot. The time interval which constitutes a slot is defined by the minimum and maximum measured inter-arrival times of messages in this slot (rounded to what we believe is the actually programmed time).
- **Inter-slot width:** Analogous to the previous feature, there are intervals between slots which are not used for sending a message and thus empty (see Fig. 7.3 for an illustration).

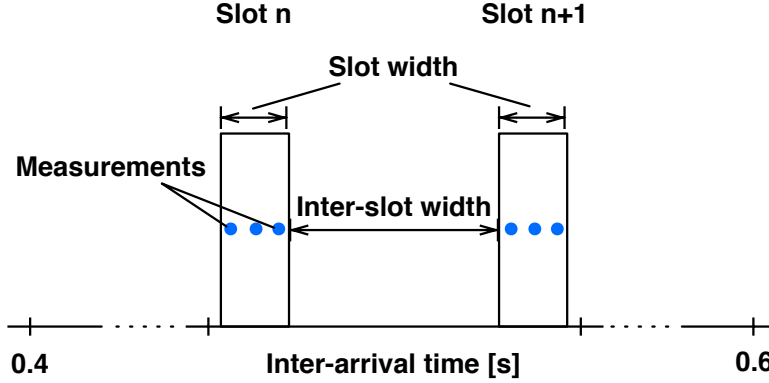


Figure 7.3: Schematic showing two slots as used by the transponder implementations, determined by measured inter-arrival times.

- **Missing slots:** Some implementations consistently do not use every fifth of their slots (concretely, those at 0.44, 0.49, 0.54 and 0.59 seconds for position and velocity messages). A subset of these uses the 0.59s slot but very sparingly.
- **No width slots:** Some implementations' first and last slots do not have a width. If these slots are chosen by the transponder, a message is sent exactly at 0.4 / 0.6 seconds respectively.
- **First slot:** Regardless of the slot pattern, transponders differ in the timing of the first slot that is being used, or in other words the actual minimum time interval Δt_{min} between two consecutive messages of the same type.
- **Last slot:** Analogous to the last point, transponders also differ in the timing of the last slot that is being used, or in other words the actual maximum time interval Δt_{max} between two consecutive messages of the same type.

7.3.2 Experimental Design

In this section, we describe our experimental setup, including the data collection process.

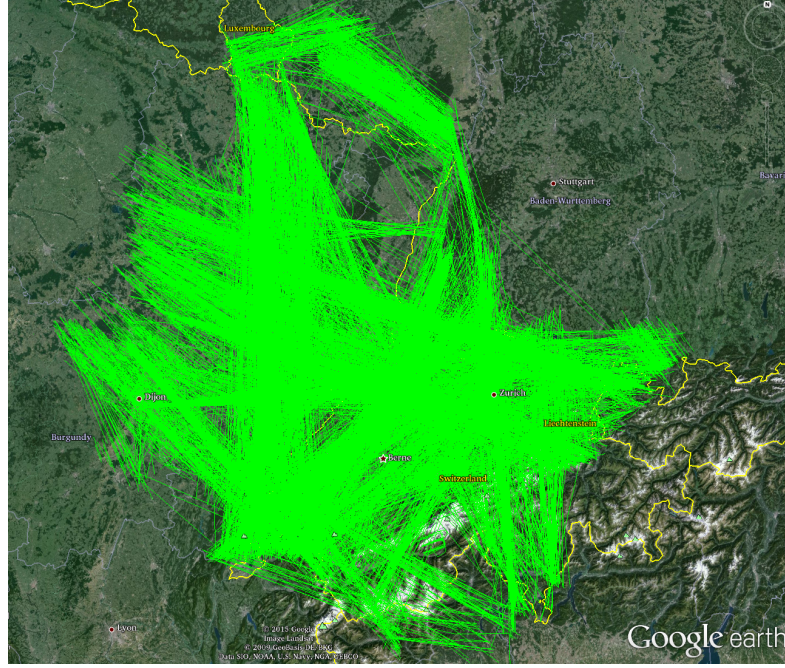


Figure 7.4: Exemplary visualization of 2910 of the flight trajectories used for our data analysis spanning roughly one day [25].

Flights	44,692
Unique ICAOs	4,997
# messages	30,772,643
Time frame	10 days

Table 7.1: Statistics about the utilized OpenSky dataset collected by a single sensor between November 9 and November 18 of 2014.

Data Collection and Hardware

As ADS-B has been in the roll-out phase for years, we can use data from actual aircraft collected in real-world wireless environments using the OpenSky project. For the present analysis (see also Table 7.1), we use a dataset that spans the period between November 9 and November 18, 2014. This dataset contains 30,772,643 ADS-B messages received from SBS-3 sensors manufactured by Kinetic Avionics. Besides the message content, they provide a timestamp of the message reception. The timestamps have a clock resolution of 50 ns. All sensors have omnidirectional antennas and can receive signals from a distance of up to 400 km. We stripped down the dataset to only use flights for our evaluation which had at least 200 received messages. The final data sample consists of 20,932 flights, a part of

	Slots	Width	Inter-slot	Missing	N/W slots	First	Last
Type 1a	39	$\pm 0.00025s$	0.005s	No	No	0.405s	0.595s
Type 1b	41	$\pm 0.00025s$	0.005s	No	Yes	0.40s	0.60s
Type 2	16	$\pm 0.001s$	0.01s	Yes	No	0.40s	0.59s
Type 3	20	$\pm 0.0005s$	0.01s	No	No	0.40s	0.59s
Type 4	17	$\pm 0.00015s$	0.0125s	No	Yes	0.40s	0.60s
Type 5	27	$+0.00016s$	0.008s	No	No	0.40s	0.608s
Type 6	81	$\pm 0.0005s$	0.0025s	No	No	0.40s	0.60s

Table 7.2: Feature combinations of different transponder implementations. Number of slots, first slot, and last slot features are given for position and velocity messages.

which is visualized in Fig. 7.4.

7.3.3 Evaluation

We analyse our data by sorting the flights according to different clusters according to their slot numbers, inter-slot distances, slot width, and missing slot behaviour. Based on these features, we discovered six main types of transponder behaviour in our dataset. Table 7.2 shows the values of the their inherent features.

We chose to further sub-divide the first type into 1a and 1b since they are very similar to each other in comparison with the other types, and differ only in additional slots at the beginning and the end. It is interesting to note that some of these implementation violate the original ADS-B specification of an inter-arrival time between 0.4 and 0.6 seconds. Some are in a clear violation such as type 5 which has its last slot at almost 0.61s, while others have their slots centred exactly at 0.4s (or 0.6s), causing one half of the slot to be outside the specified interval. This is presumably also the explanation for the “no width slots” feature found in type 1b and 4, which means the edge slots are both at exactly the edges of the defined interval.

We further found that the transponder software typically exhibits the same patterns within any of the three regular message types as shown in Fig. 7.5. More concretely, for the identification messages that are broadcasted every 5s, the described behaviours stay the same and are only spread out over an interval of 0.4s instead of 0.2s. On the other hand, we could not find any additional

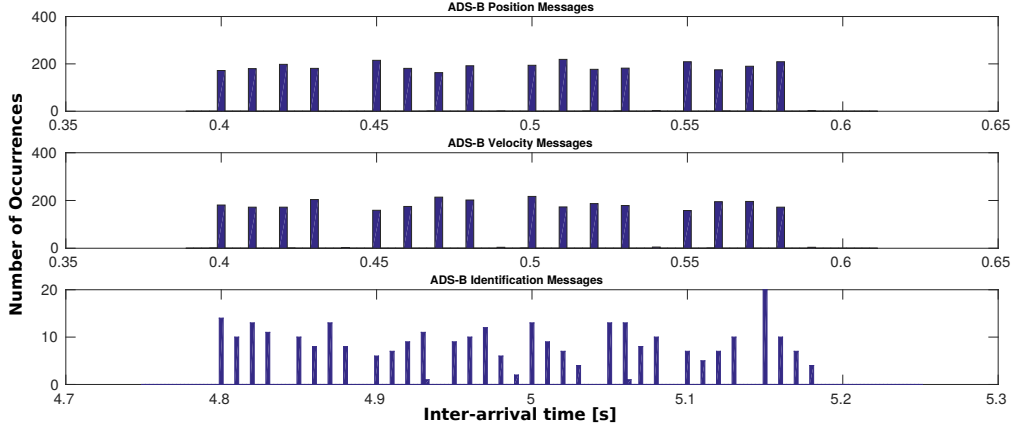


Figure 7.5: A representative illustration of the different message types. The graph shows the histograms of three time series of collected ADS-B position, velocity, and identification messages from the same flight using a type 2 transponder. The pattern characteristics are invariant across all types.

Type	1a	1b	2	3	4	5	6	Unknown
Occurrences	5246	1039	8778	5026	467	115	242	19
Share	25.1%	5.0%	41.9%	24.0%	2.2%	0.6%	1.2%	<0.1%

Table 7.3: Distribution of different transponder types in our dataset.

noticeable patterns in the inter-arrival times between different message types (e.g., between positional and velocity messages).

The distribution of the different identified transponder types is shown in Table 7.3. We can see that Types 1a, 2, and 3 make up about 90% of the market while types 1b, 4, 5, and 6 are much rarer.² At our chosen minimum level of 200 messages per flight, we have found 19 aircraft (or less than 0.1%) that did not match one of the 6 transponder type patterns. On manual inspection, the likely reasons included noisy measurements and possible transponder malfunctions. Of course, there could easily be other features or feature values that would lead to even more transponder (sub-)patterns.

²We conduct a more in-depth, not security-related, analysis of this result in the next chapter.

Feature	Stability
Slot number	99.3%
Missing slots	100%
First slot	96.5%
Last slot	94.4%
Overall classification	99.8%

Table 7.4: Time stability of the analysed features over different days when collecting 100 messages.

Feature Stability over Time

To verify the stability of our results over time, we trained all flights collected in our original dataset and looked for flights with the same ICAO identifier (i.e., the same aircraft) over the following week. With 1287 returning aircraft, the estimation of the transponder type stayed the same with 99.8% likelihood (see Table 7.4). We found that the misclassifications were generally caused by transponders moving between the very similar cluster types 1a and 1b, as the first and last slot features' accuracy was not sufficient. The last slot feature was the least stable with 94.4%, while the missing slot feature's presence did not change at all in our test.

7.3.4 Detection of Anomalies

Finally, we examine whether we can use knowledge about the transponder features of real aircraft to detect injections of false data. Our aim is to detect anomalies in the aircraft behaviour as a first step to detect potential attacks.

Attacker Model

To illustrate that it is possible to use transponder fingerprints as a defence mechanism, we assume the following simulated attacker models. Each attacker injects 200 flights into the receiving subsystem. Without loss of generality, we show that it is possible to detect attacks using one message type (positions) but the same applies to attackers using all three types; presumably the effectiveness would only increase when using more data.

- **Attacker 1** replays or creates and injects 200 messages by sending every 0.5 seconds. This assumes basic knowledge of ADS-B but not deep familiarity with the standard specifications.
- **Attacker 2** takes the same approach as attacker 1 but sends the 200 messages randomly between 0.4 and 0.6 seconds, as described in the ADS-B standard DO-260.
- **Attacker 3** replays a given aircraft using the exact same timings as received from the real aircraft but changes the broadcast ICAO identifier to a different aircraft from our sample in a bid to imitate it.

We add white Gaussian noise to the simulated timestamps of all attacker messages with a signal-to-noise ratio of $70dB$.

Attacker Detection

After simulating the three attackers defined above, we used the features based on their inter-arrival times to sort them into the seven potential transponder classes. If a message time series does not fit any of the expected classes, we assign them into the unknown class considered an anomaly or possible attack. Table 7.5 illustrates the results.

We can see that attackers 1 and 2, which are injecting ghost aircraft relatively naively without paying attention to typical ADS-B transponder implementations, are easily detected. All 200 attackers of type 1 are detected as anomaly, while 98.5% of ghost aircraft injected by attacker type 2, who injects messages uniformly in the expected interval, are detected as an unknown class. 1.5% of attacker 2's aircraft were considered an actual transponder type but not the known correct one (although there is a low probability that it would randomly be the true transponder type of a known aircraft).

Looking at the strongest attacker 3, we can also see the limitations of our approach. As her goal is to imitate real aircraft by replaying their position messages with the exact same timings as previously captured, we only observe 1.5% of

Type	True Class	False Class	Unknown Class
Attacker 1	0%	0%	100%
Attacker 2	0%	1.5%	98.5%
Attacker 3	27.0%	61.5%	1.5%

Table 7.5: Evaluation of attacker classes 1-3 using inter-arrival time-based fingerprinting.

attempts classified as unknown. A further 61.5% of attack attempts are detected because they use a different transponder type compared to the known ground truth about the imitated aircraft. This leaves 27% of successfully imitated aircraft, where the replayed transponder type corresponds with the real one. As is to be expected, this result is roughly in line with randomly picking two transponder types from the distribution shown in Table 7.3 and obtaining the same result. A more sophisticated and targeted attacker can of course take the transponder types of replayed aircraft into account and imitate the correct type every time.

7.3.5 Discussion

Our work on fingerprinting can lay the groundwork for an anomaly detection system to identify all manner of inconsistencies in the operation of air traffic communication, most notably detect potential intrusions. As we assumed, an attacker can create correctly formatted ADS-B messages in valid sequential orders and spacings making the injected aircraft indistinguishable from real ones using only standard ATC procedures.

Fingerprints of any kind can provide more, less obvious, characteristics which an attacker has to adequately mimic when inserting false data onto the wireless channel. As pointed out before, such an approach cannot provide guaranteed security. As with all types of fingerprints and anomaly detection systems, an attacker can learn the features of the system and adapt the injected messages to match the patterns expected by the attack detection system. This is no different with the data link features discussed in this work. However, we maintain their usefulness not only against naive threat agents such as hobbyists. The more features we track, the

more degrees of freedom we take away for the attacker, and consequently the more difficult it becomes to inject ghost aircraft without being noticed by the system.

7.4 Physical Layer Intrusion Detection

The second approach we examined is based on a physical layer property, namely the signal strength as obtained from the received ATC messages. The basic idea is to use RSS-based features to verify the authenticity of an aircraft and its location claims. We first define the attacker model based on their capabilities to manipulate the physical layer and their presumed course of action. We then introduce the derived features and how we combine them into a one class classification problem. Finally, we evaluate our approach by detecting the simulated attackers within a data set of legitimate aircraft tracks.

7.4.1 Attacker Model

We model the attacker's use of different RSS patterns using a single antenna. We assume all attackers are more or less stationary on the ground and attack specific sensors in transmission distance, i.e., we do not consider UAVs. Weather effects on RSS have proven negligible for our use case [24].

- **Attacker 1:** This attacker uses a straightforward constant sending strength, resulting in a Gaussian distribution due to the noisy nature of the channel. Without loss of generality, we assume the standard settings of a typical software-defined radio with a 100 mW power output and a distance of 500 m to the sensor under attack. This creates a signal with a RSS of about -65 dBm at the receiver; the standard deviation of the random noise is 3.5 dB.
- **Attacker 2:** The RSS is a random variable X , within the limits of the hardware. To simulate a random, stationary, non-adjusting attacker, we assume the RSS received at the attacked sensor to be fully random within the typical values of legitimate aircraft (in our case, the 5%/95% percentiles are -75.60 dBm and -63.54 dBm, respectively).

- **Attacker 3:** This attacker adjusts the sending strength in an attempt to be in line with the position the injected messages are representing to the attacked sensors. More concretely, the attacker knows the position of the receiver with a maximum error of 1 km (mean: 500 m) on which he bases the calculation of the distance to the claimed flight positions.

Our goal is to get an accurate read of legitimate aircraft behaviour, enabling us to detect all but the most knowledgeable, powerful and carefully carried out attacks by entities who have perfect knowledge of the system **and** its sensor locations.

7.4.2 Feature Selection

In this section, we describe the physical layer features that we select for our intrusion detection system (IDS) and how we combine them in a unified detection approach. When receiving ADS-B messages from an aircraft, the ground station can measure and store the RSS. Due to the attacker's positioning on the ground, the measurements of injected ADS-B messages are highly unlikely to match the RSS of legitimate samples. Furthermore, they should be comparably constant over time compared to aircraft covering distances of hundreds of miles in relation to the receiver. Using standard hypothesis testing, we can judge the probability whether a collected RSS sample stems from a legitimate aircraft or not.

Pearson Correlation Coefficient

In physical space, we calculate the Pearson correlation coefficient ρ between the distance (derived from the position claim in the ADS-B messages) and the RSS. Signal propagation theory (i.e., path loss) would suggest a strong negative relationship in legitimate flights, while an injection attacker, who does not adjust the sending strength in line with the claimed distance, should show no correlation. Formally, we test the null hypothesis \mathcal{H}_0 , which states that there is no association between the two variables in the population, against the alternative hypothesis

\mathcal{H}_A , stating that there is a negative association between the two variables in the population:

$$\mathcal{H}_0: \rho = 0 \quad (7.1)$$

$$\mathcal{H}_A: \rho < 0 \quad (7.2)$$

We consider a sample where \mathcal{H}_0 is rejected at the 99% significance level a legitimate flight sample and an attack if the hypothesis is accepted.

Autocorrelation Coefficient

In signal space, we use the autocorrelation coefficient (ACF) to identify attackers that are stationary and/or do not adapt their sending strength. Autocorrelation is the cross-correlation of a signal with itself. It can be used to show that a time series is not random but instead exhibits significant correlations between the original observations and the same observations shifted backwards by a discrete lag τ . The ACF helps to find repeated patterns such as periodic signals in a noisy channel. Formally, we test the null hypothesis \mathcal{H}_0 , which states that there is no autocorrelation $R(\tau)$ in the population, against the alternative hypothesis \mathcal{H}_A , saying that there is a positive autocorrelation:

$$\mathcal{H}_0: R(\tau) = 0 \quad (7.3)$$

$$\mathcal{H}_A: R(\tau) > 0 \quad (7.4)$$

We run these tests for lags $\tau = 1$ to 8 and take their mean to create a single measure for finding autocorrelation significant at the 1% level. We again consider a sample where \mathcal{H}_0 is rejected at the 99% significance level a legitimate flight sample and an attack if the hypothesis is accepted.

Detection of Multiple Antennas

Legitimate ADS-B-equipped flights send alternatingly using two separate antennas, one on top of the aircraft and one on the bottom, as specified in [57]. This setup

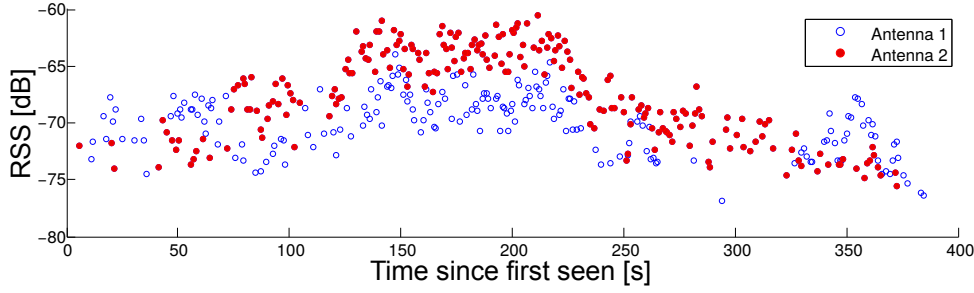


Figure 7.6: RSS samples of a flight's two separate antennas.

creates a behaviour that a sophisticated threat agent needs to mimic. Fig. 7.6 shows an example of the distinctive RSS patterns. To exploit this feature, we divide the full RSS time series into their two antenna sub-parts according to their time slots and compare various features that show only on the newly created time series. To illustrate the difference: with 300 samples per flight, we found a difference of around 1.8 dBm ($\sigma = 1.4$) in the means of the two antennas in our sample data. A single-antenna attacker, who does not adapt his sending power to mimic two antennas, is expected to exhibit no significant difference between the RSS of messages in alternating time slots. Based solely on RSS time series, we can identify other differences between a single-antenna user (i.e., an anomaly that would most likely be caused by an attacker) and messages sent out by commercial aircraft:

- The ACF of the two divided antenna time series falls much faster for legitimate aircraft than for a single-antenna attacker.
- For legitimate aircraft, even lags (2, 4, 6, 8) of the ACF of the (original) combined time series are greater than odd lags.
- Similarly, the ACF for a lag of 1 is typically higher for the divided antenna time series, while for an attacker the ACF for the divided and the combined time series are similar.

On a different note, separating the antennas first vastly improves the results of the correlation features discussed previously.

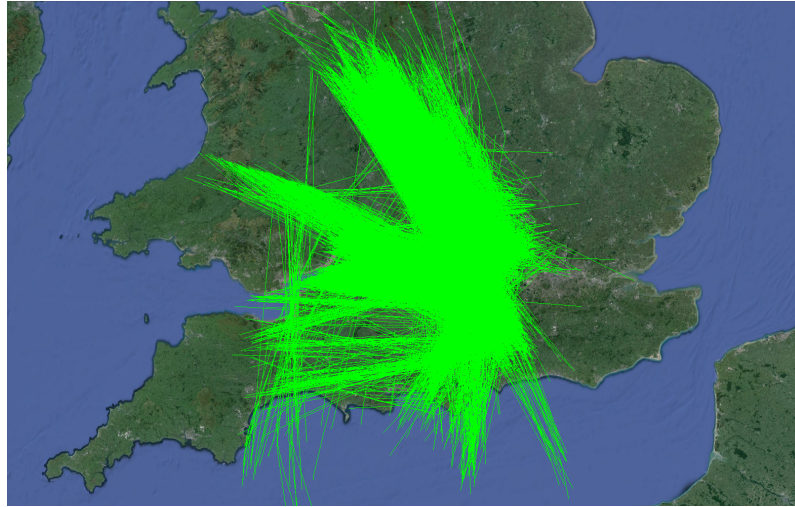


Figure 7.7: Visualization of the 7,159 flight trajectories used for our anomaly analysis.

7.4.3 Combined Anomaly Detection

To further improve our results, we combine our features in a one-class classification problem. One-class classifiers try to separate one class of data, the target data, from the rest of the feature space. Our target class is a well-sampled class of aircraft behaviour based on collected RSS data. The outlier class is unknown and online target samples are used at the time of learning. The process creates an n -dimensional classifier, where n is the number of features. For new samples, this classifier decides if they fit into the expected space or if they are rejected (i.e., classified as an anomaly worth investigating).

7.4.4 Experimental Design

First, we analyse the effectiveness of our selected features on their own, using standard hypothesis testing, before we combine them with a machine learning approach to create a more robust IDS. We employ the MATLAB toolkits Dd_Tools and PRTools³ to create data descriptions of our air traffic data. We define one-class datasets based on legitimate data collected with an ADS-B sensor and use various one-class classifiers to create descriptions which include the data.

³See http://prlab.tudelft.nl/david-tax/dd_tools.html and <http://prtools.org>.

Detection Rate [%]	Attacker 1	Attacker 2	Attacker 3	Legit Flights (FPs)
Pearson	99.8	99.9	0.2	18.6
Autocorrelation	99.6	99.4	0.3	0.1
Antenna Detection	92.6	94.0	95.5	3.9
Combined Detection	100	100	98.8	<0.01

Table 7.6: Effectiveness of the examined detection approaches. We used 7,159 legitimate flights and 143 simulated attackers for every class, with 200+ messages per flight. The percentages show the average detection rates over 5-fold cross validation.

Data

We used a data sample consisting of 7,159 flights, each flight with 200 or more received messages, collected over 24 hours and visualized in Fig. 7.7. The data collection was conducted with an OpenSky sensor installed at the top of our lab building.

For our anomaly detection approach, we test several different classifiers with 5-fold cross validation and the fraction of outliers in training set to zero (i.e., all training samples are accepted as legitimate). While the training sets are drawn from our collected sample of legitimate flights only, the separate test sets for each attacker have an added 2% of falsely-injected data (amounting to 143 flights) to be detected by the classifier. To verify our models and test our system, the RSS patterns of the attackers are simulated as described in Section 7.4.1.

7.4.5 Evaluation

Table 7.6 shows the results of the examined detection approaches. The hypothesis tests each detect attackers 1 and 2 with more than 99% probability. Especially the autocorrelation feature proves to be accurate, with few legitimate flights misclassified as false positives (0.1%). As expected, both tests fail to detect the more sophisticated attacker 3. To counter this, we analysed the distinct antenna characteristics, with which we can detect over 90% of all three attackers with a false positive rate of 3.9%. On its own, the antenna method requires 300 messages to become reliable enough, as aircraft may move in ways that can obfuscate their antenna features in the short run.

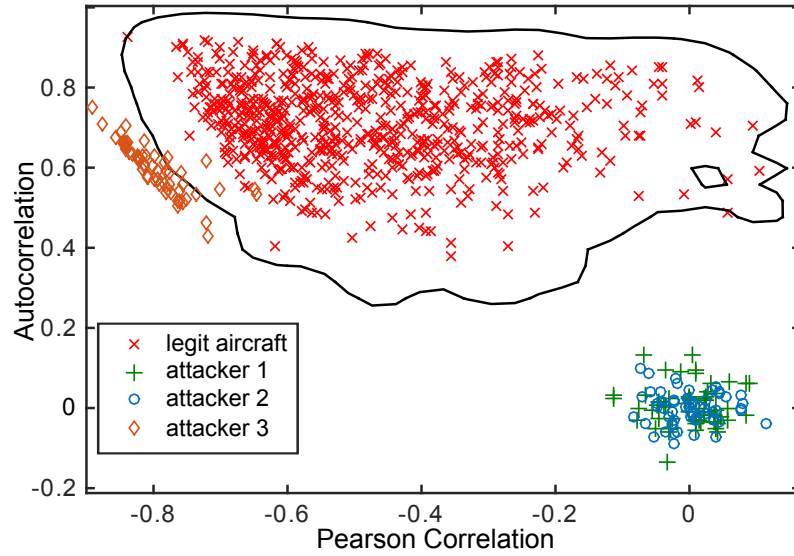


Figure 7.8: Example of a 2D-Parzen classifier used for anomaly detection with 200 collected flight samples. Red crosses are samples of legitimate flights. Attacker 1 and 2 are entirely classified as anomaly here, while attacker 3 creates few false positives.

With the combined classifier, we can accurately detect all attackers 1 and 2 without false negatives and one single false positive (less than 0.01%), using a small RSS sample of 200 messages. At the standard rate of 6.2 ADS-B messages per second, this allows detection in under 40 seconds, assuming no message loss. Even when considering a typical message loss of 30% [24], this can be achieved in less than one minute.

As illustrated in Fig. 7.8, attacker 3, who easily deceives the individual hypothesis tests, can be *too* good. He would need to introduce additional randomness and patterns similar to the spoofed aeroplanes to fall within the expected data range. This demonstrates the strength of the anomaly detection approach where the precise type of anomaly need not be known in advance.

Fig. 7.9 shows the results of the comparison between various tested classifiers, depending on the number of samples. The Parzen classifier performs best, having the lowest number of misclassified attackers. It is followed by K-Means, but the Minimax, Minimum Spanning Tree and k-Nearest Neighbours classifiers also achieve a near-zero false negative rate as 200 samples are collected, still a significant improvement on pure hypothesis testing.

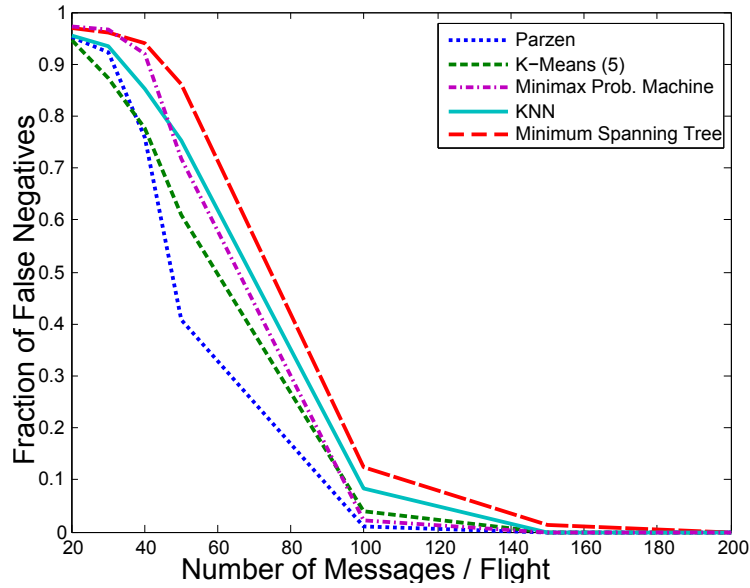


Figure 7.9: Complete classifier comparison with 5-fold cross validation. Shown are the joint false negative rates for attackers 1 + 2.

7.4.6 Discussion

As we have shown, RSS-based detection is able to reliably detect different ground-based attackers, who attempt to inject or replay ghost aircraft. While those with non-adjusting sending power strategies are straightforward to detect even with pure hypothesis testing, employing anomaly detection makes the results far more reliable in terms of sensitivity and specificity. With the added antenna features, an attacker experiences even further complexity in conducting a successful injection attempt and can be detected in acceptable time spans. In general, more samples naturally increase the confidence of the system and improve detection results at the cost of slower reaction times.

While not a countermeasure that is impossible to overcome with knowledge of sensor locations and detection strategy, schemes such as the presented one can greatly increase the security of ATC systems without introducing new technologies. Similar to the previously presented fingerprinting approach, added physical layer-based security can be obtained by a single sensor, greatly reducing deployment and maintenance costs for air navigation service providers (ANSP). It further enables the full benefits of the new ADS-B technology as the full range of the received

data can be verified and subsequently used, not only those messages from aircraft that were received by multiple sensors – as is the case with other physical layer schemes such as multilateration. This dimension of sensor requirements is analysed more thoroughly in the next section.

7.5 Lightweight Aircraft Location Verification

Besides the use of the received signal strength as in the previous section, other physical layer characteristics lend themselves towards building transparent security systems for ATC. In this section, we examine the time of flight of the signals and the difference between them. We develop systems that can verify the accuracy of the location claims of an aircraft and potentially locate the attacker, if a signal is received by enough sensors. We evaluate our system with real-world aircraft data obtained from OpenSky and compare it with the popular multilateration approach (see Section 6.2.1), which is based on the same physical layer features.

7.5.1 Attacker Model

For our location verification, we consider that potential attackers have different mobility models which can influence the temporal credibility of their positional claims as their physical positions and signal characteristics change. In this section, we provide a concrete description of the attackers' characteristics (see also Fig. 7.10 for an illustration).

Ground-based and stationary

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

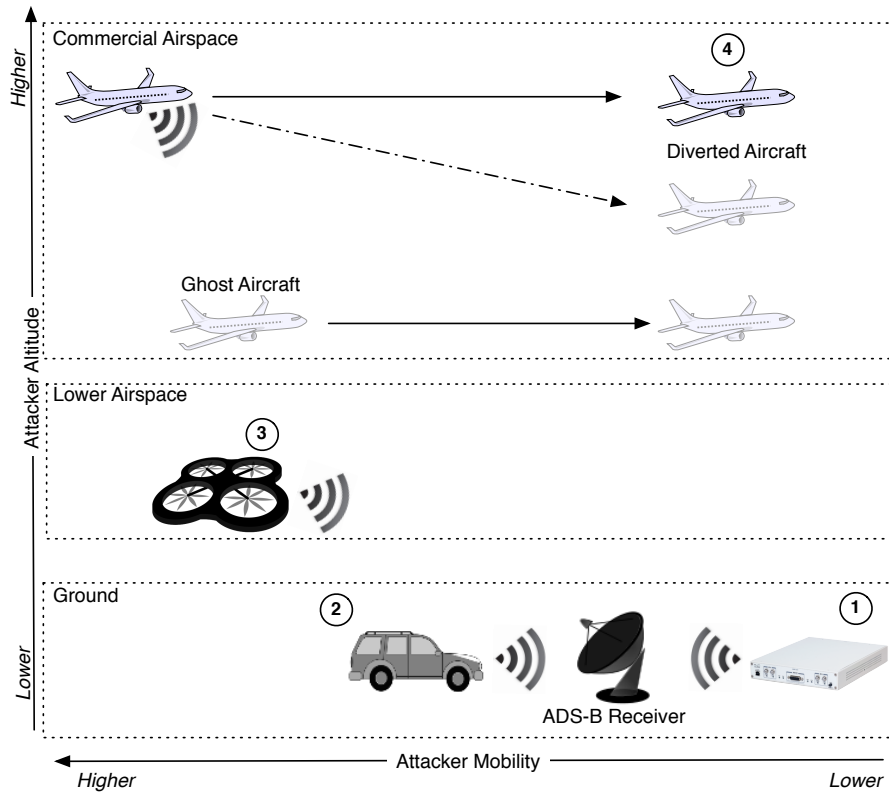


Figure 7.10: Graphical overview of the mobility models of four distinct attacker types. Attacker 1 is stationary on the ground, attacker 2 is mobile on the ground, attacker 3 has mobility up to a few hundred meter above ground. All three attackers inject a ghost aircraft onto the channel. Attacker 4 is a commercial aircraft using its legitimate transponder to send out wrong ADS-B messages to conceal its true position.

Ground-based and mobile

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

Low airspace and mobile

Attacker 3 is mobile within the limits of a typical unmanned aerial vehicle. Without loss of generality, we assume a hand-held commercial UAV, for example a standard

model working within 2 km range, up to an altitude of approximately 600 m and with a vertical top speed of 100 km/h. In general, a UAV is a versatile ADS-B sender and a much more flexible tool for an attacker than ground-based solutions. The airborne attacker seeks to emulate the physical characteristics of a commercial aircraft (or other UAV, as their use of ADS-B for navigation and collision avoidance will become standard in the future) much more closely than the previous threat models.

High airspace and mobile

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

7.5.2 Considerations about Aircraft Localization

In this section, we discuss several characteristics inherent to the aircraft localization problem. We also examine the properties of MLAT in this context, a well-established navigation technique in aviation that can independently verify aircraft positions. As we argue, while still a viable solution in some areas, the real-world applications of MLAT in the air traffic surveillance space are considerably limited and require significant improvement.

Problem Characteristics

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

- **Outdoor line-of-sight environment:** Contrary to many location estimation and sender verification problems found in academic research, the aircraft location problem is naturally outdoors. On the 1090 MHz channel, the line of sight (LoS) is a crucial factor in receiving signals. We require an outdoor LoS propagation model for our work in terms of loss and propagation.
- **Vast distances:** In *wide area* surveillance, the distances covered are naturally much larger than in more local or indoor problems. Aircraft flying at cruising altitudes (typically 35,000 feet or higher for commercial aircraft) can be observed up to the radio horizon of 400 km or more. This is orders of magnitude larger than typical indoor location problems. Our approach can easily be adapted for airport surveillance, too.
- **Few multipath effects:** At typical aircraft cruising altitudes, we experience comparably few diffractions leading to multipath effects that influence signal characteristics. This enables us to use simpler theoretical models than in more complex indoor and multipath-rich environments. Most importantly, the propagation timings between aircraft positions and sensors can be approximated easily by using the speed of light c .

The Drawbacks of Multilateration

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

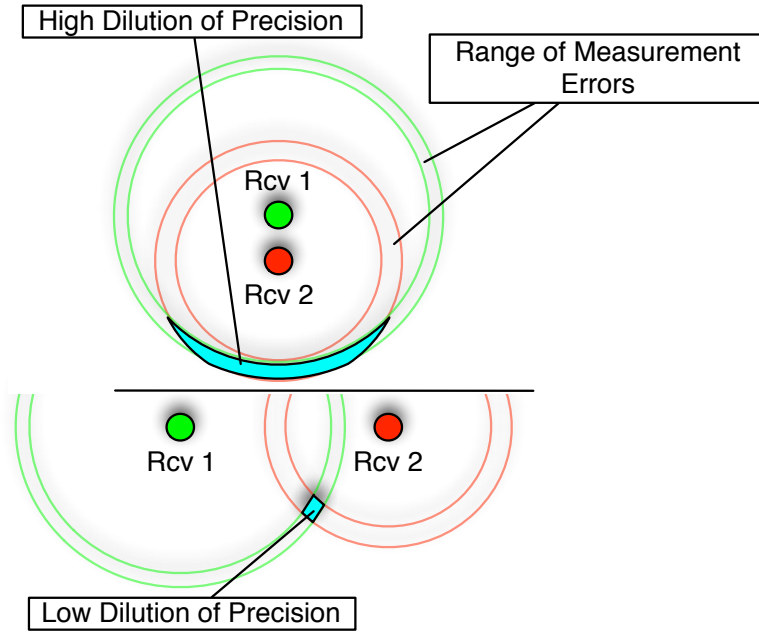


Figure 7.11: Geometric dilution of precision. The circles show the measurement errors of the respective receivers; the intersections demonstrate the area where the true location of the measured object can be found. An adverse placement of receivers in relation to the target (top) can severely affect the outcome of the localisation compared to a favourably deployment (bottom).

1. MLAT is highly susceptible to noisy environments and even relatively minor measurement errors outside a small core area. An important quality metric for a deployment and its MLAT accuracy with respect to the target object's relative position is the *geometric dilution of precision*, or GDOP (illustrated in Fig. 7.11). It describes the effect of a deployment on the relationship between the errors of the obtained TDoA measurements and their resulting impact on the final errors in the object's calculated position, or formally:

$$\Delta Location Estimate = \Delta Measurements \cdot GDOP$$

GDOP is widely used in positioning systems such as GPS, where good ratings for this multiplier are commonly considered to be below 6, with 10 to be fair and everything over 20 to be of poor quality [176].

2. Theoretically, four or more sensors are sufficient to compute a position of an object in 3D space. However, it is very difficult to get the precise altitude of an aircraft when all the receivers are on the ground (i.e., in one plane) and do not

provide sufficient elevation angle diversity. In that case, the *vertical* dilution of precision (VDOP) may be too large, so that only horizontal coordinates are calculated for aircraft surveillance and the altitude must be obtained by other means [173].

3. As a non-functional drawback, MLAT systems are very expensive: Where ADS-B needs only one receiver for accurate wide area surveillance, MLAT requires every signal to be received by at least four stations with little noise. Geographical obstacles (e.g., mountain ranges, oceans) make it even more difficult to install a comprehensive wide area system at the desired service level.
4. A determined and resourceful attacker could spoof wireless signals such that using their TDoAs for localization would result in a position of the attacker's choice [83]. While a system of many MLAT receivers can counteract this by severely restricting the attacker's freedom [31], this would also further increase the cost of multilateration.

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

Scalability and Coverage with TDoAs

One of the main goals of our design is to tackle MLAT's scalability and coverage problems. An ATC data communications network consists of a given number of sensors that are deployed outside, in a line of sight with the airspace they are expected to cover. Naturally, overlapping reception ranges between receivers are required to obtain TDoAs. If more sensors are to receive the same message, they need to be located closer together. While this increases the overlap, it also decreases the overall ADS-B coverage of the receivers. Worse even, only a small part of the MLAT coverage is usable, since GDOP causes its accuracy to deteriorate quickly.

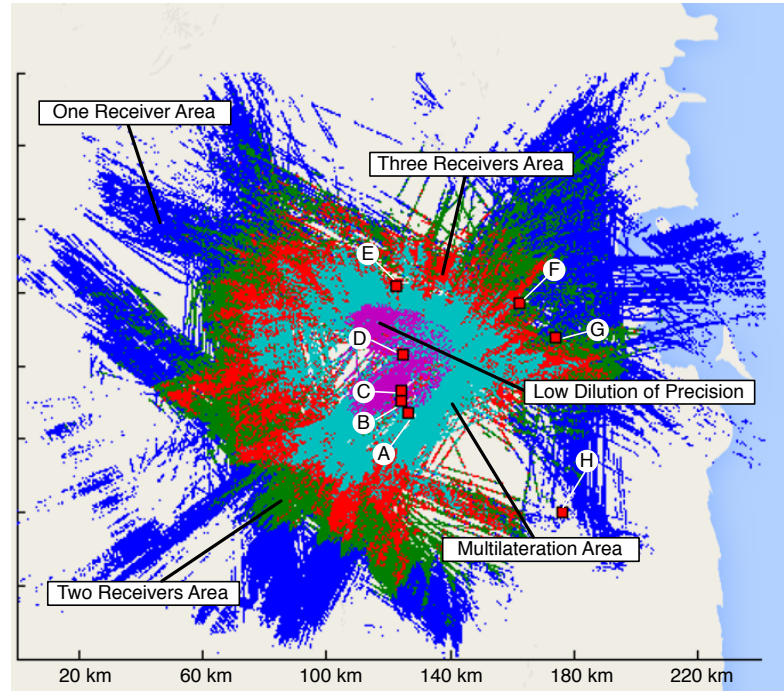


Figure 7.12: The map shows the practical reception ranges of 8 co-located OpenSky sensors. The turquoise part is the MLAT-capable area, the purple centre shows the area with acceptable (i.e., $DOP < 10$) accuracy.

Methods not suffering from GDOP and working with fewer sensors could vastly improve security compared to MLAT.

To demonstrate this fact, we analysed more than 50 million ADS-B messages from aircraft at cruising altitudes (ca. 38,000 ft) with using 8 co-located OpenSky sensors in Switzerland. Fig. 7.12 graphically illustrates the regions where messages are picked up by a given number of receivers. It also depicts the MLAT-capable area which makes up roughly 5% of the overall covered area (see Table 7.7). The area where MLAT is reliably accurate is even smaller at around 0.37% of total coverage. When we look at the relative number of messages that can be used for verification purposes, this becomes even clearer. Less than 4% of all received messages are seen by 4 or more sensors on the ground and can be used for MLAT. If we take into account dilution of precision, we are left with only 0.36% of usable messages. While these numbers concern a natural deployment under real-world constraints, this does not change significantly even in simulations with near-optimal coverage (e.g., rectangular or triangular, as discussed in [177]).

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

Table 7.7: Statistics on OpenSky dataset used for aircraft location verification. The table shows the absolute and relative number of messages collected by a given amount of sensors. The last column provides the relative area covered by that number of sensors.

Of all analysed legitimate flights for which we received more than 100 messages, 87.7% had at least 10 messages received by 2+ sensors, 65.37% by 3+ sensors and only 9.73% were MLAT-capable.

7.5.3 Designing Lightweight Aircraft Location Verification

Based on our analysis, we propose two other TDoA-based methods to verify aircraft location claims: a grid-based k-NN approach and statistical verification. Both do not suffer from dilution of precision and work with as few as 2 sensors, increasing the effective coverage of our deployment by a factor of >100 , thus vastly reducing costs.

Expected TDoAs

The key insight of our approach is the use of *expected TDoAs*. As outlined before, the time differences of arrival of a received signal between multiple sensors are a physical primitive that can be used to establish the possible location(s) of a sender. Since we know the location (claim) of an aircraft from its ATC communication, we can extend this concept and (pre-)calculate the time differences that we would *expect* to see based on any given location.

Concretely, we can use a simple outdoor LoS propagation model suitable for the aircraft location verification problem. It is straightforward to calculate the absolute propagation times of an ATC signal to the receiving ground stations by dividing the distances d_1, \dots, d_n between the sender's claim and each of the stations $1, \dots, n$

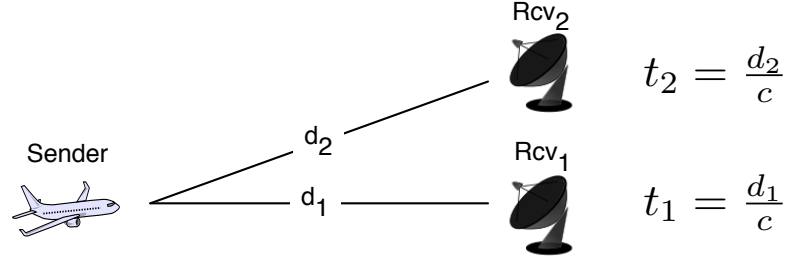


Figure 7.13: An example illustrating the calculation of expected TDOAs. The assumed distance of the sender to both receivers is multiplied by c . Subtracting the smallest time t_i from the other times gives the TDOAs relative to receiver i .

by the speed of light c (see Fig. 7.13).⁴ Subtracting the smallest resulting time t_i from the other times gives the TDOAs relative to the nearest receiver i .

We use the calculation of expected TDOAs to both verify and estimate the location claims of real aircraft and attackers.

Location Verification with Expected TDOAs

For our location verification approach, we use an offline phase to learn the distribution of errors between expected and actual TDOAs between our sensors.

In the online phase, we use the non-parametric Wilcoxon rank-sum test to check if the received sample distribution matches the expected distribution. By establishing the proximity to the expected data distribution, we can validate the sender.

Offline phase In the offline phase, we use training data to learn the deviations of real and expected TDOAs between two or more sensors. This creates a distribution of errors specific to a sensor, taking into account all the real-world noise introduced through propagation, synchronization, and measurement errors. The distributions as created by our setup are leptokurtic with a mean of 0 and a standard deviation of approximately 1 microsecond.

Online phase In the online phase, we test the likelihood that the measured TDOA of a received message is valid. The deviation between the expected TDOA based

⁴As the propagation is not happening in a vacuum, this is an approximation (however, the difference is insignificant [178]).

on its positional claim and the actual TDoAs must conform to the distribution of our collected data for any receiving sensor. We can gain stronger confidence over time by collecting more samples, effectively dealing with outliers and minimising false positives.

To check if the measurements match the expected distribution, we employ the non-parametric Wilcoxon rank-sum test to test the null hypothesis

\mathcal{H}_0 : *The sample comes from the same distribution as our training data.*

against the alternative hypothesis

\mathcal{H}_A : *The sample comes from a different distribution than our training data.*

(in other words, they are sent from a source not actually at the claimed position) at a 99.99% significance level. Compared to other distribution or location tests, the Wilcoxon test is more robust on non-normal distributions as we experience them with our setup.

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

Location Estimation with Expected TDoAs

Indoor and outdoor localization problems have been studied extensively in the literature, often in the scope of sensor networks and radar applications. Liu et al. [179] give an overview of the techniques used in wireless indoor positioning including the different algorithms (k-Nearest Neighbour, lateration, least squares and Bayesian among others) and primitives such as RSS, TDoA, time of arrival (ToA) and angle of arrival (AoA). While TDoA systems are limited in indoor environments (due to multipath effects and non-availability of time synchronization

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

```

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

```

and clocks fine-grained enough to provide good results at very short distances [180]), they offer very good performance in long-distance outdoor line-of-sight environments such as those encountered in the aircraft location problem.

In terms of algorithms, the k-Nearest Neighbours (k-NN) has proven to do very well in short-distance, indoor RSS fingerprinting compared to other methods [181], although it has been studied less in long-distance scenarios and can become computationally expensive with large databases.

Putting these findings together, we design a novel approach to locate aircraft by creating a 2D grid that contains expected TDoA measurements for each position for a given sensor deployment. For every incoming message, the nearest neighbours of the messages' TDoAs on the grid are calculated. Then the position of the signal is computed and the result is compared with the position given by the aircraft. When the estimate deviates too far from the claim, an attack is likely. Furthermore, these estimates can be used to roughly locate an attacker.

Similarly to the location verification approach, we use an offline training phase while the online phase continuously verifies new aircraft.

Offline phase Over an exemplary grid of $N \cdot M$ squares, we calculate the fingerprint vector of expected TDoAs between the deployed sensors for every square. Based on this, we create the final training set by generating every subset

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

```

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

```

of combinations with at least 2 sensors ($\sum_{i=2}^n \binom{n}{i}$, with n being the number of sensors), e.g., 26 sets overall for a 5 sensor deployment. This is required when a new online message is received by fewer than the maximum number of deployed sensors. In that case the appropriate set needs to be chosen to find the k nearest neighbours. Algorithm 1 details the offline training phase.

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

Setting the number of nearest neighbours to k , we match the received physical fingerprint $R = TDoA_1, \dots, TDoA_n$ to the saved grid fingerprint F based on their Euclidean distance

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

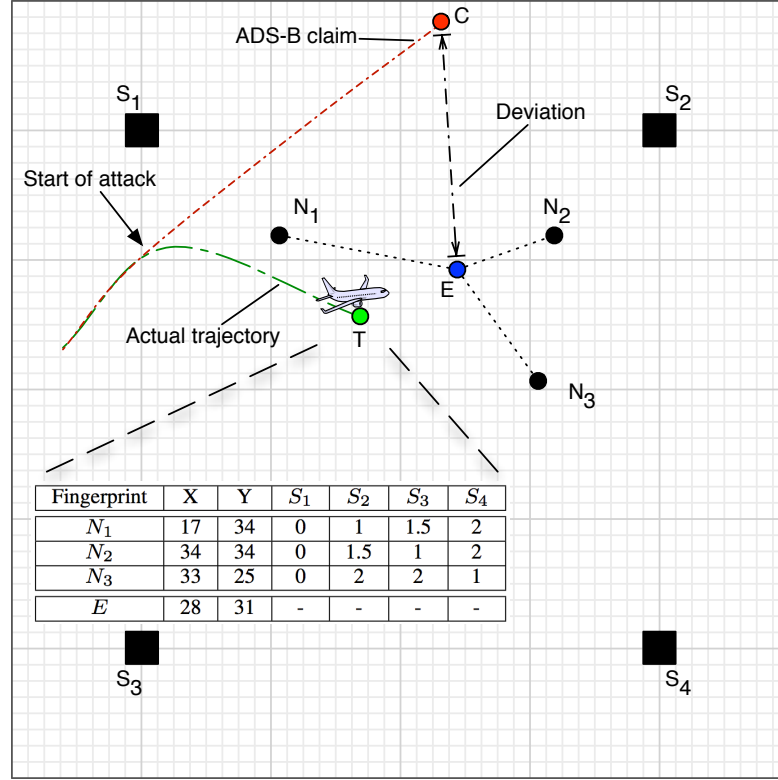


Figure 7.14: Location estimation with 3-NN in an adversarial setting where actual and claimed trajectory diverge. Using TDoA data from 4 sensors S_1, \dots, S_4 the 3 nearest neighbours N_1, N_2, N_3 found in the lookup table are averaged to obtain the location estimate E . If the deviation between E and the ADS-B claim C exceeds a threshold, an alarm is sent.

It is intuitive that in the spatial domain of our grid there are multiple neighbours that are approximately the same distance from our point of interest, hence k is an important parameter influencing the accuracy. If $k > 1$, the positions of all k neighbours are averaged by taking the mean of the longitude and the latitude. This constitutes the final estimate of the aircraft position, which is closer to the true location than any single neighbour (see Fig. 7.14 for an illustration).

7.5.4 Experimental Design

Data Collection and Hardware

As ADS-B has been in the roll-out phase for years, we can use real-world data to estimate the propagation characteristics of ADS-B messages. We do not make any assumptions on hardware features such as sending power or antennas as there

are many configurations found in different aircraft.

For our evaluation, we rely on real-world ADS-B data which we obtained from the OpenSky project. For the present analysis, we use a dataset that spans the period between 26 June 2013 and 25 June 2014. This dataset contains 53,551,672 ADS-B messages received from SBS-3 sensors manufactured by Kinetic Avionics. Besides the message content, they provide a timestamp of the message reception. From this data, we use 5 sensors that are closely located together to be able to calculate their TDoA data. The timestamps have a clock resolution of 50 ns. All sensors have omnidirectional antennas and can receive signals from a distance of up to 400 km.

Synchronization

As our low-cost SBS-3 sensors do not provide built-in synchronization (e.g., via GPS), we synchronize our data a posteriori with the help of positional ADS-B messages sent by aircraft. By using the positional information in those messages and approximating their respective propagation time, we can recover the timing offset between our ground station sensors and achieve global synchronization. We also take into account the drift of the internal clocks to improve the results. Overall, this approach enables us to achieve synchronization that is low-cost and works well with minimal requirements. More accurate and efficient synchronization using GPS could help to further improve on the accuracy of our results. However, the increased security of GNSS-free synchronization is another major advantage besides cost savings. It is obvious that in the attacker model with full access to the wireless channel, GPS-spoofing or jamming⁵ are further tools available to the attacker besides the mere injection of ADS-B messages and hence the use of GPS does not necessarily improve the overall security of the system.

Grid Design

We construct a 2D grid over a typical flight altitude of 38,000 ft (ca. 11,582 m) with a size of 2 degrees longitude and 2 degrees latitude which, due to the Earth's spherical geometry, translates to an area of ca. $150 \text{ km} \cdot 220 \text{ km} = 33,000 \text{ km}^2$. We

⁵A practical real-world threat, see e.g. [182].

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

Table 7.8: Averaged horizontal distances from the four attackers' positions to their claimed aircraft positions during the time that flight data is injected.

obtain evenly-spaced approximate squares where the number of squares (or the squares' size) is a trade-off between performance and accuracy as elaborated in the evaluation section. Of course, computation time and accuracy also depend on the size of the surveillance area. $33,000 \text{ km}^2$ are representative for wide area ATC surveillance, covering aircraft's en-route flight phase at cruising altitude.

Test Data

We use real-world flight data to test our scheme. Taking 10,443 legitimate flights with more than 100 collected messages each, we show that they are accurately verified by our system. Furthermore, we use data from various simulated attackers (due to ethical reasons, we do not implement real-world attacks) on the ground and in the air and check whether they will be verified or not. Table 7.8 shows the average simulated positions for all four attackers as described in Section 7.5.1. Using an omnidirectional antenna, each attacker injects 200 messages with the legitimate coordinates of a real flight from our sample and follows specific location patterns:

- **Attacker 1** has a fixed random horizontal position on the grid with an altitude between 0 and 500 m from which all 200 messages are sent.
- **Attacker 2** is defined by a random start position similar to attacker 1 and a random horizontal direction, moving on the ground with a speed of 50 km/h.
- **Attacker 3** has a random start position, a random altitude between 0 and 1100 m and a random horizontal direction, moving with a speed of 200 km/h.

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

Table 7.9: Results of the location verification approach dependent on number of received messages and number of sensors. The values signify the percentage of flights that have been classified as attackers.

- **Attacker 4**'s starting position is the same as the real aircraft but diverts horizontally at a random angle between 10 and 45 degrees (at cruising altitude), making attacker 4 the most difficult to detect.

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

7.5.5 Evaluation

In this section, we use the collected flight data to verify our approaches. Furthermore, we inject data from four different attackers to test the systems' resilience against intruders.

Location Verification

Table 7.9 shows the results of testing our location verification method. As we can observe, it is able to detect all attackers successfully, while minimizing false positives. For all legitimate flights, the null hypothesis is accepted when at least 30

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

Table 7.10: Horizontal errors in different grid square sizes using k-NN vs. MLAT, with 5 sensors and $k = 5$. k-NN shows a better mean accuracy than MLAT of up to 41% in our data set.

samples are collected. Indeed, only with a sample size of 10 are any false positives recorded, regardless of the number of sensors used.

For attackers 1-3, which are all relatively far away from their claimed distances (i.e., on the ground or in low airspace), \mathcal{H}_0 is typically rejected after collecting 10 or more message samples. False negatives stay in the low single digits even with TDoAs gathered by only two sensors but approach zero for 3 or more sensors. For the most powerful attacker 4, who is acting very similar to the injected position claims, a sample size of 100 is needed to detect the largest part of the injected flights, missing 4.4% of the attacks with 5 sensors.

These results offer quick attack detection: In a non-lossy environment, we can collect 100 messages in under 20 seconds. Assuming even 50% message loss, we are alerted within 40 seconds after the aircraft has diverted from its claimed course.

Location Estimation

To ensure a baseline for the accuracy of the location estimation method, we compare it with the GPS-based ADS-B position claims of legitimate flight data. We use a part of the whole data set comprising over 100,000 positional ADS-B messages where every message has been seen by 5 sensors, providing us with the necessary TDoA measurements. All location claims are on the pre-defined surveillance grid in terms

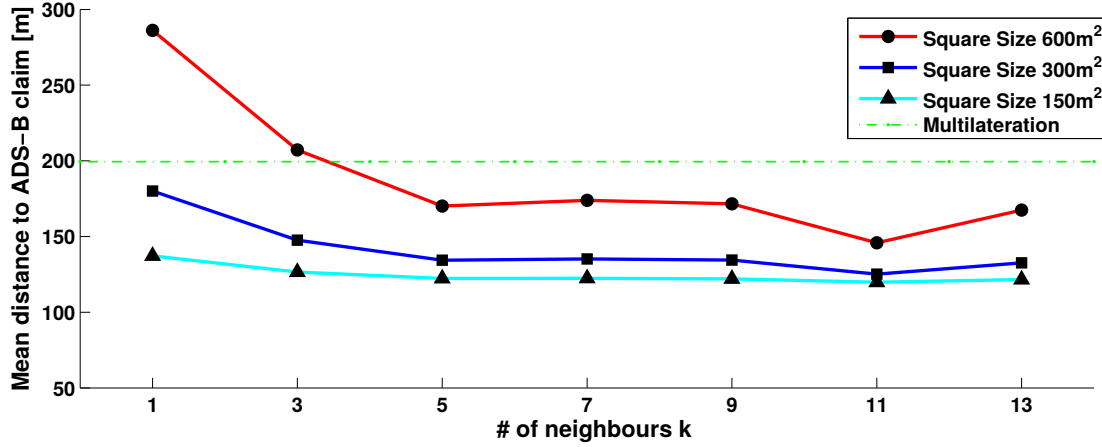


Figure 7.15: Optimal choice of neighbours k for different square sizes (MLAT as comparison).

of latitude and longitude, while their mean altitude is 11,148.8 m ($\sigma = 687.59$ m). Table 7.10 shows the location estimation quality using k -NN with squares of five different sizes over an area of 33,000 km² with $k = 5$.

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

Concerning the optimal choice of k , Fig. 7.15 illustrates the gains in accuracy when averaging a higher number of neighbours. We can see a large improvement until $k = 5$ especially for a grid square size of 600 m². Further decreases in mean accuracy are small and much less pronounced with smaller square sizes.

We also compared k -NN with a linearised MLAT algorithm using the same TDoA measurements from 5 sensors. The results show that with a 600 m² grid size, k -NN does 14.2% better than MLAT on mean errors, increasing to 41% for a 50 m² grid size. Overall, we find that k -NN does better than MLAT on noisy TDoA measurements such as those we experienced in our real-world data. Especially the more outlier-sensitive metrics RMSE and mean improve with k -NN while MLAT

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

Table 7.11: Average horizontal errors using k-NN ($k = 5$) with 150 m square size and different amounts of receivers. MLAT (5 sensors) is provided as comparison.

generally shows good median results. Since k-NN does not suffer from dilution of precision, this is to be expected as the mean GDOP in our dataset is 24.35 ($\sigma = 8.06$). Taking only “good” values below 10 into account, MLAT’s metrics are bound to improve vastly. However, doing this also decreases the number of usable messages by over 90%, reinforcing the fact that k-NN is useful in a much larger area.

The computational time is the trade-off for k-NN’s accuracy and robustness. Only with the largest square size of 600 m² is it comparable to MLAT. However, depending on the density of the airspace and the available equipment, even larger grids and longer computation times would not pose a problem in real-world settings.⁶ In scenarios where location estimation is run mainly to verify suspicious aircraft claims, it is entirely irrelevant as the examined amount of data is very small.

For our security analysis, it is furthermore important to compare the impact of sensor numbers on location estimation. Table 7.11 shows the results for the same dataset and a 150 m² grid size, if only a subset of the five sensors receives the messages. After analysing all possible subsets and averaging the results, we conclude that with only three sensors sufficient horizontal accuracy can be achieved.

Robustness One of the main advantages of k-NN over MLAT is the fact that it is much more robust to noise. To analyse the impact of noisy signal measurements, we conducted simulations testing various noise levels against the two algorithms, using a scenario similar to the one with which we collected our real-world data using OpenSky. We randomly distribute 5 sensors on a 100x100 km grid, at a height

⁶The complexity of the MLAT algorithm is constant for a fixed number of sensors, while k-NN depends on the number of squares, i.e., both the size of the monitored area and the desired accuracy.

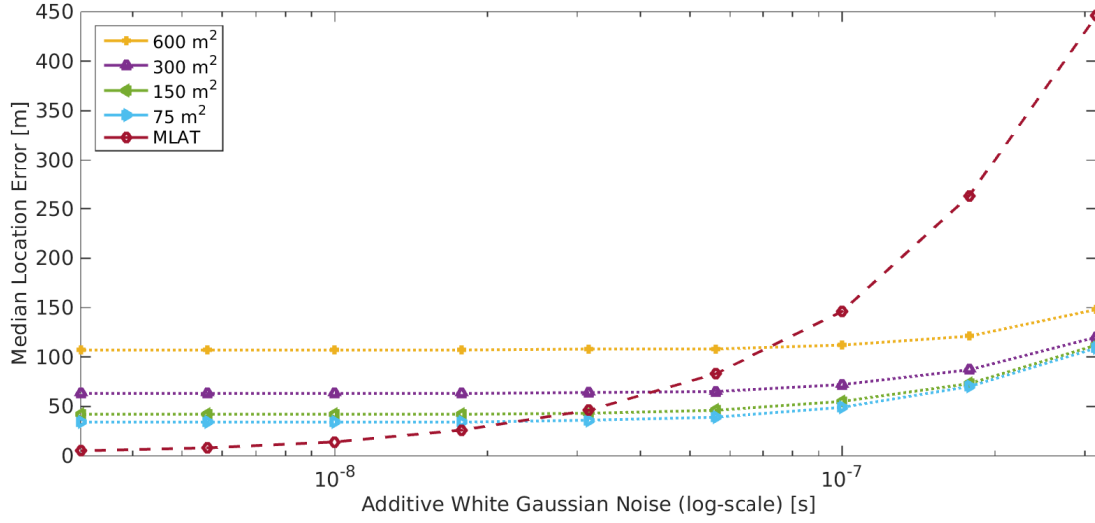


Figure 7.16: Median location error depending on the level of noise N affecting the measurements. k-NN results obtained with $k = 5$ and 4 different square sizes.

randomly distributed between 0 and 1,000 m. We create 100,000 aircraft signals sent from the grid at a height between 10,000m and 11,000m, calculate their ideal time of arrival at the 5 sensors and add white Gaussian noise $Z_i \sim \mathcal{N}(0, N)$ for each time i , where N is the variance given in seconds. We repeat this simulation 1,000 times at each variance level N to smooth out effects of the sensor placement.

From the results shown in Fig. 7.16, we can conclude that the MLAT algorithm performs very well for no or very little noise. This is not surprising for an exact method. However, the algorithm is much less robust against increased noise levels compared to k-NN. As the noise approaches a variance of $N = 10^{-7.5}s$, we can see that the median location error of MLAT surpasses the one of k-NN with $k = 5$ and $75m^2/150m^2$ square sizes. This means that at a noise level of over $31.6ns$, our k-NN approach provides superior results.

Fig. 7.17 gives further insight into the underlying reasons for the performance difference between both algorithms. It shows the median location error depending on the GDOP of the measured signal. At a noise level of $N = 10^{-7}s$, we find that MLAT only performs well when GDOP is extremely low, while k-NN is unaffected by this problem. Comparing MLAT to k-NN with various grid densities, we can see that only signals with $GDOP < 5$ can provide a low average localization error; higher GDOP

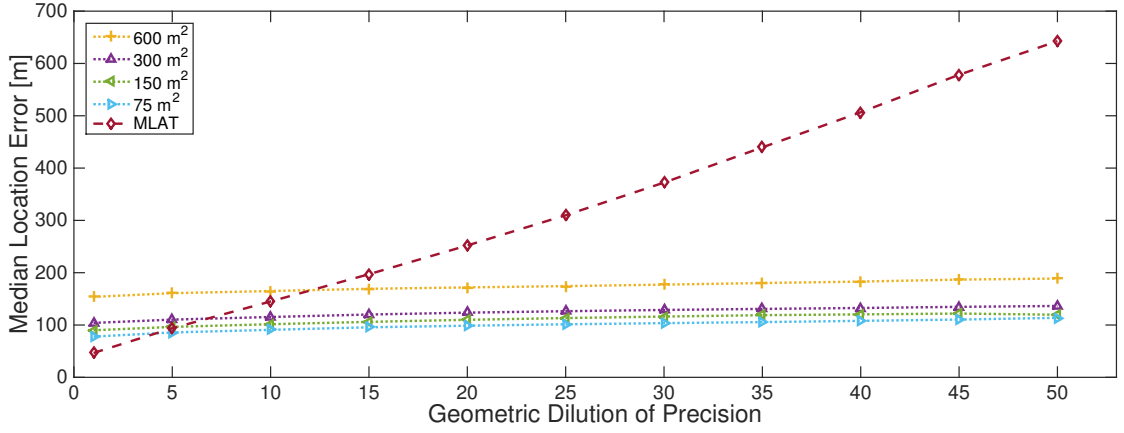


Figure 7.17: Median location error depending on the geometric dilution of precision affecting the measurements at a noise level of $N = 10^{-7}s$. k-NN results obtained with $k = 5$ and 4 different square sizes.

levels are quickly outperformed by all square sizes. Unfortunately, such a low dilution of precision is present in only a very small fraction of the potential surveillance area.

These findings illustrate that the MLAT algorithm requires extremely tight and costly synchronization and is still severely inhibited by the geometry of its receiver’s locations. In contrast, k-NN can provide effective localization with good quality even with low-cost hardware such as Kinetic Avionics SBS-3 boxes and arbitrary receiver placement. On a further note, even where the required tight synchronization is in place, it can potentially malfunction due to technical failures or adversarial influences (e.g., an attack on the GPS signals used for synchronization).

Attacker Detection We analyse the results of our attacker models who inject false ADS-B data from a different location. Through experimental analysis of the legitimate OpenSky data, we first obtain a threshold that has not been exceeded by any legitimate flight in our data set.

We found that our system should flag a given flight as illegitimate when the average deviation between ADS-B claim and k-NN estimate exceeds 1,000 m over a period of 12 messages received by 3 sensors. With this setting we encountered zero false positives in our test data, yet detect all false-data injections by attackers 1-3 within these 12 messages as their location far exceeds the threshold. Attacker 4, who starts out from the correct position, is still detected in fewer than 38 messages

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

Table 7.12: Left: Mean distances between estimates and claimed location injected by an attacker. Right: Mean distances to actual horizontal location of an attacker. k-NN ($k = 5$) with 150 m square size. k-NN accurately detects the distances between the attacker and the claim and gives a good guess about the real origin of the signal. MLAT also detects the deviations can only provide an accurate position of the aircraft-based attacker.

on average, i.e. after about 20 seconds without loss or 40 seconds assuming 50% loss. Naturally, the precise thresholds depend on equipment and scenario and should be fitted accordingly.

Besides detection of false claims, location estimation can provide a guess of the attacker's current location. Table 7.12 provides the results of this estimation for all four attacker types. Our horizontal estimate for the origin of message signals fits within approximately 2,000 m of the real location for the ground-/low airspace-based attackers. For the high-altitude attacker 4, we obtain an estimate accurate within the typical error range for legitimate flights of less than 200 m.

Table 7.12 also illustrates a major drawback of MLAT in the same scenario. While it is feasible (though costly) to build a multilateration system with good accuracy for larger surveillance areas in the sky, it is difficult to provide the same level of accuracy on the ground. Hence, while MLAT offers a similar estimate quality for the aircraft attacker 4 in our setup, is not able to provide any helpful guess for the location of all ground/low airspace attackers.

7.5.6 Discussion

We proposed and evaluated two methods of location verification. The first one, statistical and based on collected time differences of arrival between as little as two ADS-B sensors, allows us to quickly detect injected data with high certainty. Using only low-cost ADS-B sensors, we find that it outperforms MLAT in terms of range and detection speed, increasing coverage by a factor of more than 100.

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

MLAT is a popular technology in independent ATC localisation as it is difficult to attack in typical settings and TDoAs are readily available without hardware or software changes. While not perfect, it is currently considered the best security option until the community works out fundamental long-term solutions for authentication in ATC communication networks. In the meantime, it is crucial to work on increasing current security to protect air traffic against potentially devastating adversaries. We have shown that our approaches can vastly improve range, detection speed and accuracy of air traffic surveillance in real-world environments.

7.6 Comparison of Transparent ATC Security

Lastly, we want to provide a concise overview of the advantages and disadvantages of the transparent security measures discussed in this chapter: data link fingerprinting, physical layer intrusion detection, and lightweight aircraft location verification (divided further into location verification and location estimation). We identify the nine dimensions relevant for our comparison:

- **Efficacy against threat actors:** The first aspect of comparison is the security efficacy of the proposed solution. We are interested in what level of attackers can be detected and what kind of sophistication is needed to circumvent it (as defined in Section 2.3).
- **Sensitivity:** Secondly, we compare the effectiveness against ghost aircraft injections, specifically, how accurate / how high the rate of detection of such attacks is.

- **Specificity:** In contrast, reducing false positives, i.e., the flagging of non-existing attacks is crucial for any real-world system. If specificity is too low, the practicality of an approach will be greatly reduced.
- **Cost:** As cost remains a crucial factor in civil aviation, it is important to discuss how expensive the deployment of a solution is to an ANSP.
- **Ease of integration:** Another important aspect concerning the real-world applicability of a solution is the complexity of integration into current ATC technology and processes. None of the solutions require technological changes but their ease of deployment still varies.
- **Minimum number of sensors required:** Related to the cost and ease of integration is the question of how many receiving sensors are required to execute a given countermeasure.
- **Speed of detection:** Another crucial dimension is the time it takes until a confident decision on a given aircraft can be made, maximizing sensitivity and specificity in the process but still providing quick detection of potential attacks.
- **Training phase:** This aspect indicates whether the proposed countermeasure requires a training phase and whether that training requires real aircraft data obtained at the intended location of use or can be conducted offline without such data.
- **Attacker localisation:** Apart from the detection of an attack and the verification of real aircraft, it is beneficial to narrow down the location from which a potential attack has been conducted.

Table 7.13 provides a high-level overview of all aspects. We can see that the fingerprinting and RSS-based physical layer solutions are simple and cheap to implement, as they only require a single sensor and are easily integrated with current SSR systems, requiring only the collection of already existing SSR protocol

	Fingerprinting	Physical Layer RSS	Location Verification	Location Estimation
Threat Actors	Up to Hobbyist	Up to Cyber Crime / Terrorism	Up to Nation States	Up to Nation States
Sensitivity	Low	High	High	Very High
Specificity	High	High	High	Very High
Cost	Very Low	Very Low	Low	Medium
Integration	Easy	Easy	Very Easy	Very Easy
Sensors	1	1	2	3
Detection Speed	100s	40s	10-40s	6-20s
Training	Yes	Yes	Yes	Offline
Attacker Localisation	No	No	No	Possible

Table 7.13: Comparison of the proposed transparent ATC security approaches.

and meta data. They provide good protection against less sophisticated attackers and those who are not aware of these defence mechanisms being employed. They differ mostly in their sensitivity as fingerprinting is likely to miss replay attackers, whereas only relatively powerful attackers with the ability to control their sending strength and good knowledge of the sensors' locations can fool the RSS-based detection mechanisms. On the other hand, regarding specificity, both approaches have been extremely unlikely (less than 0.1%) to show false positives, i.e. detecting attacks where there are none. Both require a training phase based on real aircraft data from the deployment location and do not provide the possibility to locate potential attackers. Their detection speed is practical with about 100 seconds for fingerprinting and 40 seconds for the RSS-based approach.

On the two TDoA-based security approaches to ATC communication, we can conclude that they are highly effective with regards to threat actors, up to the level of nation states. An attacker (similar to multilateration) requires highly sophisticated means including perfect synchronization, directional antennas, and knowledge of the sensors' locations to defeat this physical security primitive. Even assuming such power, extensive deployment of sensors provide can provide a level of security that is difficult to overcome. As laid out in Section 7.5, location verification

requires fewer sensors to work than its estimation counterpart which is reflected in their cost. Both approaches can be easily integrated into existing SSR systems, especially where multilateration is considered or already deployed. Exploiting the data from several sensors, aircraft location estimation using k-NN provides the highest sensitivity and specificity in our comparison with no false positives or missed attacks and a fast detection speed of generally less than 20 seconds even for sophisticated attackers. Last but not least, it enables the user to locate an attack, which can be helpful in its termination.

Overall, we can conclude that there are several possibilities to create a transparent attack detection system for air traffic communication which are cost-efficient and effective. Due to their low cost and easy integrability, there is a good case to be made for the combination of several approaches. The fusion of several systems is a long-standing practice in ATC and in the present case, too, it offers a number of advantages, with the most important being the improvement of both sensitivity and specificity. Furthermore, it increases the reliability of the system and also the range and timeliness of attack detection, particularly in areas where few sensors can be deployed.

*Nam et secundas res splendiores facit amicitia et
adversas partiens communicansque leviores.*

*For friendship makes prosperity more shining and
lessens adversity by dividing and sharing it.*

— Cicero's *Laelius De Amicitia*

8

New Privacy Challenges in Wireless Air Traffic Communication

Contents

8.1	Introduction	141
8.2	Data Sources for Aircraft Information	143
8.3	Passive Tracking of Aircraft	144
8.3.1	Tracking using Commercial Web Services	144
8.3.2	Tracking using OpenSky	145
8.3.3	Discussion	148
8.4	Information Collection using the Data Link Layer . . .	150
8.4.1	Transponder Identification	150
8.4.2	Privacy Implications	151
8.4.3	Mitigation	153
8.5	Large-scale Event Detection	154
8.5.1	Experimental Design	154
8.5.2	Feature Selection	155
8.5.3	Evaluation	156
8.5.4	Discussion	159
8.6	Summary	160

8.1 Introduction

Open systems, whether deliberately open or not, and insecure wireless technologies enable concerns that go beyond the discussed active attacks on critical aviation

infrastructures. The combination of legacy infrastructure, disruption through SDRs, and the proliferation of public and private sensor networks, which capture air traffic communication around the globe, opens up several new challenges to the privacy of aviation users.

In previous chapters, we have analysed the impact of active attacks on ATC communication and proposed to employ transparent countermeasures powered by cheap sensor networks. However, the availability of these networks changes not only the security landscape but at the same time introduces severe consequences for the privacy of the people who use aircraft as their means of transport. As powerful actors increasingly lose their informational edge over ordinary citizens, privacy challenges are increasing as aircraft information becomes available widely and easily on the Internet.

In this chapter, we highlight some of the consequences of this shift of technological advantage that aviation used to enjoy. We postulate that while cheap air traffic communication networks will play a key role in future security and surveillance solutions, they have already created many challenges to aviation privacy which will only be exacerbated in the future. The aviation community needs to take into account these developments and address them if they want to maintain the privacy of their users.

Concretely, we discuss the changing nature of aircraft tracking by purely passive observers exploiting the openness of ATC protocols. We examine both filtered commercial web services and unfiltered networks such as OpenSky or similar private receiver installations. Concretely, we analyse the number, type, and ownership of all aircraft seen in OpenSky's receiving range. With this information, we compare how many presumably sensitive aircraft are blocked from public trackers and thus illustrate the futility of current attempts to maintain traditional privacy for aircraft owners.

We further use the same data for a proof-of-concept in the detection of large-scale events in the real world. We postulate that it is possible to use aircraft meta data to

detect unusual events happening in the coverage area of a sensor network and validate this theory using the well known World Economic Forum held in Davos each year.

8.2 Data Sources for Aircraft Information

There are various public data sources available that provide meta information on aircraft based on their identifiers (ICAO 24 bit identifier or callsign). This information typically includes the aircraft type (e.g., Airbus A320) and the owner/operator (e.g., British Airways or a governmental agency), which can be exploited for further in-depth analysis.

These public databases can broadly be divided into online and offline sources:

- **Online sources** are websites with query possibilities returning meta information and, in the case of flight tracking websites, the movement of aircraft. Here, the not-for-profit project <http://airframes.org> is the most valuable source for learning general aircraft information as it offers the most comprehensive data, including background knowledge such as pictures and historical ownership information. Secondly, we use the commercial project <http://flightradar24.com> (short: FR24) to look up aircraft metadata and flight tracks, i.e. historical movement data. FR24 states it has more than 500,000 aircraft in their database (as of April 2016).
- **Offline sources** are aircraft databases, typically in SQLite or CSV format provided by third parties. We use two of these sources for our large-scale analysis of metadata. The first database is available and constantly updated in the Planeplotter software (<http://www.coaa.co.uk/planeplotter.htm>). Our version of the SQLite database file `database.sqb` created in November 2014, contained 120,149 rows of aircraft data. The second database is available from Junzi Sun at TU Delft (<http://junzisun.com/adb/>), who has been collecting information on all visible aircraft from FR24 over a period of 8 months at the time of writing (April 2016), amounting to 89,391 rows.

Note that these sources are naturally not complete, since they rely on compiling many separate smaller databases. During our research we found that, for example, most non-commercial aircraft with Spanish registrations were entirely missing from all sources.

8.3 Passive Tracking of Aircraft

We define aircraft tracking as the act of obtaining positional information on aircraft, live or delayed, by private and actors outside of military or aviation circles. We further consider purely passive observers as outlined in Section 2.3. Motives for aircraft tracking are manifold, ranging from traditional hobbyist planespotters, over military and business interests, to criminal intent. Where traditionally the private tracking of aircraft has relied on visual means, i.e., seeing — and recognizing — the aircraft in the sky or at an airport, the technical advancements discussed in Chapter 2 have made accurate, fast and large-scale tracking of all transponder-equipped aircraft feasible and widespread.

For a private actor, there are two options to track aircraft: commercial web services or installing personal ATC sensors. While a single receiver can already provide tremendously interesting results, these are increased manifold with a larger network. As this type of network is not out of the scope of any interested actor any more, we use OpenSky to illustrate the power of such installations by comparing its view of the airspace with that of commercially available providers.

8.3.1 Tracking using Commercial Web Services

Commercial web services, which exploit the non-confidential nature of air traffic control protocols such as ADS-B or Mode S for aircraft tracking, have been available to the public for more than 10 years and have attracted considerable interest from traditional and many newly gained planespotters. This development has been met with suspicion by many aviation users, who value their privacy when using air transport, including commercial, military, and government actors.

Traditional ‘offline’-inspired solutions to maintaining privacy and secrecy for aircraft, such as the Aircraft Situation Display to Industry (ASDI) scheme,¹ aim to prevent the public display of aircraft movements by working with commercial web services. The ASDI register allows aircraft owners with security concerns to restrict the tracking of their aircraft [183]. Through ASDI, government and court orders, and also directly by aircraft operators, complying web trackers are provided with lists of aircraft identifiers, which they in turn do not publicly display. These privacy measures range from obscuring live flight information over blocking the display of historical aircraft movements, to fully denying any knowledge about the existence of an aircraft.

However, we argue that such approaches have become long obsolete in the SDR era. On the one hand, these lists are neither complete nor implementable without failure. On the other hand, it has become practical, even trivial, to get an unfiltered and accurate live picture of the airspace for any interested party. While some of the largest online flight tracking services such as FlightRadar24 comply with requests to not display private or sensitive aircraft data, there are many unregulated sources available, some of which even actively identify interesting aircraft to the public (for example <http://www.adsbexchange.com>; see also [22]).

8.3.2 Tracking using OpenSky

As discussed, commercial flight tracking websites provide live information on a large number aircraft. However, many aircraft do not get tracked — or at least displayed — by these websites, despite flying regularly.

The OpenSky Network is one option open to researchers providing unfiltered data. In this section, we show what it is capable of in terms of aircraft tracking. Again, it cannot be stressed enough that OpenSky’s basic capabilities for tracking aircraft are open to anyone with little investment of resources, proving the futility of traditional privacy solutions through blocking schemes.

¹Information on blocking requests is available at <https://www.nbaa.org/ops/security/asdi/>.

	Aircraft	[%]
Total	37,360	100
Identifiable, of which	28,657	76.7
- Fully trackable	17,905	47.9
- No tracking history, of which	10,414	27.9
- - No aircraft information	3,653	9.8
- Other (Ground)	338	0.9
Unidentifiable	8,703	23.3

Table 8.1: Breakdown of aircraft seen by OpenSky and their privacy restrictions on public tracking websites.

Blocked Aircraft

In order to understand the extent of such blocks on aircraft, we compare OpenSky’s data with that of a major commercial flight tracker. We are able to identify three different blocking categories, which are included in Table 8.1. The least stringent block is ‘no tracking history’ in which an aircraft appears on the tracking website with information such as model or operator but has no flights recorded. Next is ‘no information’ whereby the aircraft has no ancillary information as well as no tracking history, with existence being confirmed through other data sources. Finally, some aircraft were fully ‘unidentifiable’ in any available public sources. This could suggest an even stronger block but we speculate that the largest part is more likely due to the available sources not being complete. Hence, we concentrate our analysis on the more than three quarter of all aircraft for which data was available.

The data indicates that about 77% of all aircraft are identifiable via public data sources and 48% of all aircraft seen via ADS-B are tracked by the commercial tracker. Consequently, 52% have no recorded flights on the tracker, 19.7% of which still have information listed, suggesting that only their tracking history is actively blocked. An additional 8.9% have their meta information blocked on complying tracking websites but are identifiable through other sources, while about 23% could not be identified through any available means.

Table 8.2 breaks down the aircraft types along three considered levels of tracking capabilities: full tracking, no recorded flight history, and completely blocked (i.e., no meta data or indication that the aircraft is known is available). We determine

Aircraft affiliation	Number of aircraft (% total aircraft)		
	with flight history	without flight history	fully blocked
Private/business	1,473 (8.2%)	5,863 (56.3%)	2,237 (59.7%)
Military	148 (0.8%)	2,125 (20.4%)	799 (21.3%)
State-related	14 (0.1%)	35 (0.3%)	30 (0.8%)
Scheduled/Other	16,270 (90.9%)	2,486 (23.9%)	586 (15.6%)

Table 8.2: Breakdown of identifiable types of aircraft with and without flight history, and with full blocks on tracking websites used for comparison.

the type using the registrar of the aircraft. For example, if the registrar is a state actor such as government or police, we classify it as ‘state-related’. Note that the 8,703 unidentified aircraft from Table 8.1 are excluded from this table as we naturally cannot categorise them.

There are two major interest groups which seek out blocks from internet databases: state actors such as air forces and government aircraft, and aircraft privately owned by companies, institutions, or individuals.

Military and State Aircraft Aircraft in use by the military and governments have a stronger privacy requirement. Although the movements of the most prolific members of government are typically highly publicized, the exact routes and timings may need to be kept secret. Furthermore, many sensitive diplomatic missions of all types have no interest in becoming publicly known at the time (or at all). Whilst cooperation by tracking services can make it harder to track these aircraft, our results show that any listener with an SDR and publicly available metadata can still learn of their presence. Table 8.2 shows that the vast majority of the military aircraft seen by us show no flight history (2,125 compared to 148 with history).

It has to be noted that, while isolated ADS-B-equipped military aircraft are regularly spotted on tracking websites (as indicated by the 148 visible aircraft in our data), the vast majority of military aircraft are either not equipped with ADS-B or have the ability to turn it off [184]. Thus, by using other protocols such as Mode S or ACARS, the visibility of such sensitive flights can be further improved by any interested party.

State aircraft were also observed in our data, including those belonging to state leaders or monarchs.² In the majority of cases, these aircrafts' histories were not displayed by the tracking website although the discrepancy is not as large as for military aircraft (14 visible aircraft compared to 35 blocked aircraft). It is notable that most of the blocked aircraft do not display any metadata, while this is less the case for military aircraft.

Private/Business Aircraft Some users of non-scheduled aviation such as business jets or private pilots seek to protect for example their business activities or do not like to broadcast their movements to the world in general. On the former reason, for example, there have been allegations of such movements being used as information for stock trading [185]. Consequently, Table 8.2 shows that the vast majority of identified business or private jets do not have recorded flight history on tracking websites (5,863 vs. 1,473 that are trackable), implying their use of the ASDI programme or similar direct-to-website notices.

They, too, are of course easily tracked with cheap SDR hard- and software, making these blocks a mere inconvenience for the interested passive observer. Lastly, all current privacy preserving approaches are further undermined by an aircraft's use of other air traffic communication protocols such as ACARS (as shown by Smith et al. [90]) or Mode S, which also give away the presence, identity and track of any aircraft that uses them.

8.3.3 Discussion

In this section, we showed that traditional schemes to prevent the widespread and easy tracking of sensitive aircraft have become all but obsolete. On the contrary, much sensitive information is already widely available directly from leaky commercial trackers and aggregated by social media feeds and dedicated websites. We argue that such democratization of information has led to a partial erosion of power for state actors, as airborne missions become known immediately to even passive observers.

²For example, we noted one of the Air Force One aircraft multiple times, this has also been previously discussed by Costin et al. [12].

Today, plane-spotters around the world detect high-level aviation anomalies, potential incidents, and ‘interesting aircraft’ practically immediately and on a large scale, a capability previously limited to state actors. Social media accounts tracking emergency broadcasts provide instant news coverage for both the press and interested individuals, much to the chagrin of some in the traditionally closed aviation community. Hijacked aeroplanes, too, are detected easily by individuals at home and shared in real-time over Twitter while the aircraft is still in the air [22].

An exemplary case study is provided by the intelligence and security services. With increasing automation and availability of online aviation feeds, the development has gone from occasional sightings of aircraft operated by domestic security services to the large-scale and immediate detection of all transponder-equipped aircraft. An example of the implications of this technological shift is the recent uncovering of a large number of surveillance aircraft employed through front companies of the FBI, an operation that had previously gone unnoticed for some decades [22].

In military settings, this type of open surveillance using data gleaned from Mode S and ADS-B broadcasts has led to similar information leakage through the – intentional or unintentional – use of transponders during active missions. The diligent tracking of recent airborne engagements in Syria by NATO and Russian aircraft illustrate this point. As airstrikes and reconnaissance missions can easily be detected and anticipated, potentially sensitive strategic and operational information is broadcasted, and deniability of airborne actions becomes difficult, the impact of insecure civil aviation protocols on military users is growing [22].

The same problems that are causing concern for current manned surveillance aircraft apply to UAVs. As aviation authorities are expected to maintain a similar standard of rules for drones in the civil airspace, the mandatory use of ADS-B transponders will retain the broadcasts of sensitive data to anyone listening. As a concrete occurrence, Swiss border patrol drones are forced to operate with enabled ADS-B broadcasts on surveillance missions. Their targets, such as human traffickers and smugglers, can track the position of the drones on tracking services

using their smartphones and easily avoid discovery by moving only when the drones are not operating [22].

These examples illustrate the impact that merely passive threat agents have currently, besides the active attacks on wireless air traffic communication protocols discussed in the previous sections. As the information leaks described here even stem from web services that actually *comply* with existing privacy regulations, it is safe to assume that the further spread of SDR hardware and OpenSky-like capabilities will eventually render the privacy of sensitive aircraft and their location history non-existent in practice.

8.4 Information Collection using the Data Link Layer

Using fingerprinting techniques such as those described in Section 7.3, we can also collect information on aircraft without accessing any of the content of their communication. We provide a proof-of-concept in this section and analyse the different types of transponders used by aircraft seen on OpenSky.

8.4.1 Transponder Identification

ADS-B-capable transponders form part of the avionics that the plane was manufactured with or that were later retrofitted by the fleet operator. There is a wide variety in transponders manufactured and installed around the world, depending on business and regulatory environments. For new aircraft, where the transponder came with the installed avionics at delivery, the purchaser selects a whole avionics suite from the available options given by the manufacturer of the aircraft. When a transponder is retrofitted later, the options are much broader and a suitable transponder can be chosen from any that are certified for a) the operator's home nation and b) for the airframe in question. To look deeper into the matter of differences in ADS-B transponders, we used our public aircraft sources as ground truth on aircraft types and operating airlines.

Type	Manufacturers	Examples
Type 1a	Rockwell Collins	TPR 901
Type 1b	Honeywell	TRA-67A
Type 2	Honeywell	TRA-100
Type 3	Rockwell Collins, ACSS	GLU-920/925, XS-950
Type 4	Honeywell	Embraer SBAS, A380
Type 5	Garmin	G3000, G5000
Type 6	Honeywell	ERJ-190 Primus FMS

Table 8.3: Manufacturers of the various transponder behavior classes.

We further conducted a Google research exercise to find out which of the aircraft fleet are equipped with what kind of transponder. This type of data is not necessarily easily available but can eventually be inferred from cross-referencing many articles and data sheets on avionics equipment of different fleets around the web.³ Some of the required information can also be contained in existing scientific articles such as [186]. Using these sources, we found concrete information on some of the fleets we observed in our data set and used them as exemplary representatives of the whole cluster. Interestingly, the different implementations are not exclusive to a single manufacturer as seen in Table 8.3. While we do not see a direct negative impact by our work on the security or privacy of the transponder manufacturers or users, we do not publish the mappings between aircraft fleet and transponder type and behaviour at this point.

We could not establish a link between different versions of the implementation of the ADS-B standard (DO-260, DO-260A or DO-260B), which we assumed to be one reason for the variety of the installed landscape. As the behaviour of the transmission periodicity can easily be changed with a software upgrade of the responsible air data computer, this can also confound the results.

8.4.2 Privacy Implications

Naturally, the possibility of transponder and aircraft fingerprinting has some implications for flight privacy. While the ability to fingerprint specific aircraft

³An example is provided by the Lufthansa Technik website here: <http://goo.gl/VCGr4x>.

or types of aircraft is not likely considered to be a problem for scheduled airliners, this may be different for private or business aircraft.

Because of the extensive modern flight tracking possibilities such as the ones we have discussed above, there have been many concerns especially within the general aviation community, which has been asking for an effective privacy mechanism. The UAT data link of ADS-B used by some general aviation aircraft in the US offers such a privacy mechanism. More concretely, an aircraft can generate a non-conflicting, random, temporary ID to avoid consistent tracking over time by third-party services. However, this generated ID can only be used under visual flight rules while not receiving ATC services, severely limiting its usefulness. Besides this, it has been shown that the DO-282B privacy solution has serious weaknesses, as the real ID of the aircraft and its random ID are correlated [187].

Yet, even if this oversight were to be fixed, it would not close the fingerprinting-based privacy issues discussed in this work. While we did not explicitly analyse UAT, we have no reason to believe that it does not exhibit similar fingerprintable patterns. Further extension of this work, through the use of additional features derived from both physical-layer and data-link layer characteristics, could potentially increase the granularity of the approach. If this improvement leads to a level where single aircraft could be identified with some certainty, it would enable the tracking even of pseudonymous aircraft who are not broadcasting their real ICAO number.

Indeed, preliminary analysis suggests that there are many further data link features that can be exploited. Not least due to the scattered state of Mode S and ADS-B equipage at this point of the world-wide roll-out, we find that there is a large variety in the usage of message types and fields between aircraft. To illustrate this point, we collected a sample of Mode S-equipped aircraft in the Frankfurt area over the course of 90 days. While analysing our dataset of 7,587 aircraft (which have been seen for more than 10,000 seconds and sent at least 1,000 messages), we identified 62 message fields that are used to broadcast periodical content. We found 3,375 unique combinations of these message fields, suggesting that up to 44.5% of aircraft may be uniquely identified using their message types and contents

alone. While some of these message types such as those providing data about an aircraft's capabilities and instrument settings can change over time, these numbers suggests that there may be enough potential information in the transponder data of aircraft to pose a severe privacy problem.

Business intelligence As a secondary consideration, OpenSky's dataset can provide an interesting picture of the current ADS-B transponder market. Such data is not necessarily easily available and can prove interesting for competitors or market researchers.⁴ The data can, for example, easily be broken down into segments, showing the proliferation of certain transponder types or manufacturers in different countries or regions. Alternatively, it would be possible to analyse trends over time.

8.4.3 Mitigation

Where fingerprinting is considered a challenge to privacy, there is little in terms of quick solutions that could be done to mitigate this problem currently. Although the task is daunting, creating a strategy to solve the privacy challenges is possible and should be multi-pronged. Companies would need to provide software updates to all their transponders which would need to be applied by airlines and private pilots. To make such updates effective, however, the different suppliers are required to agree on a common implementation of their ADS-B message system. Defining the standard DO-282B [58] more rigorously by the RTCA would help, although changing standards is generally a lengthy process in many technical areas.

Furthermore, it is not clear if any single implementation of randomly chosen message intervals offers a better performance from a networking perspective in the given scenario, as the interval slots chosen by aircraft are independent of each other. Yet, performance is something which should be thoroughly analysed before making a decision as the 1090 MHz channel is notoriously overloaded with message loss rates of up to 90% in crowded airspaces [24].

⁴Although there are paid options offering some of this data, e.g. the Aviation Week Intelligence Network <http://awin.aviationweek.com>.

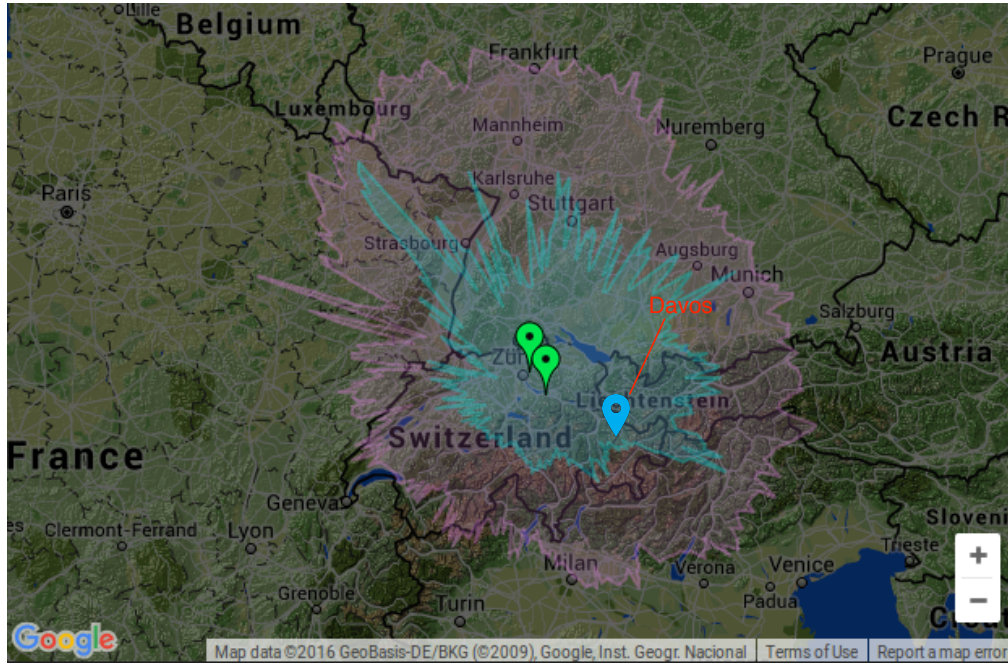


Figure 8.1: Coverage of the closest two sensors in the Davos, Switzerland area in 2016.

8.5 Large-scale Event Detection

Besides the tracking of individual aircraft, it is possible to use aircraft data, in our case from OpenSky, in an attempt to detect unusual events happening in the coverage area of a sensor network. In this section, we discuss one approach that can successfully detect large-scale events and validate it using the well known World Economic Forum (WEF) in Davos as an example. Such open source information has enjoyed popular application in many areas, including private and public intelligence services, which utilize it for purposes of open source intelligence (OSINT) [188].

8.5.1 Experimental Design

Similar to our approach concerning the previously discussed privacy-related issues, we combine OpenSky’s sensor data with the same publicly available databases. We use their information on 24-bit ICAO identifiers, aircraft types, and airlines.

To detect the WEF in Eastern Switzerland, we chose to use data from the two sensors closest to the area of interest for the first 65 days of each of the last three years. In 2014, these were the SBS-3 receivers 30783 and 30788. In 2015, only 30788

was covering the area and in 2016, we had the Radarcape receivers 80596247, and 80602915. Fig. 8.1 shows the coverage area for 2016. It is arguably an advantage to have different data sources for each year, as potential overfitting effects of our model seem unlikely, if it proves accurate regardless of the variation in the input data.

For each year (or event detection cycle), we retrieved the accumulated ICAO identifiers of all aircraft seen in the surveillance area for each separate day. By cross-referencing it with our databases, we derived distinct features from these lists, which were used to detect anomalies in the respective time series. In our case, we were searching for unusual patterns over a period of several days that would point us to the correct date of the WEF.

8.5.2 Feature Selection

We examine 7 potentially interesting features, both in their absolute values and their relative value compared to the mean of the dataset:

- **Number of all seen aircraft:** The total number of aircraft can give an indication of the overall flight activity in the surveillance area.
- **Number of distinct business aircraft:** The number of typical business-class aircraft can indicate events where many attendees use private and privately chartered aircraft (typically high-level executives).
- **Number of distinct military-related aircraft:** The number of military aircraft in the surveillance area can indicate a heightened security requirement around a local event but also increased military activity elsewhere in the globe.
- **Number of distinct government-related aircraft:** Similarly, an unusual fluctuation in the number of government-related aircraft can be an indication of significant diplomatic events or other state actor activities.

- **Number of distinct helicopters:** Helicopters as opposed to aircraft are often used for special purposes such as medical or police movements, or as fast private transport. Hence, a fluctuation in normal levels can indicate an unusual event.
- **Number of blocked aircraft:** As we have shown, blocked aircraft, i.e. aircraft which are not displayed by major online flight trackers, are often private, business, military, government-related or other aircraft, with a relevant interest in privacy. Thus, an increase in this feature in a given timespan can indicate the existence of events of interest.
- **Number of unknown aircraft:** Similarly, as a subset of the former feature, aircraft not available in any of the public databases available to us are of increased interest. This type of secrecy could be further indication of relevant events.

8.5.3 Evaluation

We now evaluate whether there are significant anomalies in the time series of our features to detect a large-scale event. Our ground truth consists of the dates of the WEF: 22-25 January 2014, 21-24 January 2015, and 20-23 January 2016.

Fig. 8.2 illustrates 3 of the 7 discussed features over a time span of 65 days from December 1, 2013 collected by two OpenSky sensors located in Switzerland,⁵ while Fig. 8.3 does the same for the blocked and business aircraft features over the respective time periods of 2015 and 2016, starting from the beginning of the year.

By analysing the time series, we can detect outliers (i.e., peaks significantly above the long-term mean) even by simple visual inspection. As shown in the area highlighted in red, the two highest peaks in terms of absolute business aircraft seen coincide with the first and the last day of the WEF in that year.

Table 8.4 shows each feature's deviation per day from the long-term median in 2016 (for the full data, see Appendix A). Quantitatively, we find a distinct increase of

⁵Unfortunately, due to technical problems with OpenSky at the time, there is no further data available for February 2014.

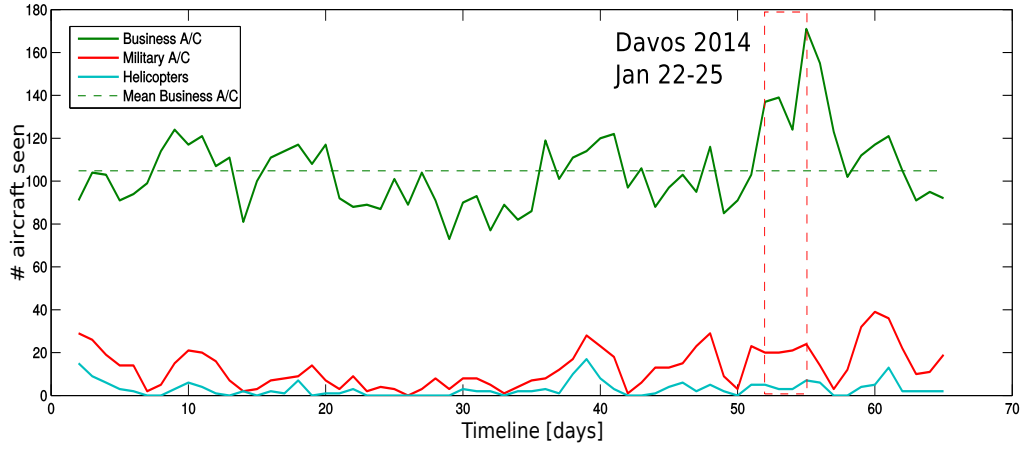


Figure 8.2: Illustration of three time-series features (absolute number of aircraft) in the Eastern Swiss surveillance area of OpenSky around the time period of Davos 2014.

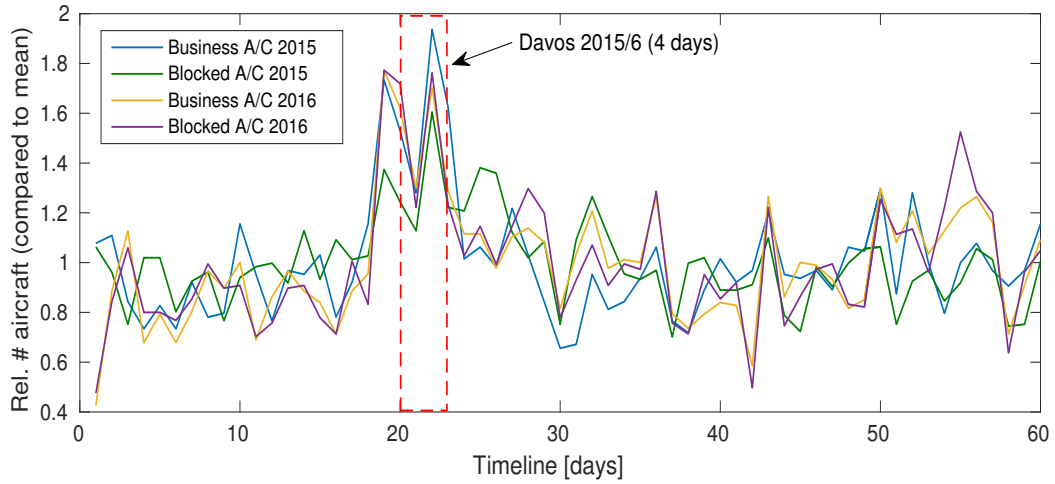


Figure 8.3: Illustration of two time-series features (relative number of aircraft) in the Eastern Swiss surveillance area of OpenSky around the time period of Davos 2015 and 2016, respectively.

up to 77% from the mean business aircraft activity on the day before the WEF 2016, which also means a 46% increase in activity over the next highest peak recorded outside of the WEF time period. Similarly, an up to 77% increase from the mean activity and 25% over the next highest peak make the appearance of blocked aircraft a very strong indicator for unusual events of the WEF type. It should be noted that these two strongly useful features have a large overlap in terms of the aircraft they represent, as business aircraft are typically blocked from public flight trackers as discussed in Section 8.3.2 (the cross-correlation is 0.95 for the 2016 sample).

Day	Relative Feature Values							Absolute Feature Values						
	Bus	Hel	Mil	Gov	All	Unk	Blo	Bus	Hel	Mil	Gov	All	Unk	Blo
1	0.42	0.80	0	1.40	0.91	0.77	0.48	36	1	0	3	1384	8	44
2	0.89	0.80	0.37	1.40	1.15	0.68	0.84	77	1	3	3	1746	7	78
3	1.13	0	0.37	1.40	1.19	0.77	1.06	98	0	3	3	1805	8	98
4	0.68	1.61	0.37	2.79	1.05	1.16	0.80	59	2	3	6	1589	12	74
5	0.79	1.61	0.74	0.47	0.96	0.87	0.80	68	2	6	1	1457	9	74
6	0.68	1.61	0.99	0.47	1	1.36	0.77	59	2	8	1	1526	14	71
7	0.81	2.41	1.48	1.40	0.99	0.87	0.85	70	3	12	3	1504	9	79
8	0.97	0	1.11	1.40	1.05	1.07	0.99	84	0	9	3	1600	11	92
9	0.90	0	0.25	1.40	1.01	0.97	0.90	78	0	2	3	1535	10	83
10	1.01	0	0	0	1.06	0.97	0.91	87	0	0	0	1608	10	84
11	0.68	0	0.74	0.93	0.96	0.58	0.70	59	0	6	2	1460	6	65
12	0.87	0	0.99	0.47	0.90	0.68	0.76	75	0	8	1	1362	7	70
13	0.96	2.41	1.48	0	0.92	0.68	0.90	83	3	12	0	1400	7	83
14	0.89	1.61	1.11	0.47	0.94	0.87	0.91	77	2	9	1	1423	9	84
15	0.83	0.80	1.36	1.40	1.02	0.58	0.78	72	1	11	3	1556	6	72
16	0.72	0.80	0.74	0	0.99	0.68	0.71	62	1	6	0	1510	7	66
17	0.89	0	0.37	0.93	1.06	1.07	1.01	77	0	3	2	1607	11	93
18	0.96	2.41	1.36	0.93	0.97	0.77	0.83	83	3	11	2	1477	8	77
19	1.77	1.61	0.74	0.47	0.88	1.07	1.77	153	2	6	1	1336	11	164
20	1.62	0.80	1.73	0.47	0.99	1.65	1.72	140	1	14	1	1500	17	159
21	1.30	2.41	1.24	0.93	0.98	0.48	1.22	112	3	10	2	1490	5	113
22	1.70	3.21	1.36	0.47	1.11	1.65	1.76	147	4	11	1	1683	17	163
23	1.27	1.61	0.74	0.93	1.02	0.87	1.23	110	2	6	2	1553	9	114
24	1.12	1.61	0.49	1.40	1.06	0.87	1.03	97	2	4	3	1610	9	95
25	1.12	0.80	1.11	0.93	1.05	1.65	1.15	97	1	9	2	1598	17	106
26	0.98	2.41	1.48	0.47	0.97	1.55	0.99	85	3	12	1	1479	16	92
27	1.11	2.41	1.73	0.93	0.97	1.26	1.15	96	3	14	2	1480	13	106
28	1.15	3.21	1.73	2.33	0.96	1.55	1.30	99	4	14	5	1463	16	120
29	1.09	3.21	2.35	1.40	1.09	1.36	1.20	94	4	19	3	1662	14	111
30	0.81	0	0.37	1.86	1.04	0.87	0.78	70	0	3	4	1575	9	72
31	1.04	0	0.25	0.47	1.06	0.58	0.93	90	0	2	1	1612	6	86
32	1.20	0.80	1.61	0.93	1.08	1.26	1.07	104	1	13	2	1645	13	99
33	0.97	0	1.36	0.47	0.95	1.07	0.91	84	0	11	1	1441	11	84
34	1.02	0.80	1.73	0	0.93	1.65	0.99	88	1	14	0	1405	17	92
35	0.98	0	1.11	0	0.99	1.07	0.97	85	0	9	0	1505	11	90
36	1.25	0.80	1.73	0.47	1.10	1.84	1.29	108	1	14	1	1665	19	119
37	0.80	0.80	0.49	2.33	1.05	0.87	0.76	69	1	4	5	1589	9	70
38	0.73	1.61	0.49	0.93	1.05	1.07	0.71	63	2	4	2	1596	11	66
39	0.80	0	0.99	0.93	1	1.07	0.95	69	0	8	2	1515	11	88
40	0.84	0	1.48	1.86	0.91	0.48	0.85	73	0	12	4	1387	5	79
41	0.83	0	0.99	0.47	0.81	1.45	0.92	72	0	8	1	1229	15	85
42	0.59	0	0.37	0.47	0.44	0.58	0.50	51	0	3	1	669	6	46
43	1.26	0	1.73	1.40	1.05	0.39	1.22	109	0	14	3	1598	4	113
44	0.86	1.61	0.49	0.93	1.07	0.58	0.75	74	2	4	2	1625	6	69
45	1.01	0	0.25	0.47	1.05	0.68	0.86	87	0	2	1	1599	7	80
46	0.98	0	0.74	1.86	1.04	0.97	0.97	85	0	6	4	1584	10	90
47	0.94	0	0.49	1.40	0.91	0.77	0.99	81	0	4	3	1386	8	92
48	0.82	0	0.87	1.40	0.76	0.39	0.83	71	0	7	3	1150	4	77
49	0.86	2.41	0.87	1.40	0.76	0.39	0.82	74	3	7	3	1148	4	76
50	1.31	0.80	1.24	0.47	1.15	1.16	1.25	113	1	10	1	1741	12	116
51	1.09	0.80	0.62	1.86	1.13	1.74	1.11	94	1	5	4	1709	18	103
52	1.19	0	0.25	0.47	1.15	0.97	1.14	103	0	2	1	1747	10	105
53	1.01	0	1.73	0.47	1.08	0.97	0.96	87	0	14	1	1640	10	89
54	1.13	0.80	2.10	0.93	0.99	1.65	1.22	98	1	17	2	1497	17	113
55	1.23	1.61	2.23	1.40	0.97	2.13	1.52	106	2	18	3	1479	22	141
56	1.27	0	1.24	1.40	1.01	0.77	1.29	110	0	10	3	1534	8	119
57	1.16	4.82	1.48	0.47	1.07	0.97	1.20	100	6	12	1	1621	10	111
58	0.72	0.80	0.12	1.40	1.06	0.39	0.64	62	1	1	3	1614	4	59
59	0.90	1.61	0.49	1.86	1.07	1.26	0.96	78	2	4	4	1631	13	89
60	1.08	0	0.74	1.86	1.02	0.58	1.05	93	0	6	4	1549	6	97

Table 8.4: The table shows the absolute values and the relative differences to the long-term mean values of each feature from 1 January, 2016 to 1 March, 2016. The WEF took place on day 20-23 (in bold).

On the other hand, we do not notice significant and repeatable changes in the surveillance area during the WEF in the other examined features. Concretely, there is little to no absolute change in helicopters, government and military aircraft, overall flight activity, or unknown aircraft during the WEF compared to its run-up and aftermath. It has to be noted that the first three of these are very low in volume, making even small absolute changes large in relative terms and limiting their informative value considerably.

8.5.4 Discussion

Large-scale event detection using air traffic communication data can provide potential insights into business movements and political events, even when the actually transmitted aircraft identifiers are pseudonymised or unknown. For example, as long as the mapping to the type of aircraft is available, events attracting many business or other interesting aircraft can reliably be detected. Even more trivial, comparing one's own sensor vision of the airspace with the one provided by public trackers in order to detect non-displayed aircraft (blocked for whatever reason) is possible without this mapping or any other previously obtained information. These low requirements make open source event detection achievable for even the most low-resource actor.

There are some natural pitfalls in our analysis as presented here: First, the data quality and consistency needs to be ensured. Variations in the reception quality of sensors (caused, for example, by construction or other disturbances) may distort the underlying data. However, we have shown that our approach works consistently over three years with different dates, different sensors, and different locations.

Secondly, this type of open-source information collection and anomaly detection also cannot tell us what event exactly is happening but it is a first step in a typical open-source intelligence process which collects open source data to first extract information and subsequently intelligence [188].

Lastly, the additional use of information gained from other aviation technologies such as Mode S or ACARS can further increase the accuracy and precision of the presented results (as indicated in [90] for the case of ACARS).

8.6 Summary

The work presented in this chapter shows that aviation is facing severe new privacy challenges in relation to the wireless communication used for air traffic. Similar to the security issues discussed in this thesis, the privacy concerns and their severity are influenced by the enabling SDR technology.

While some of the results are preliminary, it is clear that the use of relatively simple and publicly accessible means allows any inclined person to perform effective tracking of aircraft movements. This tracking ranges from the level of the individual aircraft to large-scale movement correlations and it cannot be prevented by current privacy-preserving solutions.

While for the security vulnerabilities we presented several quickly-deployable solutions in this dissertation, which are themselves based on cheap COTS hard- and software, the case is different for the discussed privacy concerns. Here, transparent solutions may prove more difficult. Instead, the aviation community must consider these challenges and begin to develop new communication technologies that serve the needs of its various classes of users and stakeholders.

One never notices what has been done; one can only see what remains to be done.

— Marie Curie [189]

9

Summary & Future Work

Contents

9.1	Summary of Results	161
9.2	Future Work	162
9.3	Final Conclusions	164

9.1 Summary of Results

This work aims to take a systematic view of the security of wireless communication technologies in aviation. We first analysed the current aviation environment as a whole with regards to its communication usage and attitude towards wireless security. Until now, security analyses and attacks on air traffic communication often focused on isolated protocols and ignored crucial domain experience.

By integrating the security knowledge from the academic and hacker communities, technology standards, and the opinions of international aviation experts, we provided a detailed overview of the technologies, their vulnerabilities, and existing attacks.

We further examined the aviation community's awareness concerning wireless systems security and collected expert opinions on the safety impact of potential attacks. Our results motivate the need to reassess the attack risks under realistic system models and the development of appropriate short- and long-term countermeasures.

Recognizing the need for cost-efficient, easily deployable attack detection systems, this thesis proposed, implemented, and evaluated several effective methods of detecting injections and manipulations of wireless ATC systems. While they still require practical testing in real-world aviation system environments, we believe they can significantly improve the current security until long-term solutions have been found.

However, it is important to view the full consequences of the proliferation of SDR-based sensor networks, which are able to listen and store all unencrypted aviation communications around the world. While we clearly laid out that such networks can be part of the solution and help securing the airspace, they also highlight new privacy issues and intelligence extraction possibilities. While not immediately safety-related, the impact of the privacy concerns we considered in the previous chapter will need to be considered by the aviation community. As such, we hope that it can serve as a starting point for future research and help to make informed decisions about the type and the confidentiality of the information that is broadcast by all aviation users.

9.2 Future Work

In the larger space of air traffic communication, there are many areas in which future work is required to secure the airspace in the longer term. Among these, with relevance to this dissertation, we identify the following five as most urgent:

- **Raising awareness:** To increase the security and privacy of the aviation system, awareness of cyber security issues among aviation circles and governments is a key factor. Only by raising awareness can the necessary research and development happen, enabling the responsible bodies to address the problem, and preventing the exploitation of existing vulnerabilities in the future. Without such awareness concerning the criticality of existing vulnerabilities, the necessary change will likely not come about before a real-world accident occurs.

Likewise, informing developers and users of aviation communication about the existing privacy issues is imperative so they are not provided with a false sense of security. Widespread knowledge about the non-confidentiality of aviation communication technologies can reduce reliance on these technologies for sensitive data and even modify user behaviour. Increased awareness can also factor into risk assessments and improve user response in case of a real breach.

- **Adapting regulations:** Further, we argue that top-down regulations are crucial in an industry such as aviation that is very cost-conscious and where actions are often taken only when required by regulators. Tying in with the point about awareness, governments and authorities need to be put in a knowledgeable position to issue the necessary regulations, and they should further consider the effect of their actions – or inaction – on the future security of the air traffic communication system. On the other hand, regulations for end users (of SDRs or tracking websites) are likely to be ineffective, due to their already widespread availability as well as their accessibility and the ease of sharing across (regulatory) borders.
- **Real-world penetration testing:** To gauge the full impact of attacks on all wireless technologies used in aviation, penetration testing of the systems as used in practice is required. While attacks on any single technology are trivial, little is known about the concrete effects in the real world. Many of the deployed ATC systems are highly proprietary and essentially acting as a black box between the reception of wireless messages and, for example, their final display on ATC radar screens. A thorough practical investigation will lay the foundation for the final two points on this agenda.
- **Development of secure new protocols:** Considering the decade-long development and certification cycles, research on protocols that include security by design is required as quickly as possible even though it will only pay in the long-term. Existing examples of security designs and analyses for the ADS-B

protocol as outlined in Chapter 6 can inform the directions of such future research, both for ATC-related protocols and information services.

Besides increasing the safety of the airspace, new protocols can also provide improvements for the issues of aviation privacy and secrecy as discussed in Chapter 8. With proper design and implementation of pseudonymous identifiers, most of the relevant information leakage could be reduced to the level of previous, non-technologically enhanced, plane-spotting days. Thus, the impact of SDRs and sensor networks which exploit unencrypted air traffic communication or potential side-channels could be minimized. This point concerns in particular military, governmental, and private aviation.

- **Deployment of attack detection solutions:** Last but not least, we have argued in this work that it is important to deploy defence strategies that do not require modifications of existing infrastructure and protocols. Attack detection methods, such as the one described in Chapter 7, rely on cyber-physical defences such as improved localization protocols, statistical analysis, machine learning, and physical-layer security. They can be deployed within a small time window and work transparently, thus immediately improving the security of communications for aviation users without a costly and time-consuming overhaul of existing systems.

9.3 Final Conclusions

A systematic awareness of the existing issues in wireless networking is maybe the most important factor contributing towards safer skies in the future. With the trend going towards more automated data networks communication, we strongly believe that aviation should catch up with the state of the art in wireless security to maintain its excellent safety record and reputation in the future. One survey comment noted that regulations are crucial in an industry such as aviation which is very cost-conscious and that actions were typically taken only when required by regulators.

In terms of academic research, it is just as crucial that future security developments do not ignore the domain-specific knowledge and requirements of aviation, which is something we have tried to outline throughout this work. Focusing on isolated problems or technologies without taking the whole system into account will inevitably lead to impractical solutions dismissed by the aviation community.

In the future, aviation will require newly developed secure solutions for all applications and the security community must be strongly involved from the start. As one aviation expert summarized their answers to our survey: “*These questions [on security] are silly. Remember aviation is behind 30 years.*”

However, as outlined in the beginning of this thesis, the rise of software-defined radios and easy accessibility of wireless air traffic communication will not be the only paradigm change threatening the safety of the future airspace. Other problems such as the safe integration of UAV need to be addressed, and further, yet unforeseen, disruptions are bound to happen in the future (e.g., the potential threat of quantum computing to newly developed cryptography-based communication protocols). As such, aviation needs to draw the right conclusions from current developments and develop processes that adapt more quickly to these challenges, both on the technical and the human side.

Appendices

A

Large-Scale Event Detection Data 2014-2015

This Appendix provides the tables with the complete data for 2014 and 2015 as discussed in Section 8.5.

Day	Relative Feature Values							Absolute Feature Values						
	Bus	Hel	Mil	Gov	All	Unk	Blo	Bus	Hel	Mil	Gov	All	Unk	Blo
1	1.18	0.32	0.53	0.60	1.03	0.87	0.96	108	2	5	1	2060	35	325
2	1.24	0.63	1.38	0	1.08	0.90	1.21	114	4	13	0	2153	36	408
3	1.11	1.26	1.60	1.81	0.96	0.77	0.96	102	8	15	3	1914	31	324
4	1.09	0.63	1.60	1.81	0.98	0.87	1.08	1	4	15	3	1954	35	365
5	1.02	0.32	1.38	2.41	0.98	0.97	1.01	94	2	13	4	1963	39	343
6	1.10	0.32	0.53	2.41	1.08	1.52	1.24	101	2	5	4	2162	61	418
7	0.77	0.32	0.32	0	1	1.12	0.97	71	2	3	0	1996	45	329
8	1.01	0	0.32	0	1.05	0.80	0.92	93	0	3	0	2099	32	311
9	1.08	1.11	1.06	0	1.05	1.12	1.16	99	7	10	0	2102	45	392
10	1.14	1.74	1.70	0.60	1.05	1.22	1.24	105	11	16	1	2090	49	420
11	1.18	2.05	0.64	1.21	1.04	0.97	1.16	108	13	6	2	2086	39	394
12	1.10	0.95	1.38	3.02	1.05	1.10	1.20	101	6	13	5	2091	44	405
13	1.19	0.16	0.64	1.21	1.15	1.30	1.17	109	1	6	2	2291	52	397
14	0.88	0.63	0.21	0	1.10	1.05	0.90	81	4	2	0	2199	42	306
15	0.92	0.47	0.74	1.81	1.07	0.82	0.97	84	3	7	3	2130	33	329
16	0.86	0	0.11	0.60	1.02	1.05	1	79	0	1	1	2041	42	337
17	0.66	0	0.32	1.21	0.77	0.60	0.76	61	0	3	2	1542	24	257
18	0.52	0	0	1.81	0.54	0.30	0.54	48	0	0	3	1070	12	184
19	0.78	0	0.21	0	0.94	0.90	1.02	72	0	2	0	1870	36	345
20	1.05	0.79	0.32	0.60	1.12	1.32	1.31	96	5	3	1	2244	53	442
21	0.93	0.16	0.53	1.81	1.04	0.67	0.96	85	1	5	3	2086	27	324
22	0.73	0.32	0.11	0.60	1.02	0.87	0.80	67	2	1	1	2041	35	269
23	0.95	0.95	0.64	0.60	1.02	0.85	1.01	87	6	6	1	2034	34	342
24	0.72	1.26	0.32	1.81	0.77	0.67	0.80	66	8	3	3	1534	27	270
25	0.63	0.63	0.32	0.60	0.79	0.67	0.70	58	4	3	1	1584	27	236
26	0.90	0	0.21	0	0.96	0.57	0.74	83	0	2	0	1917	23	249
27	0.82	0.95	0.53	1.21	1.10	1.05	1	75	6	5	2	2203	42	339
28	0.88	0	0.64	1.21	1.03	0.55	0.77	81	0	6	2	2052	22	261
29	1.24	0.16	0.43	2.41	1.07	0.62	0.87	114	1	4	4	2144	25	293
30	0.99	0.32	1.49	0	1.09	1.15	1.12	91	2	14	0	2178	46	378
31	1.09	4.27	1.06	0.60	0.96	1.07	1.02	1	27	10	1	1921	43	346
32	1.13	3	1.91	1.21	1.01	1.10	1.03	104	19	18	2	2015	44	347
33	1.15	2.21	1.49	0.60	0.99	1.07	1.06	106	14	14	1	1986	43	359
34	1.22	0.95	1.60	0	1.09	1.30	1.08	112	6	15	0	2185	52	365
35	0.87	0.32	0.21	0	1.04	1.15	0.86	80	2	2	0	2080	46	290
36	1.04	0.79	0.43	1.81	1.06	1.07	0.89	95	5	4	3	2115	43	3
37	0.80	0.32	1.60	0.60	1	0.95	0.88	73	2	15	1	1990	38	299
38	0.81	0.79	0.53	1.21	0.88	1.10	0.95	74	5	5	2	1751	44	323
39	0.96	1.42	1.38	0	0.97	1	1.10	88	9	13	0	1943	40	372
40	0.85	2.05	1.81	0	0.93	0.90	0.82	78	13	17	0	1860	36	278
41	1.13	1.11	2.02	2.41	1.03	0.85	1	104	7	19	4	2054	34	338
42	0.73	0.32	0.32	0.60	0.97	1.40	0.79	67	2	3	1	1928	56	267
43	0.87	0.32	0.21	0	0.99	1.27	0.73	80	2	2	0	1978	51	247
44	1	1.58	1.28	5.43	0.98	1.40	1.04	92	10	12	9	1950	56	352
45	1.27	2.69	0.96	3.02	0.94	0.97	1.26	117	17	9	5	1884	39	425
46	1.32	2.05	1.06	0.60	1	1.75	1.32	121	13	10	1	1995	70	446
47	1.15	1.42	1.81	2.41	0.98	0.95	1.16	106	9	17	4	1963	38	393
48	1.67	2.05	2.23	0.60	1.13	1.37	1.48	153	13	21	1	2252	55	501
49	1.56	1.74	1.17	0	1.05	1	1.01	143	11	11	0	2096	40	340
50	1.21	0.95	0.43	0	1.05	1.32	1.08	111	6	4	0	2094	53	365
51	0.98	0.32	1.28	0	1.03	1.10	1.10	90	2	12	0	2059	44	371
52	1	1.58	2.02	0	0.94	1.27	1.11	92	10	19	0	1887	51	376
53	1.08	0.95	2.34	0.60	0.95	1.25	1.14	99	6	22	1	1906	50	385
54	1.17	3.79	2.23	1.81	0.99	1.27	1.08	107	24	21	3	1979	51	365
55	1.06	1.42	1.81	1.21	1.07	1.07	1.02	97	9	17	2	2141	43	344
56	0.89	1.58	1.28	1.21	1.04	0.77	0.79	82	10	12	2	2082	31	268
57	0.97	0.32	0.53	3.62	1.03	0.70	0.89	89	2	5	6	2061	28	302
58	0.84	2.37	1.49	0	1	0.95	0.99	77	15	14	0	26	38	334
59	0.85	1.26	1.49	1.21	0.89	1.20	0.89	78	8	14	2	1777	48	3
60	0.97	0.47	1.06	0.60	0.96	0.82	1	89	3	10	1	1908	33	339

Table A.1: The table shows the absolute values and the relative differences to the long-term mean values of each feature from 8 December, 2013 to 5 February, 2014. The WEF took place on day 46-49 (in bold).

Day	Relative Feature Values							Absolute Feature Values						
	Bus	Hel	Mil	Gov	All	Unk	Blo	Bus	Hel	Mil	Gov	All	Unk	Blo
1	0.64	0	0.37	3.39	0.88	0.83	0.87	41	0	2	3	993	4	121
2	1.08	1.61	0.56	2.26	1.14	0.83	1.06	69	1	3	2	1296	4	147
3	1.11	0	0	1.13	1.15	0.41	0.96	71	0	0	1	1306	2	133
4	0.85	0	0.19	1.13	1.07	0.21	0.75	54	0	1	1	1217	1	104
5	0.74	1.61	0	0	1.03	1.03	1.02	47	1	0	0	1165	5	141
6	0.81	1.61	0.56	0	0.97	1.03	1.02	52	1	3	0	11	5	141
7	0.74	1.61	1.31	1.13	0.95	0.21	0.80	47	1	7	1	1076	1	111
8	0.92	4.82	1.31	1.13	1.02	0.62	0.93	59	3	7	1	1159	3	128
9	0.78	0	1.12	1.13	1.08	0.62	0.96	50	0	6	1	1228	3	133
10	0.80	0	0.56	3.39	1.09	0.21	0.77	51	0	3	3	1239	1	106
11	1.16	0	0.37	0	1.05	0.21	0.94	74	0	2	0	1192	1	130
12	0.95	1.61	2.43	0	1.02	0.62	0.98	61	1	13	0	1150	3	136
13	0.77	1.61	0.56	2.26	0.93	0.62	1	49	1	3	2	1052	3	138
14	0.97	1.61	1.31	0	0.90	0.41	0.92	62	1	7	0	1022	2	127
15	0.95	1.61	1.31	0	1	1.45	1.13	61	1	7	0	1137	7	156
16	1.03	0	1.31	0	1	1.03	0.93	66	0	7	0	1127	5	129
17	0.77	0	0.19	0	1.02	1.45	1.09	49	0	1	0	1160	7	151
18	0.92	1.61	0.19	0	1.01	0.83	1.01	59	1	1	0	1140	4	140
19	1.16	3.21	0.56	3.39	0.95	0.62	1.03	74	2	3	3	1080	3	142
20	1.71	1.61	0.94	3.39	0.91	1.24	1.37	109	1	5	3	1026	6	190
21	1.52	4.82	2.25	3.39	0.90	1.03	1.24	97	3	12	3	1017	5	172
22	1.28	1.61	2.06	1.13	0.94	1.03	1.13	82	1	11	1	1068	5	156
23	1.93	1.61	2.62	3.39	1.08	0.83	1.61	123	1	14	3	1226	4	222
24	1.61	3.21	0.56	1.13	1	0.41	1.22	103	2	3	1	1128	2	169
25	1	0	0.94	2.26	1.06	1.45	1.21	64	0	5	2	12	7	167
26	1.06	1.61	1.31	1.13	1.03	2.69	1.38	68	1	7	1	1162	13	191
27	0.99	0	2.25	1.13	0.93	2.27	1.36	63	0	12	1	1048	11	188
28	1.21	3.21	1.68	1.13	0.89	1.65	1.12	77	2	9	1	17	8	155
29	1.03	0	1.31	1.13	0.92	1.03	1.02	66	0	7	1	1037	5	141
30	0.85	0	0.56	0	1.03	1.65	1.08	54	0	3	0	1172	8	150
31	0.66	0	0.37	0	0.99	1.03	0.75	42	0	2	0	1122	5	104
32	0.67	0	0.19	1.13	1.06	2.69	1.09	43	0	1	1	1195	13	151
33	0.95	0	1.50	0	1	2.69	1.27	61	0	8	0	1138	13	175
34	0.81	3.21	0.75	0	0.91	1.45	1.11	52	2	4	0	1036	7	153
35	0.85	0	0.94	0	0.91	1.24	0.95	54	0	5	0	1028	6	132
36	0.94	0	0.94	0	0.96	0.83	0.93	60	0	5	0	1086	4	129
37	1.06	0	0.94	1.13	1.02	1.03	0.97	68	0	5	1	1153	5	134
38	0.77	0	0.37	2.26	1.02	0.41	0.70	49	0	2	2	1153	2	97
39	0.72	0	0.19	3.39	1.06	1.86	1	46	0	1	3	1198	9	138
40	0.89	1.61	1.12	0	0.96	2.27	1.02	57	1	6	0	1083	11	141
41	1.02	0	1.68	0	0.92	1.03	0.89	65	0	9	0	1046	5	123
42	0.92	1.61	2.06	3.39	0.90	0.83	0.89	59	1	11	3	1022	4	123
43	0.97	0	0.75	1.13	0.94	0.41	0.91	62	0	4	1	1059	2	126
44	1.21	0	0.19	1.13	1.04	0.83	1.10	77	0	1	1	1175	4	152
45	0.95	0	0.19	1.13	1.03	0.41	0.79	61	0	1	1	1169	2	109
46	0.94	0	0.37	2.26	1.05	0.41	0.72	60	0	2	2	1193	2	1
47	0.97	0	0.94	1.13	0.99	0.83	0.98	62	0	5	1	1124	4	136
48	0.89	1.61	0.94	0	0.95	1.03	0.90	57	1	5	0	1075	5	125
49	1.06	3.21	2.25	1.13	0.95	1.03	1	68	2	12	1	1075	5	138
50	1.05	1.61	1.68	0	1	1.24	1.06	67	1	9	0	1133	6	146
51	1.30	3.21	2.25	2.26	1.10	0.62	1.06	83	2	12	2	1249	3	147
52	0.88	0	0	0	1.07	0.83	0.75	56	0	0	0	1216	4	104
53	1.28	1.61	0.19	0	1.08	0.62	0.93	82	1	1	0	1225	3	128
54	1	0	0.75	0	1.01	0.83	0.97	64	0	4	0	1140	4	134
55	0.80	0	1.68	0	0.91	1.03	0.85	51	0	9	0	1033	5	117
56	1	0	1.50	0	0.95	1.03	0.92	64	0	8	0	1080	5	127
57	1.08	3.21	1.50	0	1	1.24	1.06	69	2	8	0	1129	6	146
58	0.97	0	1.12	0	1.07	1.03	1.01	62	0	6	0	1214	5	140
59	0.91	0	0.19	0	1.05	0.41	0.74	58	0	1	0	1192	2	103
60	0.97	0	0.37	0	1.10	0.62	0.75	62	0	2	0	1245	3	104

Table A.2: The table shows the absolute values and the relative differences to the long-term mean values of each feature from 1 January, 2015 to 1 March, 2015. The WEF took place on day 21-24 (in bold).

References

- [1] Boeing. *Boeing Current Market Outlook 2015-2034*. Tech. rep. Accessed June 2016. 2015. URL: <http://www.boeing.com/commercial/market/>.
- [2] US Department of Transportation. *Unmanned Aircraft System (UAS) Service Demand 2015- 2035*. Tech. rep. Accessed June 2016. 2013. URL: <https://fas.org/irp/program/collect/service.pdf>.
- [3] Cary R Spitzer, Uma Ferrell, and Thomas Ferrell. *Digital Avionics Handbook*. 3rd ed. CRC Press, 2014.
- [4] Nancy Moran and Gerrit De Vynck. “WestJet Says It Never Sent Hijack Alarm, Wasn’t in Danger”. In: *Bloomberg* (Jan. 2015). Accessed June 2016. URL: <http://goo.gl/gSy2oa>.
- [5] Alison Williams. “Jets vanishing from Europe radar linked to war games”. In: *Reuters* (June 2014). Accessed June 2016. URL: <http://goo.gl/qKUrRp>.
- [6] Heather Kelly. “Researcher: New air traffic control system is hackable”. In: *Cable News Network (CNN)* (July 2012). Accessed June 2016. URL: <http://goo.gl/5naCSS>.
- [7] Kim Zetter. “Air Traffic Controllers Pick the Wrong Week to Quit Using Radar”. In: *Wired* (July 2012). Accessed June 2016. URL: <http://www.wired.com/2012/07/adsb-spoofing/>.
- [8] Krishna Sampigethaya, Radha Poovendran, and Linda Bushnell. “A Framework for Securing Future e-Enabled Aircraft Navigation and Surveillance”. In: *AIAA Infotech@Aerospace Conference*. Apr. 2009, pp. 1–10.
- [9] Donald McCallie, Jonathan Butts, and Robert Mills. “Security analysis of the ADS-B implementation in the next generation air transportation system”. In: *International Journal of Critical Infrastructure Protection* 4.2 (Aug. 2011), pp. 78–87.
- [10] Brad Haines. *Hacker + Airplanes = No good can come of this*. Presented at DEFCON 20. Las Vegas, USA, July 2012.
- [11] Hugo Teso. *Aircraft hacking: Practical aero series*. Presented at The Fourth Annual Hack in the Box Security Conference in Europe (HITBSECCONF2013). Amsterdam, NL, Apr. 2013.
- [12] Andrei Costin and Aurélien Francillon. “Ghost is in the Air(traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices”. In: *Black Hat USA*. July 2012, pp. 1–12.

- [13] Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. “Experimental analysis of attacks on next generation air traffic communication”. In: *International Conference on Applied Cryptography and Network Security (ACNS)*. Springer. 2013, pp. 253–271.
- [14] Devin Lundberg, Brown Farinholt, Edward Sullivan, Ryan Mast, Stephen Checkoway, Stefan Savage, Alex C Snoeren, and Kirill Levchenko. “On The Security of Mobile Cockpit Information Systems”. In: *Proceedings of the 21st ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM. Nov. 2014, pp. 633–645.
- [15] Paul Marks. “Air traffic system vulnerable to cyber attack”. In: *New Scientist* (Sept. 2011). Accessed June 2016. URL: <https://goo.gl/G0FQTs>.
- [16] Steve Henn. “Could The New Air Traffic Control System Be Hacked?”. In: *National Public Radio (NPR)* (Aug. 2012). Accessed June 2016. URL: <http://goo.gl/pfJn61>.
- [17] Andy Greenberg. “Next-Gen Air Traffic Control Vulnerable To Hackers Spoofing Planes Out Of Thin Air”. In: *Forbes* (July 2012). Accessed June 2016. URL: <http://goo.gl/1uxBXw>.
- [18] Mark Clayton. “Malaysia Airlines Flight MH370: Are planes vulnerable to cyber-attack?”. In: (Mar. 2014). Accessed June 2016. URL: <http://goo.gl/xXn8eM>.
- [19] Neil McAllister. “FAA: ‘No, you CAN’T hijack a plane with an Android app’”. In: *The Register* (Apr. 2013). Accessed June 2016. URL: <http://goo.gl/nUoFP3>.
- [20] Phil Polstra and Captain Polly. *Cyber-hijacking Airplanes: Truth or Fiction?* Presented at DEFCON 22. Las Vegas, USA, Aug. 2014.
- [21] Gavin Walker. “Is air traffic control a soft target for hackers?”. In: *NATS Blog* (Oct. 2013). Accessed June 2016. URL: <http://goo.gl/zDy3rE>.
- [22] Martin Strohmeier, Matthew Smith, Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. “Assessing the Impact of Aviation Security on Cyber Power”. In: *8th International Conference on Cyber Conflict (CyCon)*. NATO CCD COE. 2016, pp. 223–241.
- [23] Martin Strohmeier, Matthias Schäfer, Rui Pinheiro, Vincent Lenders, and Ivan Martinovic. “On Perception and Reality in Wireless Air Traffic Communication Security”. In: *IEEE Transactions on Intelligent Transportation Systems* PP.99 (2016), pp. 1–20.
- [24] Martin Strohmeier, Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. “Realities and Challenges of NextGen Air Traffic Management: The Case of ADS-B”. In: *IEEE Communications Magazine* 52.5 (2014), pp. 111–118.
- [25] Martin Strohmeier and Ivan Martinovic. “On Passive Data Link Layer Fingerprinting of Aircraft Transponders”. In: *Proceedings of the First ACM Workshop on Cyber-Physical Systems-Security and/or Privacy (CPS-SPC)*. ACM. 2015, pp. 1–9.

- [26] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. “Intrusion Detection for Airborne Communication using PHY-Layer Information”. In: *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*. Springer. 2015, pp. 67–77.
- [27] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. “Lightweight location verification in air traffic surveillance networks”. In: *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security (CPSS)*. ACM. 2015, pp. 49–60.
- [28] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. “A Localization Approach for Crowdsourced Air Traffic Communication Networks”. In: *arXiv preprint arXiv:1610.06754* (Oct. 2016).
- [29] Martin Strohmeier, Ivan Martinovic, Markus Fuchs, Matthias Schäfer, and Vincent Lenders. “OpenSky: A Swiss army knife for air traffic security research”. In: *34th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. IEEE/AIAA. Sept. 2015, pp. 1–14.
- [30] Matthias Schäfer, Martin Strohmeier, Vincent Lenders, Ivan Martinovic, and Matthias Wilhelm. “Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research”. In: *Proceedings of The 13th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. Apr. 2014, pp. 83–94.
- [31] Nils Ole Tippenhauer, Christina Pöpper, Kasper Bonne Rasmussen, and Srdjan Čapkun. “On the Requirements for Successful GPS Spoofing Attacks”. In: *Proceedings of the 18th ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, Oct. 2011, pp. 75–86.
- [32] Mohamed Mahmoud, Alain Pirovano, and Nicolas Larrieu. “Aeronautical communication transition from analog to digital data: A network security survey”. In: *Elsevier Computer Science Review* 11 (May 2014), pp. 1–29.
- [33] Sean S. Swords. *Technical History of the Beginnings of Radar*. 1st ed. London, UK: The Institution of Engineering and Technology, 1985.
- [34] David Adamy. *EW 101: A first course in electronic warfare*. Artech House, 2001.
- [35] Farid Dowla. *Handbook of RF and wireless technologies*. Newnes, 2003.
- [36] Eric Blossom. “GNU radio: tools for exploring the radio frequency spectrum”. In: *Linux Journal* 2004.122 (2004), p. 4.
- [37] Marko Wolf, Moritz Minzlaff, and Martin Moser. “Information Technology Security Threats to Modern e-Enabled Aircraft: A Cautionary Note”. In: *AIAA Journal of Aerospace Information Systems* 11.7 (2014), pp. 447–457.
- [38] Victoria Bryan and Peter Maushagen. “Flightradar24 finds not just planespotters flocking to its website”. In: *Reuters* (May 2015). Accessed Jun 2016. URL: <http://www.reuters.com/article/us-airlines-flightradar-idUSKBN0031Q720150518>.
- [39] Benjamin Hart. “Flight Radar Shows Planes Avoiding Ukraine In Aftermath Of Malaysia Airlines Crash”. In: *Huffington Post* (July 2014). Accessed June 2016. URL: http://www.huffingtonpost.com/2014/07/17/flight-radar-ukraine-crash-malaysia_n_5596574.html.

- [40] Gwyn Topham. “Malaysian Airlines plane mystery: how can a flight disappear off radar?” In: *The Guardian* (Mar. 2014). Accessed June 2016. URL: <https://www.theguardian.com/world/2014/mar/10/malaysia-airlines-plane-mystery-disappear-off-radar>.
- [41] Keith Stouffer, Joe Falco, and Karen Scarfone. “Guide to industrial control systems (ICS) security”. In: *NIST special publication* (2011), pp. 800–882.
- [42] Chris. W. Johnson. “You Outsource the Service but Not the Risk: Supply Chain Risk Management for the Cyber Security of Safety Critical Systems”. In: *34th International System Safety Conference (ISSC 2016)*. 2016.
- [43] Gerald L Dillingham, Gregory C Wilshusen, and Nabajyoti Barkakati. *Air Traffic Control: FAA Needs a More Comprehensive Approach to Address Cybersecurity As Agency Transitions to NextGen*. Tech. rep. GAO-15-370. 2015.
- [44] David Hollis. “Cyberwar case study: Georgia 2008”. In: *Small Wars Journal* 7 (Jan. 2011), pp. 1–10.
- [45] Pierre Augustin Caron de Beaumarchais, Jacques Joseph Marie Decroix, et al. *Oeuvres complètes de Voltaire: Commentaires sur Corneille*. Vol. 44. Carez, Thomine et Fortic, 1822.
- [46] *International Standards and Recommended Practices, Annex 10: Aeronautical Telecommunications*. 2nd ed. Volume III: Communication Systems. International Civil Aviation Organization (ICAO). 2007.
- [47] *International Standards and Recommended Practices, Annex 10: Aeronautical Telecommunications*. 2nd ed. Volume V: Aeronautical Radio Frequency Spectrum Utilization. International Civil Aviation Organization (ICAO). 2001.
- [48] Merrill I. Skolnik. *Radar Handbook*. 3rd ed. Electronics electrical engineering. The McGraw-Hill Companies, 2008.
- [49] *International Standards and Recommended Practices, Annex 10: Aeronautical Telecommunications*. 4th ed. Volume IV: Surveillance and Collision Avoidance Systems. International Civil Aviation Organization (ICAO). 2007.
- [50] *Minimum Operational Performance Standards for 1090 MHz Extended Squitter Automatic Dependent Surveillance – Broadcast (ADS-B) and Traffic Information Services – Broadcast (TIS-B)*. Tech. rep. DO-260B (with Corrigendum 1). RTCA, Inc., Dec. 2011.
- [51] *Standards and Recommended Practices for the Universal Access Transceiver (UAT)*. Revision 5.0. International Civil Aviation Organization (ICAO). Apr. 2005.
- [52] *Manual on the Universal Access Transceiver (UAT): Detailed Technical Specifications*. 1st ed. Revision 4.1. International Civil Aviation Organization (ICAO). June 2005.
- [53] *Minimum Operational Performance Standards (MOPS) for Aircraft VDL Mode 2 Physical Link and Network Layer*. Tech. rep. DO-281B. RTCA, Inc, Mar. 2012.
- [54] Patricia Massimini, James E Dieudonne, Leone C Monticone, Dean F Lamiano, Edward Brestle, et al. “Insertion of controller-pilot data link communications into the National Airspace System: is it more efficient?” In: *18th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. Vol. 1. 1999, pp. 1–6.

- [55] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. “On the Security of the Automatic Dependent Surveillance-Broadcast Protocol”. In: *IEEE Communications Surveys & Tutorials* 17.2 (2015), pp. 1066–1087.
- [56] RTCA Inc. *Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast (ADS-B)*. DO-242A (including Change 1). Dec. 2006.
- [57] RTCA Inc. *Minimum Operational Performance Standards for 1090 MHz Extended Squitter Automatic Dependent Surveillance – Broadcast (ADS-B) and Traffic Information Services – Broadcast (TIS-B)*. DO-260B with Corrigendum 1. Dec. 2011.
- [58] RTCA Inc. *Minimum Operational Performance Standards for Universal Access Transceiver (UAT) Automatic Dependent Surveillance – Broadcast*. DO-282B with Corrigendum 1. Dec. 2011.
- [59] John R. Barry. *Wireless Infrared Communications*. Boston: Kluwer Academic, 1994.
- [60] Matthias Schäfer, Martin Strohmeier, Matthew Smith, Markus Fuchs, Rui Pinheiro, Vincent Lenders, and Ivan Martinovic. “OpenSky Report 2016: Facts, Figures and Trends in Wireless ATC Communication Systems”. In: *35th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. Sept. 2016.
- [61] Marion Loren Wood and Richard Wright Bush. *Multilateration on Mode S and ATCRBS Signals at Atlanta’s Hartsfield Airport*. Tech. rep. ATC-260. MIT Lincoln Laboratory, Jan. 1998.
- [62] *Air/Ground Character-Oriented Protocol Specification*. Tech. rep. 618-7. ARINC, June 2013.
- [63] *Minimum Operational Performance Standards for Traffic Alert and Collision Avoidance Systems II (TCAS II)*. Tech. rep. DO-185B. RTCA, Inc., June 2008.
- [64] *Minimum Operational Performance Standards for Flight Information Services Broadcast (FIS-B) with Universal Access Transceiver (UAT)*. Tech. rep. DO-358. RTCA, Inc., Mar. 2015.
- [65] *Surveillance and Broadcast Services Description Document*. Tech. rep. SRT-047, Revision 1. U.S. Department of Transport and Federal Aviation Administration, Oct. 2011.
- [66] International Civil Aviation Organization (ICAO). *2013–2028 Global Air Navigation Plan*. Tech. rep. 2013.
- [67] IEEE Std 802.16-2009. “IEEE Standard for local and metropolitan area networks Part 16: Air interface for fixed and mobile broadband wireless access systems”. In: (May 2009). Revision of IEEE Std 802.16-2004.
- [68] Nikos Fistas. *AeroMACS Briefing / Update*. Tech. rep. European Organization for the Safety of Air Navigation (Eurocontrol), Jan. 2016.
- [69] SESAR. *AeroMACS safety and security analysis*. Tech. rep. P15.02.07 D08. Jan. 2014.
- [70] Johann Wolfgang von Goethe. *Maximen und Reflexionen*. Ed. by Helmut Koopmann. Munich, Germany: Deutscher Taschenbuch Verlag and C.H.Beck, 2006.

- [71] Matthias Wilhelm, Ivan Martinovic, Jens B Schmitt, and Vincent Lenders. “Short paper: reactive jamming in wireless networks: how realistic is the threat?” In: *Proceedings of the fourth ACM conference on Wireless network security (WiSec)*. ACM. 2011, pp. 47–52.
- [72] Christina Pöpper, Nils Ole Tippenhauer, Boris Danev, and Srdjan Čapkun. “Investigation of signal and message manipulations on the wireless channel”. In: *Proceedings of the 16th European Conference on Research in Computer Security (ESORICS)*. Springer. 2011, pp. 40–59.
- [73] Matthias Wilhelm, Jens B Schmitt, and Vincent Lenders. “Practical message manipulation attacks in IEEE 802.15.4 wireless networks”. In: *Proceedings of the 16th international GI/ITG conference on Measurement, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance (MMB/DFT)*. 2012, pp. 29–34.
- [74] Tim Stelkens-Kobsch, Andreas Hasselberg, Thorsten Mühlhausen, and Nils Carstengerdes. “Towards a more secure ATC voice communications system”. In: *34th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. IEEE. 2015, pp. 1–9.
- [75] Romano Fantacci, Simone Menci, Luigia Micciullo, and Laura Pierucci. “A secure radio communication system based on an efficient speech watermarking approach”. In: *Security and Communication Networks* 2.4 (2009), pp. 305–314.
- [76] Mohamed Slim Ben Mahmoud, Alain Pirovano, Christophe Guerber, Nicolas Larrieu, and José Radzik. *Aeronautical Air-Ground Data Link Communications*. John Wiley & Sons, 2014.
- [77] Naval Air Systems Command and Naval Air Warfare Center. *Electronic Warfare and Radar Systems Engineering Handbook*. 2nd ed. Point Mugu, CA: Naval Air Warfare Center, 1999.
- [78] Aeronautical Surveillance Panel. *Draft Doc9924 Guidance Material for the measurement of All-Call Reply Rates*. Tech. rep. 2013.
- [79] Eurocontrol. *Updated Work on 5 - Final Report on Electromagnetic Environmental Effects of, and on, ACAS*. Tech. rep. Accessed June 2016. Aug. 2009. URL: <http://goo.gl/QqMZDT>.
- [80] European Aviation Safety Agency. *Results from EASA technical investigation on the radar detection losses in June 2014 in Central Europe*. Tech. rep. Jan. 2015.
- [81] Patrick J Martone and George E Tucker. “Candidate requirements for multilateration and ADS-B systems to serve as alternatives to secondary radar”. In: *20th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. 2001, pp. 1–12.
- [82] Christos Rekkas and Melvyn Rees. “Towards ADS-B implementation in Europe”. In: *2008 Tyrrhenian International Workshop on Digital Communications - Enhanced Surveillance of Aircraft and Vehicles*. IEEE, Sept. 2008, pp. 1–4.
- [83] Daniel Moser, Patrick Leu, Vincent Lenders, Aanjhan Ranganathan, Fabio Ricciato, and Srdjan Capkun. “Investigation of multi-device location spoofing attacks on air traffic control and possible countermeasures”. In: *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking (MobiCom)*. ACM. 2016, pp. 375–386.

- [84] International Civil Aviation Organization (ICAO). *Guidance Material: Security issues associated with ADS-B*. Tech. rep. Montreal, QC, Canada, 2014.
- [85] Krishna Sampigethaya, Radha Poovendran, Sudhakar Shetty, Terry Davis, and Chuck Royalty. “Future e-enabled aircraft communications and security: The next 20 years and beyond”. In: *Proceedings of the IEEE* 99.11 (2011), pp. 2040–2055.
- [86] Curtis Risley, James McMath, and Captain Brian Payne. “Experimental encryption of aircraft communications addressing and reporting system (ACARS) aeronautical operational control (AOC) messages”. In: *20th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. Oct. 2001, pp. 1–8.
- [87] *DataLink Security, Part 1 - ACARS Message Security*. Tech. rep. 823P1. ARINC, Dec. 2007.
- [88] Frank Leipold. *Expert Dialogue on Real-time Monitoring of Flight Data - Session 5: Views of Airlines and Pilots*. Tech. rep. May 2014.
- [89] Mary Kirby. “How Ryanair monitors health of Boeing 737s without ACARS”. In: *Runway Girl Network* (Sept. 2014). Accessed June 2016. URL: <http://www.runwaygirlnetwork.com/2014/09/06/how-ryanair-monitors-health-of-boeing-737s-without-acars/>.
- [90] Matt Smith, Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. “On the security and privacy of ACARS”. In: *IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS)* (2016), pp. 1–27.
- [91] *Minimum Aviation System Performance standards (MASPS) for Flight Information Service - Broadcast (FIS-B)*. Tech. rep. DO-267A. RTCA, Inc, Apr. 2004.
- [92] Richard Wagner. *IFF, Combat ID and the Information Domino Effect*. Tech. rep. Canadian Department of National Defence, 2009, pp. 1–22.
- [93] U.S. Congress, Office of Technology Assessment. *Who Goes There: Friend or Foe?* Washington, DC: U.S. Government Printing Office, 1993.
- [94] Larry Kenney, Joe Dietrich, and Jerry Woodall. “Secure ATC surveillance for military applications”. In: *IEEE Military Communications Conference (MILCOM)*. IEEE, Nov. 2008, pp. 1–6.
- [95] RTCA Inc. *Proposed Change to DO-181D and ED-73C for Higher Squitter Rates at Lower Power*. Special Committee 209 ATCRBS / Mode S Transponder MOPS Maintenance. Apr. 2007.
- [96] David Adamy. *EW 102: A second course in electronic warfare*. Norwood, MA: Artech House, 2004.
- [97] David Adamy. *EW 103: Tactical battlefield communications electronic warfare*. Norwood, MA: Artech House, 2008.
- [98] Charles J Middleton. *Risk Assessment Planning for Airborne Systems: An Information Assurance Failure Mode, Effects and Criticality Analysis Methodology*. Tech. rep. Wright-Patterson Air Force Base, Ohio: Air Force Institute of Technology, 2012.
- [99] Domenic Magazu III. “Exploiting the Automatic Dependent Surveillance-Broadcast system via false target injection”. MA thesis. Wright-Patterson Air Force Base, Ohio: Air Force Institute of Technology, 2012.

- [100] Meng Yue. “Security of VHF Data Link in ATM”. In: *Aeronautical Telecommunications Network: Advances, Challenges, and Modeling*. Ed. by M. Sarhan Musa and Zhijun Wu. CRC Press, 2015, pp. 65–92.
- [101] Pietro Pierpaoli, Magnus Egerstedt, and Amir Rahmani. “Altering UAV flight path by threatening collision”. In: *34th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. 2015, pp. 1–10.
- [102] Federal Aviation Administration (FAA). *FAR/AIM 2015: Federal Aviation Regulations/Aeronautical Information Manual*. Aviation Supplies & Academics Inc, 2014.
- [103] Bureau d’Enquetes et d’Analyses. *Final report on the accident on 1st June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France flight AF 447 Rio de Janeiro–Paris*. Tech. rep. Paris, France, 2012.
- [104] John Scardina. *Overview of the FAA ADS-B link decision*. Tech. rep. Federal Aviation Administration, June 2002.
- [105] Pameet Singh and Peter Sandborn. “Obsolescence driven design refresh planning for sustainment-dominated systems”. In: *The Engineering Economist* 51.2 (2006), pp. 115–139.
- [106] Mark G Ballin. “Next Frontier: Sharing the Airspace with Increased Autonomy”. In: *Digital Avionics Handbook*. Ed. by Cary R Spitzer, Uma Ferrell, and Thomas Ferrell. 3rd ed. CRC Press, 2014, pp. 779–794.
- [107] Corinne Mulley and John D Nelson. “Interoperability and transport policy: the impediments to interoperability in the organisation of trans-European transport systems”. In: *Journal of Transport Geography* 7.2 (1999), pp. 93–104.
- [108] Markus Franke. “Competition between network carriers and low-cost carriers—retreat battle or breakthrough to a new level of efficiency?” In: *Journal of Air Transport Management* 10.1 (2004), pp. 15–21.
- [109] International Civil Aviation Organization. *Review report of the thirteenth meeting of Automatic Dependent Surveillance-Broadcast (ADS-B) study and implementation task force*. Tech. rep. Beijing, 2014.
- [110] Sathya S Silva, Luke Jensen, and Robert J Hansman Jr. “Safety Benefit of Automatic Dependent Surveillance-Broadcast Traffic and Weather Uplink Services”. In: *Journal of Aerospace Information Systems* 12.8 (2015), pp. 579–586.
- [111] Mica R Endsley. “Toward a theory of situation awareness in dynamic systems”. In: *Human Factors: The Journal of the Human Factors and Ergonomics Society* 37.1 (Mar. 1995), pp. 32–64.
- [112] Mitchell J Nathan, Kenneth R Koedinger, and Martha W Alibali. “Expert blind spot: When content knowledge eclipses pedagogical content knowledge”. In: *International Conference on Cognitive Science*. USTC Press, 2001, pp. 644–648.
- [113] Jordi Bonet i Armengol. “El Templo de la Sagrada Familia Nuevas aportaciones al estudio de Gaudí”. In: *Loggia, Arquitectura & Restauración* 9 (2000), pp. 22–29.
- [114] Mark Luk, Adrian Perrig, and Bram Whillock. “Seven cardinal properties of sensor network broadcast authentication”. In: *Proceedings of the fourth ACM workshop on Security of ad hoc and sensor networks (SASN)*. ACM, 2006, pp. 147–156.

- [115] Adrian Perrig and Doug Tygar. *Secure Broadcast Communication in Wired and Wireless Networks*. New York, NY: Springer Science & Business Media, 2003.
- [116] Boris Danev, Davide Zanetti, and Srdjan Capkun. “On physical-layer identification of wireless devices”. In: *ACM Computing Surveys* 45.1 (Nov. 2012), pp. 1–29.
- [117] Boris Danev, Heinrich Luecken, Srdjan Capkun, and Karim El Defrawy. “Attacks on physical-layer identification”. In: *Proceedings of the third ACM conference on Wireless network security (WiSec)*. 2010, pp. 89–98.
- [118] Jeyanthi Hall, Michel Barbeau, and Evangelos Kranakis. “Enhancing intrusion detection in wireless networks using radio frequency fingerprinting.” In: *Proceedings of the 3rd IASTED International Conference on Communications, Internet and Information Technology (CIIT)*. 2004, pp. 201–206.
- [119] Roel Maes and Ingrid Verbauwhede. “Physically unclonable functions: A study on the state of the art and future research directions”. In: *Towards Hardware-Intrinsic Security*. Springer, 2010, pp. 3–37.
- [120] Srinivas Devadas, Edward Suh, Sid Paral, Richard Sowell, Tom Ziola, and Vivek Khandelwal. “Design and Implementation of PUF-Based "Unclonable" RFID ICs for Anti-Counterfeiting and Security Applications”. In: *IEEE International Conference on RFID*. IEEE, Apr. 2008, pp. 58–64.
- [121] Mario Strasser, Christina Pöpper, Srdjan Capkun, and Mario Cagalj. “Jamming-resistant Key Establishment using Uncoordinated Frequency Hopping”. In: *IEEE Symposium on Security and Privacy (S&P)*. IEEE, May 2008, pp. 64–78.
- [122] Christina Pöpper, Mario Strasser, and Srdjan Capkun. “Jamming-resistant Broadcast Communication without Shared Keys”. In: *Proceedings of the USENIX Security Symposium*. 2009, pp. 231–247.
- [123] Yao Liu, Peng Ning, Huaiyu Dai, and An Liu. “Randomized Differential DSSS: Jamming-Resistant Wireless Broadcast Communication”. In: *Proceedings of the 29th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)* (Mar. 2010), pp. 1–9.
- [124] Cindy Finke, Jonathan Butts, Robert Mills, and Michael Grimaila. “Enhancing the security of aircraft surveillance in the next generation air traffic control system”. In: *International Journal of Critical Infrastructure Protection* 6.1 (Mar. 2013), pp. 3–11.
- [125] Kyle D Wesson, Todd E Humphreys, and Brian L Evans. *Can Cryptography Secure Next Generation Air Traffic Surveillance?* Tech. rep. Radionavigation Laboratory, University of Austin at Texas, Jan. 2014.
- [126] Ken Samuelson, Ed Valovage, and Dana Hall. “Enhanced ADS-B Research”. In: *IEEE Aerospace Conference*. IEEE, 2006, pp. 1–7.
- [127] Richard V Robinson, Mingyan Li, Scott A Lintelman, Krishna Sampigethaya, Radha Poovendran, David Von Oheimb, and Jens-Uwe Bußer. “Impact of Public Key Enabled Applications on the Operation and Maintenance of Commercial Airplanes”. In: *7th AIAA Aviation Technology Integration, and Operations Conference (ATIO)*. 2007, pp. 1–10.

- [128] Bryan Parno and Adrian Perrig. “Challenges in Securing Vehicular Networks”. In: *Workshop on Hot Topics in Networks (HotNets-IV)*. 2005, pp. 1–6.
- [129] Junqi Zhang and Vijay Varadharajan. “Wireless sensor network key management survey and taxonomy”. In: *Journal of Network and Computer Applications* 33.2 (Mar. 2010), pp. 63–75.
- [130] Adrian Perrig, Ran Canetti, J Doug Tygar, and Dawn Song. “The TESLA broadcast authentication protocol”. In: *RSA CryptoBytes* 5 (2005).
- [131] Adrian Perrig, Robert Szewczyk, J D Tygar, Victor Wen, and David E Culler. “SPINS: Security Protocols for Sensor Networks”. In: *Wireless networks* 8 8.5 (2002), pp. 521–534.
- [132] Donggang Liu, Peng Ning, Sencun Zhu, and Sushil Jajodia. “Practical broadcast authentication in sensor networks”. In: *The Second Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous)*. IEEE. 2005, pp. 118–129.
- [133] Jason J. Haas, Yih-Chun Hu, and Kenneth P. Laberteaux. “Real-World VANET Security Protocol Performance”. In: *IEEE Global Telecommunications Conference (GLOBECOM)*. IEEE, Nov. 2009, pp. 1–7.
- [134] Andreas Savvides, Heemin Park, and Mani B. Srivastava. “The bits and flops of the n-hop multilateration primitive for node localization problems”. In: *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications (WSNA)*. New York, New York, USA: ACM, 2002, pp. 112–121.
- [135] J Garcia Herrero, JA Besada Portas, FJ Jimenez Rodriguez, and JR Casar Corredra. “ASDE and multilateration mode-S data fusion for location and identification on airport surface”. In: *The Record of the 1999 IEEE Radar Conference*. IEEE. 1999, pp. 315–320.
- [136] Julio C. Siu. “ICAO Concepts and References Regarding ADS-B, Multilateration and Other Surveillance Techniques”. In: *ICAO/FAA Workshop on ADS-B and Multilateration Implementation*. Sept. 2011.
- [137] Stefan Brands and David Chaum. “Distance-Bounding Protocols”. In: *Advances in Cryptology — EUROCRYPT ’93: Workshop on the Theory and Application of Cryptographic Techniques Lofthus, Norway, May 23–27, 1993 Proceedings*. Ed. by Tor Helleseeth. Berlin, Heidelberg: Springer, 1994, pp. 344–359.
- [138] Jerry T. Chiang, Jason J. Haas, and Yih-Chun Hu. “Secure and precise location verification using distance bounding and simultaneous multilateration”. In: *Proceedings of the second ACM conference on Wireless network security (WiSec)*. New York, New York, USA: ACM, 2009, pp. 181–192.
- [139] Jerry T Chiang, Jason J Haas, Jihyuk Choi, and Yih-Chun Hu. “Secure location verification using simultaneous multilateration”. In: *IEEE Transactions on Wireless Communications* 11.2 (2012), pp. 584–591.
- [140] Joo-Han Song, Victor WS Wong, and Vincent CM Leung. “Secure location verification for vehicular ad-hoc networks”. In: *IEEE Global Telecommunications Conference (GLOBECOM)*. IEEE. 2008, pp. 1–5.

- [141] Aanjan Ranganathan, Nils Ole Tippenhauer, Boris Škorić, Dave Singel, and Srdjan Čapkun. “Design and Implementation of a Terrorist Fraud Resilient Distance Bounding System”. In: *Proceedings of the 17th European Conference on Research in Computer Security (ESORICS)*. 2012, pp. 415–432.
- [142] Nils Ole Tippenhauer and Srdjan Čapkun. “ID-Based Secure Distance Bounding and Localization”. In: *Proceedings of the 14th European Conference on Research in Computer Security (ESORICS)*. 2009, pp. 621–636.
- [143] Rudolph Emil Kalman. “A new approach to linear filtering and prediction problems”. In: *Journal of basic Engineering* 82.1 (1960), pp. 35–45.
- [144] Dieter Fox, Jeffrey Hightower, Lin Liao, Dirk Schulz, and Gaetano Borriello. “Bayesian filtering for location estimation”. In: *IEEE pervasive computing* 2.3 (2003), pp. 24–33.
- [145] Jimmy Krozel, Dominick Andrisani, Mohammad A Ayoubi, Takayuki Hoshizaki, and Chris Schwalm. “Aircraft ADS-B Data Integrity Check”. In: *AIAA 4th Aviation Technology, Integration and Operations (ATIO) Forum*. 2004, pp. 1–11.
- [146] Eric Chan-Tin, Victor Heorhiadi, Nicholas Hopper, and Yongdae Kim. “The Frog-Boiling Attack: Limitations of Secure Network Coordinate Systems”. In: *ACM Transactions on Information and System Security (TISSEC)* 14.3 (2011), p. 27.
- [147] Young Hwan Chang, Qie Hu, and Claire J Tomlin. “Secure Estimation based Kalman Filter for Cyber-Physical Systems against Adversarial Attacks”. In: *arXiv preprint arXiv:1512.03853* (2015).
- [148] Brandon Kovell, Benjamin Mellish, Thomas Newman, and Olusola Kajopaiye. *Comparative analysis of ADS-B verification techniques*. Tech. rep. The University of Colorado, Boulder, May 2012.
- [149] Tim Leinmuller, Elmar Schoch, and Frank Kargl. “Position verification approaches for vehicular ad-hoc networks”. In: *IEEE Wireless Communications* 13.5 (Oct. 2006), pp. 16–21.
- [150] Kai Zeng, Kannan Govindan, and Prasant Mohapatra. “Non-Cryptographic Authentication and Identification in Wireless Networks”. In: *IEEE Wireless Communications* (2010), pp. 1–8.
- [151] Suman Jana and Sneha Kumar Kasera. “On Fast and Accurate Detection of Unauthorized Wireless Access Points Using Clock Skews”. In: *IEEE Transactions on Mobile Computing* 9.3 (Mar. 2010), pp. 449–462.
- [152] Suhas Mathur, Wade Trappe, Narayan Mandayam, Chunxuan Ye, and Alex Reznik. “Radio-telepathy: Extracting a Secret Key from an Unauthenticated Wireless Channel”. In: *Proceedings of the 14th ACM international conference on Mobile computing and networking (MobiCom)*. 2008, pp. 128–139.
- [153] Junxing Zhang, Sneha K. Kasera, and Neal Patwari. “Mobility Assisted Secret Key Generation Using Wireless Link Signatures”. In: *Processdings of the 29th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. IEEE, Mar. 2010, pp. 1–5.

- [154] Qian Wang, Hai Su, Kui Ren, and Kwangjo Kim. “Fast and scalable secret key generation exploiting channel phase randomness in wireless networks”. In: *Proceedings of the 30th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. IEEE, Apr. 2011, pp. 1422–1430.
- [155] Jie Xiong and Kyle Jamieson. “SecureAngle: Improving Wireless Security Using Angle-of-Arrival Information”. In: *Proceedings of the Ninth ACM SIGCOMM Workshop on Hot Topics in Networks*. 2010, p. 11.
- [156] Di Qiu, Dave De Lorenzo, Sherman Lo, Dan Boneh, and Per Enge. “Physical Pseudo Random Function in Radio Frequency Sources for Security”. In: *Proceedings of the 2009 International Technical Meeting of The Institute of Navigation (ION ITM)*. 2009, pp. 26–28.
- [157] Christine Laurendeau and Michel Barbeau. “Insider attack attribution using signal strength-based hyperbolic location estimation”. In: *Security and Communication Networks* 1.4 (2008), pp. 337–349.
- [158] Christine Laurendeau and Michel Barbeau. “Probabilistic localization and tracking of malicious insiders using hyperbolic position bounding in vehicular networks”. In: *EURASIP Journal on Wireless Communications and Networking* 2009.1 (Feb. 2009), 2:1–2:13.
- [159] Marc Viggiano, Edward Valovage, Ken Samuelson, Dana Hall, et al. *Secure ADS-B authentication system and method*. US Patent 7,730,307. June 2010.
- [160] Leonard Schuchman. *Automatic Dependent Surveillance System Secure (ADS-S)*. US Patent 7,876,259. Jan. 2011.
- [161] Ziliang Feng, Weijun Pan, and Yang Wang. “A data authentication solution of ADS-B system based on X.509 certificate”. In: *27th International Congress of the Aeronautical Sciences (ICAS)*. 2010, pp. 1–6.
- [162] Maxim Raya and Jean-Pierre Hubaux. “The security of vehicular ad hoc networks”. In: *Proceedings of the third ACM workshop on Security of ad hoc and sensor networks (SASN)*. New York, New York, USA: ACM, 2005, pp. 1–11.
- [163] Maxim Raya and Jean-Pierre Hubaux. In: *Journal of Computer Security* 15.1 (2007), pp. 39–68.
- [164] Adrian Perrig, Ran Canetti, J. D. Tygar, and Dawn Song. “Efficient Authentication and Signing of Multicast Streams over Lossy Channels”. In: *IEEE Symposium on Security and Privacy (S&P)*. 2000, pp. 56–73.
- [165] Mohamed Hamdy Eldefrawy, Muhammad Khurram Khan, Khaled Alghathbar, and Eun-Suk Cho. “Broadcast authentication for wireless sensor networks using nested hashing and the Chinese remainder theorem.” In: *Sensors* 10.9 (Jan. 2010), pp. 8683–95.
- [166] Yih-Chun Hu and Kenneth P Laberteaux. “Strong VANET Security on a Budget”. In: *Proceedings of Workshop on Embedded Security in Cars (ESCAR)*. Vol. 06. 2006, pp. 1–9.
- [167] A Smith, R Cassell, T Breen, R Hulstrom, and C Evers. “Methods to Provide System-wide ADS-B Back-Up, Validation and Security”. In: *25th IEEE/AIAA Digital Avionics Systems Conference (DASC)*. 2006, pp. 1–7.

- [168] Leon Purton, Hussein Abbass, and Sameer Alam. “Identification of ADS-B System Vulnerabilities and Threats”. In: *Proceedings of the Australian Transport Research Forum*. Canberra, Australia, Oct. 2010, pp. 1–16.
- [169] Jerry Johnson, Holger Neufeldt, and Jeff Beyer. “Wide area multilateration and ADS-B proves resilient in Afghanistan”. In: *IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS)*. 2012.
- [170] Regina Kaune, Christian Steffes, Sven Rau, Wolfgang Konle, and Juergen Pagel. “Wide area multilateration using ADS-B transponder signals”. In: *15th International Conference on Information Fusion (FUSION)*. IEEE. 2012, pp. 727–734.
- [171] Paul Thomas. “North sea helicopter ADS-B/MLat pilot project findings”. In: *Tyrrhenian International Workshop on Digital Communications-Enhanced Surveillance of Aircraft and Vehicles (TIWDC/ESAV)*. 2011, pp. 53–58.
- [172] Anastasios Daskalakis and Patrick Martone. “A technical assessment of ADS-B and multilateration technology in the Gulf of Mexico”. In: *Proceedings of the 2003 IEEE Radar Conference*. IEEE, 2003, pp. 370–378.
- [173] Gaspare Galati, Mauro Leonardi, Pierfrancesco Magarò, and Valerio Paciucci. “Wide area surveillance using SSR mode S multilateration: advantages and limitations”. In: *European Radar Conference (EURAD)*. 2005, pp. 225–229.
- [174] Greg Welch and Gary Bishop. “An introduction to the Kalman filter”. In: *Proceedings of the Siggraph Course* (2001).
- [175] Gorkem Kar, Hossen Mustafa, Yan Wang, Yingying Chen, Wen Yuan Xu, Marco Gruteser, and Tam Vu. “Detection of on-road vehicles emanating GPS interference”. In: *Proceedings of the 21st ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM. 2014, pp. 621–632.
- [176] MR Mosavi and H Azami. “Applying Neural Network Ensembles for Clustering of GPS Satellites”. In: *International Journal of Geoinformatics* 7.3 (2011), pp. 7–14.
- [177] Wint Yi Poe. “Design Problems in Large-Scale, Time-Sensitive WSNs”. PhD thesis. TU Kaiserslautern, Germany, 2013.
- [178] Nathan J Gomes, Paulo P Monteiro, and Atilio Gameiro. *Next generation wireless communications using radio over fiber*. Chichester, UK: John Wiley & Sons, 2012.
- [179] Hui Liu, Houshang Darabi, Pat Banerjee, and Jing Liu. “Survey of wireless indoor positioning techniques and systems”. In: *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 37.6 (2007), pp. 1067–1080.
- [180] Paramvir Bahl and Venkata N Padmanabhan. “RADAR: An in-building RF-based user location and tracking system”. In: *Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. Vol. 2. IEEE, 2000, pp. 775–784.
- [181] Azat Rozyyev, Halabi Hasbullah, and Fazli Subhan. “Combined K-Nearest Neighbors and Fuzzy Logic Indoor Localization Technique for Wireless Sensor Network”. In: *Research Journal of Information Technology* 4.4 (2012), pp. 155–165.

- [182] Ali Jafarnia-Jahromi, Ali Broumandan, John Nielsen, and Gérard Lachapelle. “GPS vulnerability to spoofing threats and a review of antispooofing techniques”. In: *International Journal of Navigation and Observation* 2012 (2012), pp. 1–16.
- [183] Federal Aviation Administration. *Limiting aircraft data displayed via Aircraft Situation Display to Industry (ASDI)*. Accessed Jun 2016. 2016. URL: <https://www.fly.faa.gov/ASDI/asdi.html>.
- [184] International Civil Aviation Organization. *Guidance material on advice to military authorities regarding ADS-B data sharing*. Tech. rep. 2012.
- [185] David Gloven and David Voreacos. *Dream Insider Informant Led FBI From Galleon to SAC*. Accessed Jun 2016. 2012. URL: goo.gl/Mf6hbJ.
- [186] Busyairah Syd Ali, Arnab Majumdar, Washington Y Ochieng, and Wolfgang Schuster. “ADS-B: The Case for London Terminal Manoeuvring Area (LTMA)”. In: *Tenth USA & Europe Air Traffic Management Research and Development Seminar (ATM2013)*. 2013, pp. 1–10.
- [187] Krishna Sampigethaya, S Taylor, and Radha Poovendran. “Flight privacy in the NextGen: Challenges and opportunities”. In: *IEEE Integrated Communications, Navigation and Surveillance Conference (ICNS)*. IEEE. 2013, pp. 1–15.
- [188] Robert Steele. “Open source intelligence”. In: *Handbook of Intelligence Studies*. Ed. by Loch K Johnson. New York, NY: Routledge, 2007, pp. 129–147.
- [189] Eve Curie. *Madame Curie: A Biography*. Da Capo series in science. New York, NY: Perseus Books Group, 2001.