Measuring Operational Realities of Security and Privacy for Deployed Avionics

Matthew Smith

St Catherine's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2018

In memory of
ARTHUR JOHN SMITH
1925-2018

Acknowledgements

This thesis would not have been possible without the help of many people. My work has been supported by the Centre for Doctoral Training in Cyber Security, and funded in various parts by EPSRC, armasuisse, the Department of Computer Science at Oxford and St Catherine's College.

Many thanks go to my supervisor, Ivan Martinovic, and to my colleague Martin Strohmeier without whom I probably would not have even come across aviation security research. As well as the day-to-day guidance, they have allowed me to spend four years exploring a fascinating topic. Further thanks go to Martin for being an excellent co-author on many occasions. I would like to thank Andrew Simpson and Kasper Rasmussen for acting as Transfer and Confirmation assessors providing invaluable feedback. Similarly, I thank Chris Johnson and Andrew Martin for acting as assessors for DPhil, both of whom provided highly valuable guidance and feedback.

A large part of my research was carried out with a brilliant international team of Daniel Moser, Matthias Schäfer, Vincent Lenders and the wider OpenSky community. Over the course of my DPhil, they have provided expertise and helped me to develop ideas into papers.

Although a DPhil is a solo endeavour, I have been lucky to be part of a brilliant research group. On top of being surrounded by others doing excellent work, the group has given me many good times and has always felt like a great place to be. Thanks also go to David Hobbs and Maureen York for always being on hand to resolve problems, CDT or not.

Outside of my work, I have been provided welcome distraction by many friends. At St Catherine's, I was lucky enough to be part of an active MCR and spend (probably too much) time as a member of the Boat Club. The people I met through the Boat Club have positively defined a significant part of my time in Oxford. Beyond the University bubble, many thanks go to my cycling club, the Cowley Road Condors, and of course my friends from before I started my DPhil. I would especially like to thank Chris Chamberlain for being there throughout.

I would like to thank my family for all the love and support along the way. Were it not for the many years of help and encouragement from my Mum and Dad, and from afar by my Nan, I would not be writing this. Finally, to Ceyda, who I was lucky enough to meet during my time at Oxford. Thank you—you made even the hardest days of writing a thesis much brighter with your love and happiness.

Abstract

Despite being renowned as an exceptionally safety-conscious industry, aviation has been slow to address the cyber security threat. A critical point has been reached whereby systems which were designed many years ago are in wide use, but lack meaningful security measures. Meanwhile, it has become easy to acquire and use tools which enable potential attackers to listen to—and even tamper with—these systems. Like many safety-critical industries, aviation lacks the ability to rapidly redeploy systems. This creates a situation where known-vulnerable systems must be kept and, even worse, heavily relied upon.

The work in this thesis focusses on two topics: analysis of a well-established and heavily used general-purpose avionic communication system, and a first look at a method to analyse and prepare for attacks caused by a lack of security measures on avionics.

For the former, we focus on the Aircraft Communications Addressing and Reporting System (ACARS). We show that it has very few deployed security solutions, and the instances such solutions have been used are weak. As a consequence, we demonstrate the impact of the lack of meaningful confidentiality protection for non-commercial aviation actors—military, government and business aircraft. We show that even when efforts are made to protect privacy elsewhere, they stand a significant chance of leaking data via normal ACARS usage.

Moving to the cockpit, the second topic attempts to begin to address one of the current unknowns in the area of aviation cyber security—how attacks on avionic systems might affect the way the aircraft is flown. In this, we used a flight simulator to create the cockpit-based effects of attacks on three important systems. Using these, we created scenarios in which the aircraft was under attack and invited 30 Airbus A320 pilots to take part.

Whilst the current state of security and privacy in aviation is far from ideal, we believe that methods to provide security in the near- and long-term are achievable. Privacy in the near-term is somewhat harder, but steps towards it in the longer term are underway.

Contents

List of Figures

List of Tables

List of Abbreviations

AAC . . . . . Airline Administrative Control.

AAL . . . . . Above Aerodrome Level.

ACARS . . . . Aircraft Communications Addressing and Reporting System.

ACAS . . . . . Airborne Collision Avoidance System.

ADF . . . . . Automatic Direction Finder.

ADS-B . . . . Automatic Dependent Surveillance Broadcast.

AGL . . . . . Above Ground Level.

AMS . . . . . ACARS Message Security.

ANS . . . . . Air Navigation Services.

ANSP . . . . . Air Navigation Service Provider.

AOA . . . . . ACARS over AVLC.

AOC . . . . . Airline Operational Control.

APT . . . . . Advanced Persistent Threat.

ASDI . . . . . Aircraft Situation Display to Industry.

ATC . . . . . Air Traffic Control.

ATCO . . . . . Air Traffic Control Operator.

ATIS . . . . . Aerodrome Terminal Information Service.

ATM . . . . . Air Traffic Management.

ATZ . . . . . Aerodrome Traffic Zone.

AVLC . . . . . Aviation VHF Data Link Control.

BARR . . . . . Blocked Aircraft Registration Request.

CAA . . . . . Civil Aviation Authority.

CFIT . . . . . Controlled Flight into Terrain.

CMU . . . . . Communication Management Unit.

CNS . . . . . Communications, Navigation and Surveillance.

CPA . . . . . Closest Point of Approach.

CPDLC . . . . Controller-Pilot Data Link Communications.

CPNI . . . . . Centre for the Protection of National Infrastructure.

CSMA . . . . . Carrier-Sense Multiple Access.

CTA . . . . . Control Area.

CTR . . . . . Control Zone.

D-ATIS . . . . Digital Aerodrome Terminal Information Service.

DME . . . . . Distance Measuring Equipment.

EGPWS . . . Enhanced Ground Proximity Warning System.

ETA . . . . . Estimated Time of Arrival.

EU . . . . . European Union.

FAA . . . . . Federal Aviation Administration.

FAF . . . . . Final Approach Fix.

FANS . . . . . Future Air Navigation System.

FIR . . . . . Flight Information Region.

FMCW . . . . Frequency-Modulated Continuous Wave.

FSK . . . . . Frequency Shift Keying.

GDPR . . . . . General Data Protection Regulation.

GPWS . . . . Ground Proximity Warning System.

GS . . . . . Glideslope.

HF . . . . . High Frequency.

HFDL . . . . . High Frequency Data Link.

IAF . . . . . Initial Approach Fix.

ICAO . . . . . International Civil Aviation Organization.

IFR . . . . . Instrument Flight Rules.

ILS . . . . . Instrument Landing System.

IMC . . . . . Instrument Meteorological Conditions.

LOC . . . . . Localiser.

METAR . . . Meteorological Terminal Air Report.

NBAA . . . . . National Business Aviation Association.

NOTAM . . . Notice to Airmen.

OOOI . . . . . Out-Off-On-In.
PACARS . . . Protected ACARS.
PAPI . . . . . Precision Approach Path Indicator.
PIREP . . . . Pilot Report.
POA . . . . . Plain Old ACARS.
PSK . . . . . Phase Shift Keying.
RA . . . . . Resolution Advisory.
RNAV . . . . . Area Navigation.
RSS . . . . . Received Signal Strength.
RX . . . . . Reception.
SATCOM . . . Satellite Communications.
SBD . . . . . Short Burst Data.
SDR . . . . . Software Defined Radio.
SHF . . . . . Super High Frequency.
SID . . . . . Standard Instrument Departure.
SRA . . . . . Surveillance Radar Approach.
STAR . . . . . Standard Instrument Arrival.
TA . . . . . Traffic Advisory.
TAF . . . . . Terminal Aerodrome Forecast.
TAWS . . . . . Terrain Avoidance and Warning System.
TCAS . . . . . Traffic Collision Avoidance System.
TFMS . . . . . Traffic Flow Management System.
TOC . . . . . Top of Climb.
TX . . . . . Transmission.
UHF . . . . . Ultra High Frequency.
UIR . . . . . Upper Information Region.
USRP . . . . . Universal Software Radio Peripheral.
VFR . . . . . Visual Flight Rules.
VDL . . . . . VHF Data Link Mode 2.
VHF . . . . . Very High Frequency.
VMC . . . . . Visual Meteorological Conditions.
VNAV . . . . . Vertical Navigation.
VOR . . . . . VHF Omnidirectional Range.

A bicycle ride around the world begins with a single pedal stroke.

— Scott Stoll

1 Introduction

Contents

Aviation is oft-held as a triumph in safe systems; 2017 saw the safest year on record for commercial aviation, with just 10 fatal accidents causing the loss of 79 lives, compared to 16 accidents and 303 lives in 2016 [1]. As an industry, it is known for its ability to react to incidents by improving safety and producing methods to reduce the chance of recurrence. One way it does this is through the introduction of more complex avionics—electronic systems specific to aviation—to help pilots and air traffic controllers alike avoid unsafe situations. Safety is of utmost importance in avionics, with failure or unexpected behaviour being unacceptable [2].

Modernisation is currently a major focus in aviation, with the domain as a whole looking to leverage new and existing technology in order to improve operations. This effort is driven by three aims: to lower costs, increase safety and reduce environmental impact.

In Europe, modernisation is managed by the *Single European Sky Air Traffic Management (ATM) Joint Undertaking*, or SESAR JU, and is led by both the

European ATM organisation Eurocontrol and the European Commission [3]. In the US, the programme is called NextGen and is led and managed by the Federal Aviation Administration (FAA) [4]. Internationally, this is coordinated by the Global Air Navigation Plan, maintained by the International Civil Aviation Organization (ICAO) [5]. In the case of SESAR JU, the project is aiming to achieve a range of savings per flight by 2035, including:

- 4-8 minutes less fuel burn,
- 0.79-1.6 tonnes less CO_2,
- 1-3 minutes fewer departure delay, and
- Halving the cost of air navigation services (ANS) [6].

On top of this, the program aims to improve safety by 'factor 10', i.e. no increase in accidents despite the increase in air traffic. Both existing and future avionics will enable these changes.

A Changing Threat

Physical security is a well-known challenge for the industry. Between lasers shone into the cockpit, passengers without tickets illicitly boarding aircraft, stowaways and a number of terrorist attacks, it has responded by adapting procedures to improve robustness [7–10]. In the meantime, cyber security has rapidly become a problem for all businesses; aviation has seen some of the biggest data breaches to date, with significant numbers of British Airways and Cathay Pacific passengers affected [11, 12].

Whilst aviation has been effective in producing safe systems, it has only recently begun to address the challenge of securing them. This has resulted in many of today's avionics not having effective security mechanisms, if any at all. Partly due to the extensive safety requirements and certification process, system development takes a long time; in the extreme case of designing a new aircraft, this can take up to five years [13]. As a result, even though some technologies at the forefront

of modernisation are insecure, they cannot be quickly replaced. Since many of these systems are tied into the safe operation of airspace, this is leading to a situation where systems without adequate security mechanisms are forming the basis of future aviation infrastructure.

One major challenge for security in aviation relates to the requirement for mobility. Aircraft travel quickly over long distances, but need up-to-date information. To achieve this, many avionics communicate over long distances using unprotected wireless links. In the past, producing hardware to receive, demodulate and decode these communications as a third party was the work of a specialist, and likely to be expensive. This prevented many potential attackers from being a threat. However, the advent of affordable software-defined radios (SDRs) has shifted this dynamic.

SDRs are a type of general radio hardware which can be reconfigured through software instead of needing hardware changes. This makes them versatile and allows users to experiment more easily. Once expensive scientific hardware, SDRs are now readily available. The most stark example is the RTL-SDR, a £10 digital television USB stick which can be repurposed as an SDR capable of reception [14]. The hobbyist community has thrived in this respect, producing a wide range of open source, freely available software capable of receiving aviation signals, so far with the focus being on collecting positional data. Through this, less skilled users can download and use tools to collect this data with little knowledge themselves. However, the same technology is the basis for a new threat which both exploits the data received and uses malicious signals to interfere with avionic systems.

Aviation Security Research

The rise of SDRs have spurred research into the security of avionic systems, initially through work in the hacking community by Costin, Haines and Teso [15–17]. Often, the aviation industry was quick to decry the work; in the case of Costin's research, the FAA did not engage, insisting that protections were already in place—though it is unclear precisely how these protections work [18]. The attitude of the industry towards security research has improved somewhat in recent years, to the extent

that the US Department for Homeland Security announced work culminating in remote access to the systems of a retired Boeing 757 [19]. However, this is still an outlier, with much of the ongoing security research kept internal.

So far, academic research has focussed on surveillance technologies. One such technology is Automatic Dependent Surveillance-Broadcast (ADS-B), which forms a major part of airspace modernisation by enabling regular, accurate positional updates from aircraft. Works by Schäfer, Strohmeier and McCallie have highlighted a range of ways in which an attacker can interfere with the system [20–22].

Furthermore, the threat is not just active. Passive data collection has become commonplace in aviation through flight trackers such as Flightradar24, FlightAware or the OpenSky Network [23–25]. Some aircraft operators have sought privacy for a long time and are trying to counteract its perceived loss due to flight trackers [26]. Recent work has shown that in some cases, this data can be used to infer confidential business activities and inter-state summits [27]. Outside of the academic sphere, a number of news articles have used flight data extensively. For example, work by Aldhous identifies 'spy' aircraft by their movements, or the Geneva Dictator Alert by Pilet, which tweets when aircraft from authoritarian governments—so labelled by the US Central Intelligence Agency—land at Geneva airport [28, 29].

We are at a point where the true extent of the security—and consequently privacy—problems faced by aviation are unknown. Many avionic systems are yet to be investigated from a security and privacy point of view. The work in this thesis aims to expand upon existing knowledge by looking at the security of yet-uninvestigated communications, navigation and surveillance avionics.

The research questions which motivate this work are as follows:

1. Does the lack of use of security mechanisms on avionic data links affect aircraft operator privacy?

2. In cases where confidentiality protection is used on avionic data links, is it effective and does it provide additional privacy protections?

3. Does the lack of security mechanisms extend to navigation and surveillance systems, and does this enable attacks?

4. Can the impact of attacks on safety-critical systems be mitigated by flight crew training?

To address the first two questions we look at the Aircraft Communications Addressing and Reporting System (ACARS), which provides air-to-ground text-based message exchange worldwide for many types of aircraft. We will assess whether avionic data link usage has grown in such a way that privacy is now a problem, and the ways in which aircraft operators are trying to protect this.

For the third and fourth questions, we investigate representative navigation and surveillance systems, chosen due to their criticality in the safe operation of an aircraft. We focus on collision avoidance, landing and terrain proximity warning systems, all of which comprise important parts of modern aviation. As part of this, we outline theoretical attacks on these systems and use simulation to assess what their impact might be.

1.1 Contributions

In the course of this work, we make a number of contributions which are based on publications, or projects which are intended for publication in due course.

- In Chapter 5, we provide a review on the concept of privacy in aviation for non-commercial aircraft, and the extent of protection it provides. This builds on work in the paper *Undermining Privacy in the Aircraft Communications Addressing and Reporting System (ACARS)* published in the Proceedings of Privacy Enhancing Technologies (PoPETS) 2018 [30], and contributed to Strohmeier's paper, *The Real First Class? Inferring Confidential Corporate Mergers and Government Relations from Air Traffic Communication*, published at IEEE European Symposium on Security and Privacy (EuroS&P) 2018 [27].

- We measure how non-commercial aircraft use data link in such a way as to reveal a significant amount of location data despite efforts to protect privacy otherwise. This work forms Chapter 6 and expands upon the paper *Undermining Privacy in the Aircraft Communications Addressing and Reporting System (ACARS)* published in the Proceedings of Privacy Enhancing Technologies (PoPETS) 2018 [30].

- Regarding data link security, we provide analysis of deployed cryptography on the ACARS link, primarily in use by business aircraft with an apparent privacy requirement. This work forms Chapter 7 and includes research from *Economy Class Crypto: Exploring Weak Cipher Usage in Avionic Communications via ACARS*, published at Financial Cryptography and Data Security 2018 [31].

- Moving to navigation and surveillance systems, Chapter 8 provides detail on a first-look study of using simulation to assess the impact of realistic and feasible attacks on avionics. Specifically, we do this to understand whether flight crew actions mitigate these attacks, to what extent, and whether this approach is valuable for training. This chapter includes work from *Safety vs. Security: Attacking Avionic Systems with Humans in the Loop*, available as a preprint on arXiv [32].

As a methodological contribution, we take a measurement approach to understand how systems are used in practice and by extension what this means for security, privacy and safety. This is useful for domains such as aviation in which systems are in use for decades post-deployment, meaning that actual usage can vary from that which was intended. We believe this provides a framework for future research into security and privacy of avionic systems under a changing threat model and shifting environment. We discuss the approach as a reflection in Chapter 9.

1.2 Outline

The work in this thesis takes the following structure:

- Chapter 2 provides a background on avionic systems relevant to this work, including security research where it exists.

- Chapter 3 considers the attacker models used in the thesis, as well as discussing how technology has changed to enable many types of attacker to be a threat to aviation.

- Chapter 4 describes the workings of ACARS in detail, reviewing its security stance and our methods used to collect data from it. We also provide an overview of our survey on industry opinions of avionic data link security and privacy.

- Chapter 5 presents the current state of privacy in aviation, specifically with regard to how non-commercial aircraft are attempting to protect their movements against the emergence of public flight trackers.

- Chapter 6 covers a measurement study on how usage of the ACARS data link leaks location data for non-commercial aircraft who appear to be privacy sensitive.

- Chapter 7 investigates usage of proprietary encryption by non-commercial aircraft on the ACARS link, focussing on a system used primarily by business aircraft which is readily breakable.

- Chapter 8 moves to consider attacks on critical flight systems and attempts to assess whether crew actions mitigate their effects using flight simulation. In this, we describe the theoretical attacks, evaluate their apparent impacts, and discuss whether simulation is a valuable tool for training against such attacks.

- Chapter 9 concludes the thesis, providing a summary of results, a framework measurement approach to assessing security and privacy, recommendations based on the work and potential next steps.

1.3 Ethical & Legal Considerations

The nature of this work involves researching systems relating to critical infrastructure which could have severe consequences if attacked. Whilst we consider publishing this work an important part of the process of resolving security issues, we have taken care to not reveal details which would enable threats to carry out attacks who otherwise would not have been able to.

From a legal point of view, we have been careful to avoid collecting or transmitting data in such a way that would contravene local laws, and have ensured that we have not provided tools which attackers could use to achieve this either.

Throughout the work, we upheld strong ethical conduct, ensuring to engage with the industry and regulators where possible to share our work. This has involved following responsible disclosure processes when necessary, and is identified in the text. Furthermore, where appropriate we have sought ethical review and approval from our local ethics committee, particularly when handling personally identifiable information or conducting studies with human participants. We include reference numbers for ethics approvals where relevant.

We approached both the National Business Aircraft Association (NBAA) and Federal Aviation Administration (FAA) for comment on content relevant to Chapter 6 but at the time of writing have received no reply.

It is by riding a bicycle that you learn the contours of a country best, since you have to sweat up the hills and coast down them.

— Ernest Hemingway

2 Background

Contents

Aviation is a highly specialised field, using a multitude of bespoke concepts and technologies. This chapter provides the necessary background on the field, systems and categories of aircraft in order to contextualise the topics covered in this thesis.

2.1 Key Aviation Concepts

As a domain, aviation is well-defined and regulated for safety. This heavily impacts technology design and usage, especially with regard to the systems researched in this thesis. To provide context, we cover some fundamentals in this section.

2.1.1 Phases of Flight

A flight can be divided into many phases, but we are particularly interested in those between takeoff and landing. We illustrate these in Figure 2.1 and use the common taxonomy as defined by ICAO [33].

- **Takeoff** is from initially applying power to the aircraft on the departure runway to the point of leaving the runway.
- **Initial Climb** is the first part of the climb after takeoff, before engine power is reduced from takeoff thrust.
- **En-route** is the majority of the airborne portion of flight, further including:
  - *Climb to cruise*, which is the remainder of the climb up to the determined cruise altitude,
  - *Cruise* at a predetermined altitude, usually for most of the flight,
  - *Descent* to the *Initial Approach Fix* (IAF), which marks the start of the approach,
  - *Holding* wherein a route is followed to keep the aircraft in a particular place (known as a *holding pattern*), before further instructions are given.
- **Approach** is from the end of the descent, or the exit of a holding pattern, to the aircraft being about to touch down on the runway (namely the *flare*), of which the following are of particular interest:
  - *Initial Approach* is from the IAF to the *Final Approach Fix* (FAF) which is the final checkpoint before landing,

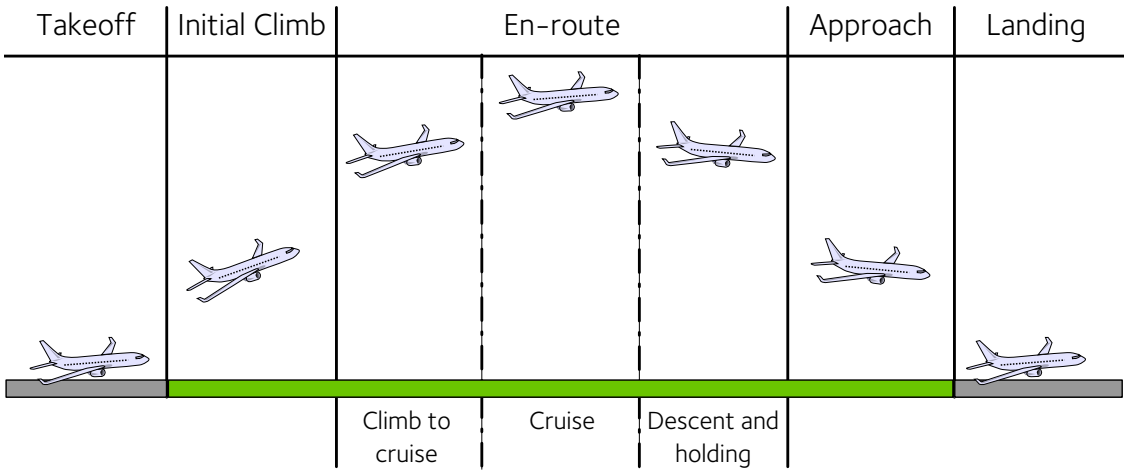

Figure 2.1: Illustration of the phases of flight as defined by ICAO. Phases are labelled at the top of the diagram.

  - *Final Approach* is from FAF to the landing flare,
  - *Missed Approach* is where crew opt to abort the approach, which can also be referred to as a *go-around*.

- **Landing** is from the flare of the aircraft to the aircraft leaving the runway, stopping on the runway, or having to abort the landing and performing a *touch-and-go* landing.

2.1.2 Air Traffic Management

Airspace is a constantly evolving system affected by factors like weather, activity on the ground and passengers. Central coordination in the form of Air Traffic Management (ATM) exists to manage this and operate at a high level of safety.

Airspace worldwide is divided into Flight Information Regions (FIRs), with Upper Information Regions (UIRs) above them [35]. It is further divided into *controlled* and *uncontrolled* airspace. Controlled airspace is where Air Traffic Control (ATC) services are used to manage traffic, performing tasks such as maintaining aircraft separation and clearing aircraft to fly certain routes or altitudes—aircraft flying here must obey ATC instructions.

FIRs are split into airways (providing routes across airspace), Control Areas (CTA), Terminal Control Area (TCA, also known as Terminal Manoeuvring Area,

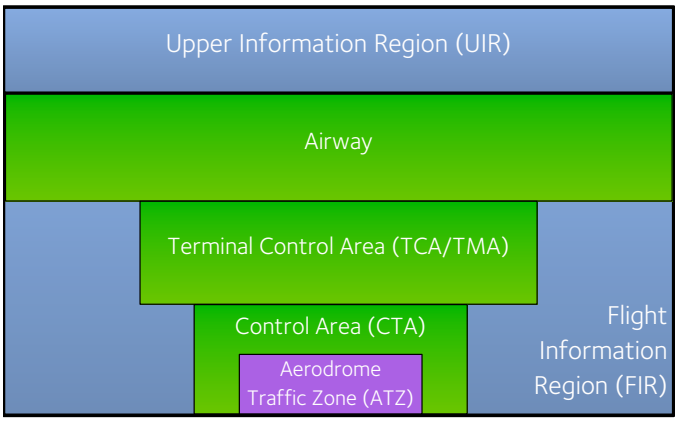

Figure 2.2: Representation of airspace sections around an airport, based on the diagram in [34].

TMA) and Control Zones (CTR). Around an airport, a small area called the Aerodrome Traffic Zone (ATZ) will be defined.[1] Uncontrolled airspace forms the remainder of the available area, and some ATC services may be provided, such as aircraft proximity warnings and weather information [36]. A representation of way airspace is split can be seen in Figure 2.2.

Zones in both FIRs and UIRs of airspace are managed by individual controllers. Whilst in a given zone, aircraft will communicate with the controller, and when they move to a new zone they perform a handoff to the next controller. This will typically involve changing a radio frequency.

2.1.3 Airline Operational and Administrative Control

Particularly for larger fleets, management and maintenance requires regular reporting and logging as well as mid-flight updates on weather or preferred routes. Airline Operational Control (AOC) and Airline Administrative Control (AAC) covers these tasks, which rely on communications between an aircraft and the ground. On top of communication with the airline, aircraft may also directly communicate with other partners, such as engine providers or airport ground staff. Importantly, AOC and AAC communication objectives are defined by the organisations themselves. This means that whilst the underlying technology is the same across aircraft, the method and means of communications can vary significantly.

[1] We do not elaborate on each type here, however full explanations can be found in [34].

2.1.4 Flight Regulations

An aircraft can be flown under two different sets of rules: Visual Flight Rules (VFR) and Instrument Flight Rules (IFR). Weather conditions and categories of airspace place conditions on which sets of rules can be used.

Visual Flight Rules. Simplistically, this could be defined as flying when navigation and control can be done with a visual reference. The weather conditions under which VFR can be flown are referred to as Visual Meteorological Conditions (VMC) [37]. Naturally, this is not possible in periods of poor visibility or heavy cloud. Only a small portion of the airspace is available for aircraft and pilots who are capable of flying solely under VFR.

Instrument Flight Rules. Rather than relying on visual references, IFR instead uses information provided by instruments so that flight can continue when no visual reference is present [38]. This means that flying in poor weather or at night is possible. Most categories of airspace require some level of IFR capability. Importantly, this type of flying requires pilots to trust their instruments and identify when they are malfunctioning.

Portions of the phases of flight defined above are mapped to instrument procedures. Standard Instrument Departures (SIDs) cover takeoff, initial climb and part of the climb to cruise, defining how aircraft should leave a runway and the route it should take to a given point. Standard Instrument Arrivals (STARs) cover part of the descent, holding, initial and final approach into a runway at a given airport. The SID and STAR used by an aircraft will depend on the route it is taking and the runways in use at a given airport.

2.2 Aircraft Categories

Broadly, aircraft can be divided into different categories based on their stakeholders or operators. Doing so allows us to reason about groups of users which have

similar privacy or security requirements. We divide them into commercial, business, military, state (including government) and hobbyists.

2.2.1 Commercial

Most public interaction with aviation is through commercial aircraft, which carries passengers and freight between airports. In the UK, commercial aviation moves a significant number of passengers; from January to October 2018, London Heathrow catered for an average of 6.7 million passengers per month [39]. Aircraft in this category usually have over 20 seats and are operated by commercial airlines. Typical short-haul aircraft are the Airbus A320 or the Boeing 737; long-haul aircraft examples include the Airbus A380 or the Boeing 777.

2.2.2 Business

Business stakeholders typically fly jets capable of 4–20 passengers. Gulfstream's G-range or Bombardier's Learjet and Challenger aircraft are popular choices. There are also business airliners based on commercial airframes produced by Boeing and Airbus, which in their VIP and corporate versions constitute the high-end of the market capable of carrying 100+ passengers. Business flights can either be commercial on-demand services (e.g. aircraft chartering) or, if the aircraft is owned by the operator and used without hire, can be counted under general aviation. They are used to transport personnel to meetings, conferences or other gatherings.

2.2.3 Military

We classify any aircraft operated by an armed force as military aircraft. Military stakeholders operate differently to commercial aircraft. A typically military fleet will consist of some civilian aircraft adapted for military purpose or used for transport, and a set of military-specific aircraft. These aircraft are able to operate in ways civilian aircraft cannot; they use military-specific communications systems and are permitted to turn off some systems that other categories aside from state are not [40]. Military aircraft types that use civilian technologies range from modified

airliners and business jets to tankers and multi-role transport aircraft (e.g. the Boeing C17 Globemaster III), but not fighter jets/combat aircraft.

2.2.4 State

Air transport for state officials differs between countries. In some states, the task falls to the flag-carrier airline, in others to the military, and many heads of state own private aircraft. For example in the UK, the Royal Family and government use state-owned, military-operated aircraft [41]. Regardless of the operator, these are often typical business aircraft, from small Gulfstream or Bombardier jets to larger Airbus or Boeing jets for bigger delegations, and tend to operate in similar ways to civilian aircraft.

2.2.5 Hobbyists & Unpowered Aircraft

The remaining portion of air traffic can be classified as hobbyists or unpowered. Aircraft in this category can take many forms, from light aircraft like a Cessna Citation, to gliders or hot air balloons. As well as being a range of different types of aircraft, they each have different security and safety systems. For example, aircraft of this size typically have fewer equipment requirements so may not visible to ATC via all surveillance mechanisms. In turn, this means that some aircraft in this group are restricted in the type of airspace in which they can fly. Due to the wide range of equipment in use by these aircraft, we do not consider them further.

2.3 Avionic Systems

We now move to look at systems relevant to the work in this thesis. Commonly, avionic systems are divided into communications, navigation and surveillance (CNS); we look at some examples of each. We also cover a hybrid category, safety nets, in which systems might use components from multiple CNS categories.

2.3.1 Communications

Since ATM is active and constantly changing, having direct communication channels between aircraft and the controllers is very important. This takes two forms: voice-based, and data link based.

Voice

Voice communication is used extensively in aviation—primarily for ATC contact—and is part of the minimum equipment that the aircraft must have to take off [42]. Much of this communication is done over VHF channels in the 118.00–136.975 MHz range, though some operators may use HF or SATCOM voice in remote areas [42]. With the amount of air traffic increasing—especially around major hub airports like London Heathrow or Amsterdam Schipol—VHF voice channels have become incredibly congested. In Europe, this has caused VHF channel spacing to need to be reduced from 25 KHz to 8.3 KHz in Europe [43]. As a result, there is a movement towards using data links instead of voice for a range of applications. Some operators may use VHF voice communications for company tasks, such as to report maintenance issues or identify faults. This is airline-dependent and are an additional service such as ARINCDirect [44].

Security Research. Since voice is so well-established as a communication means it is often seen as the fallback option, for example in emergency situations [45]. This has motivated security research over the years, though deployment of security mechanisms appears limited. The primary approach is watermarking in which aircraft identifiers are computed as a watermark and embedded in a non-speech part of the signal [46–48]. This helps to identify the aircraft despite a noisy channel, and makes it harder to spoof the signal. In [49] a system is proposed to protect voice integrity and provide user authentication through encoding data in the voice signal. Another approach is based on voice biometrics; for example, in [50, 51] both use voice fingerprints and a stress indicator to attempt to identify intruders, but rely on having a fingerprint database of pilots. In [52], the authors extend

upon the watermarking idea and propose to use initial pilot communications with ATC to register a basic voice biometric, against which future transmissions are checked. In terms of privacy, a talk at DEF CON 20 used ATC voice to identify movements of private aircraft [53].

Avionic Data Links

For the purposes of this thesis, we focus on one of the most widely used aviation data links, the Aircraft Communications Addressing and Reporting System (ACARS), which can be carried by a number of subnetworks or data links:

- VDL Mode 1 (i.e. ACARS only, also referred to as Plain Old ACARS, POA),
- VHF Data Link Mode 2, a general purpose VHF link (VDL2),
- Satellite communications (SATCOM),
- High Frequency Data Link (HFDL).

Some of these links—VDL2, SATCOM and HFDL—form the Future Air Navigation Service (FANS) and provide services besides ACARS [54]. These include Controller-Pilot Data Link Communications (CPDLC) and Automatic Dependent Surveillance Contract (ADS-C). Since non-POA components of FANS can also carry ACARS messages, we will discuss ACARS, CPDLC and ADS-C as independent systems.

Aircraft Communications Addressing and Reporting System

A general purpose avionic data link designed in 1978, ACARS has since seen widespread adoption and is now used for many purposes, with messages primarily consisting of free-text [55]. These messages can then be transmitted over POA, SATCOM, HFDL or VDL2 links. Much of the network is managed by ACARS service providers in a similar fashion to mobile cell networks, to whom users pay a fee. We cover ACARS in more technical detail in Chapter 4.

Originally, it was intended to serve as a system to exchange Out, Off, On, In (OOOI) messages, more specifically:

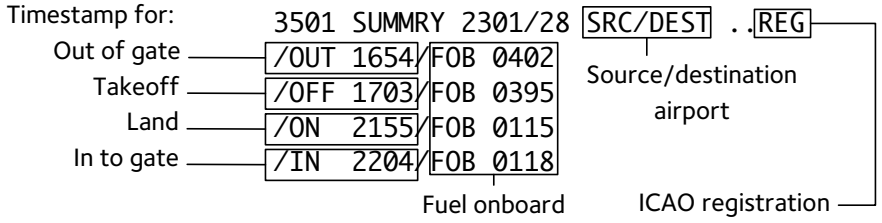

Figure 2.3: Labelled Out-Off-On-In message as seen on the ACARS link.

- *Out* – the aircraft leaves the stand at the departure airport airport,
- *Off* – the aircraft takes off from the departure airport,
- *On* – the aircraft lands at the arrival airport,
- *In* – the aircraft arrives at the gate at the arrival airport.

This allowed airlines to monitor their fleet and time keep crew members, without relying on paper record keeping. An example of this type of message can be seen in Figure 2.3.

Broadly, usage matches to ATC and AOC tasks, with a significant portion of the traffic falling under the latter. We provide some examples matched to phases of flight in Figure 2.4. Example AOC usage might be:

- Transferring flight plans to the cockpit, especially upon changes happening,
- Providing take off configuration data, including fuel, cargo and passenger loading,
- Maintenance tasks including automatic fault reporting or engine health monitoring,
- Passenger-related information sharing, e.g. requesting wheelchairs or connecting flight gates,
- Free-text exchange, used for discussion between crew and ground staff.

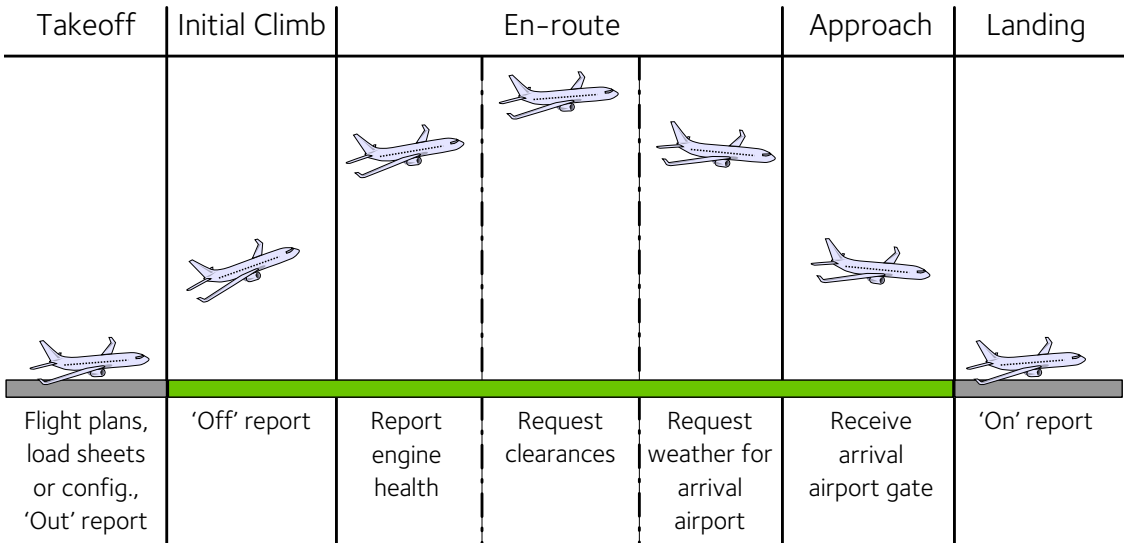

Figure 2.4: Phases of flight matched to example data link message types, provided at the bottom of the diagram.

In terms of ATC, usage is similar to voice communications however in practice, ATC over ACARS is not widely used. One ATC use case which has seen deployment is in arranging oceanic crossing clearances, for example at the Shannon entry point off the western coast of Ireland [56]. We cover ACARS security research in detail in Chapter 4.

Controller-Pilot Data Link Communications

Whilst ACARS does provide some ATC capabilities, it does not have a well-defined structure to handle the range of commands and exchanges needed. Controller-Pilot Data Link Communications (CPDLC) provides such a scheme, allowing crew and controllers to have ATC exchanges over data link. This is particularly useful in congested areas, and has seen deployment in Europe. By 2020, some parts of European airspace will require the use of CPDLC [57].

Although CPDLC is very expressive, voice is currently still used as the preferred option for time-critical communications. This is especially the case since a range of challenges exist in relying on data link ATC; these include incorrect message reception and two controllers issuing differing ATC commands at once, which does not tend to happen with voice-based ATC [58]. Furthermore, since it is carried on the standard, unsecured data links, there is potential for integrity loss and

authenticity issues—this is sufficiently established that one ANSP maintains a list of incidents arising from CPDLC or more generally, FANS [59].

Security Research. Recently, work on CPDLC security has become more active, coinciding with its wider usage. In [60], the authors assess the threat to CPDLC from different types of attacks, and present a range of potential of cryptographic schemes to address these threats. Whilst not specifically CPDLC, work in [61] identifies that VDL2—one of the carriers of CPDLC—is susceptible to a number of attacks, including injection of arbitrary signals.

2.3.2 Navigation

Another major usage of wireless communications in aviation is for navigation, which we now outline. These can provide numerous services such as distance measuring and direction finding.

Non-directional Beacon

One of the more simplistic navigation aids is the non-directional beacon (NDB), which is a ground station transmitting a radio signal with equal strength in all directions [62]. The partner to the NDB onboard an aircraft is the automatic direction finder (ADF). This uses the received direction of the radio signals from the NDB to provide a radio compass [63]. By using two antenna—a loop antenna for signal directionality, and a sense antenna for signal strength—it provides a directional reference for pilots.

An NDB will transmit in the low- to mid-frequency bands (200-1750 KHz), and range varies depending on the location and intended use. For example, NDBs for over-land transit will have a shorter range than those allowing transit over water; typical ranges are within 50 nmi to 400 nmi, with powers between 10 W and 2000 W. In order to identify a specific NDB station, each transmits a two or three letter Morse code identifier.

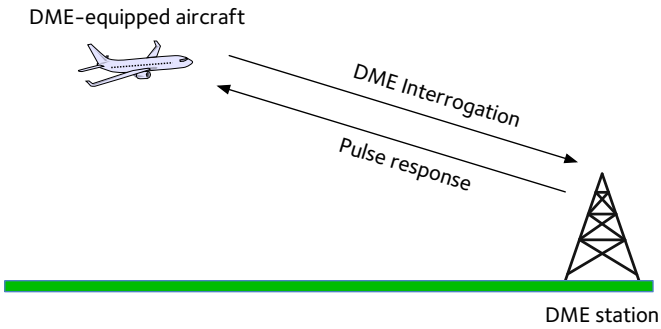

Figure 2.5: Representation of DME call-response system, with the aircraft interrogating the ground station.

VHF Omnidirectional Radio Range

One of the main radio navigation aids is the VHF omnidirectional radio range (VOR). In use since the 1960s, it can be used in a wide range of conditions and for a range of purposes including runway approaches and orientation [64]. As with NDB, VOR infrastructure consists of ground-based *radio beacons*. They transmit in the 108.00–117.95 MHz range, and due to the use of VHF as a carrier, they have a range defined by

$$r = \sqrt{1.5a} \tag{2.1}$$

where r is VHF range in nautical miles (nmi), and a is altitude in feet (ft). Frequencies of VORs are selected in such a way to minimise overlap and interference. They are used widely by all sections of aviation, and can serve as a back up to GPS-based navigation.

Distance Measuring Equipment

Operating in 962–1213 MHz, Distance Measuring Equipment (DME) allows an aircraft to measure effective range between itself and a DME ground station [65]. It does this in a call-response fashion, as indicated in Figure 2.5; the aircraft interrogates on a given frequency, to which the ground station responds with a synchronous answer. Using the round trip time, the aircraft DME transponder can calculate the slant distance between the aircraft and the DME ground station using

$$d = \frac{t_{rtt}}{2c} - t_{delay} \tag{2.2}$$

where t_{rtt} is the round trip time in seconds (s), t_{delay} is the ground station processing delay in seconds (s), d is slant distance in metres (m) and c is the speed of electromagnetic wave transmission through air (metres per second, m/s).

Since DME takes a secondary radar approach, ground stations have a limited capacity before becoming saturated; this is usually around 100 aircraft. Frequencies are chosen so that the chance of interference is minimised. The range of a DME station is calculated as per Equation 2.1, since DME is a line-of-sight system.

Instrument Landing System

Whilst VOR, DME and NDB provide navigation aid in the air, the Instrument Landing System (ILS) helps crew to land the aircraft in a consistent, workload-controlled manner. Composed of vertical and lateral guidance, it provides ground-based navigation from the arrival phase of flight through to just before landing.

Vertical guidance is in the form of the *glideslope*, which provides a slope to a point approximately 300 m beyond the near end of the runway on approach. [2] The signal is transmitted from the side of the runway to avoid antenna in the path of aircraft [66]. The slope will usually be on path to the touchdown zone—nominally the target for wheels down—with the glideslope delivering the aircraft to a point where the final part of the approach can be flown visually. Typically, the slope is 3°, however this is airport and approach dependent.[3] This achieved using overlapping lobes, as shown in Figure 2.6b, and transmitted on an Ultra High Frequency (UHF) carrier in the range 328.60–335.40 MHz [69]. A large upper lobe is modulated at 90 Hz and the lower, smaller lobe is modulated at 150 Hz. The overlapping region is configured to be the required glideslope. Due to the closeness of the lower lobe with the ground, reflections may cause false glideslopes at a steeper angle than the true slope.

[2] This end of the runway as viewed from the aircraft is known as the threshold

[3] Some approaches may be slightly more, or less. Extreme examples exist such as London City Airport at 5,5°, which requires special training [67]. A video of this can be seen in [68].

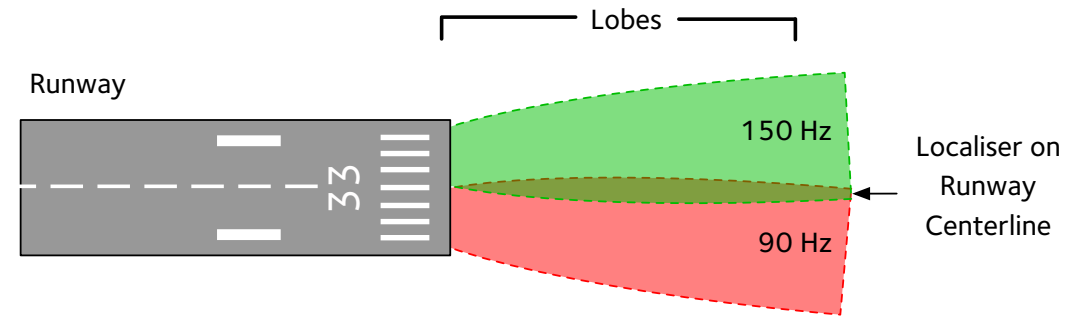

(a) Localiser (LOC) operation, with overlapping lobes indicating the runway centreline.

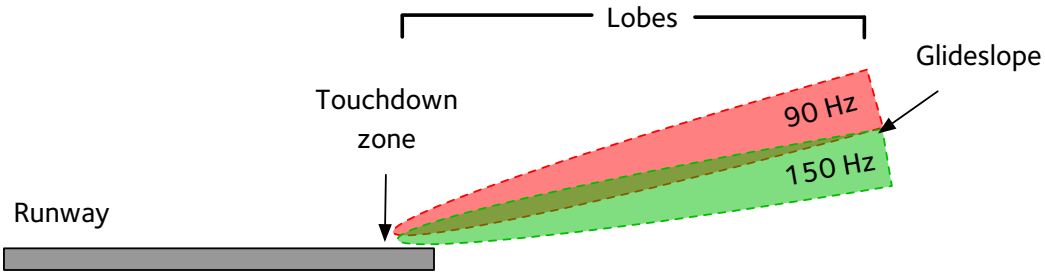

(b) Diagram of Glideslope (GS) operation, with overlapping lobes indicating the nominal glideslope.

Figure 2.6: Illustration of Instrument Landing System (ILS) components, the localiser and glideslope.

Providing lateral guidance, the localiser is calibrated out from the centreline of the runway for up to 20 nmi [70]; we illustrate this in Figure 2.6a. It uses frequencies in the range 108.10–111.95 MHz to carry two lobes, one either side of the centreline [71]. As with the glideslope, the overlapping region of these lobes identifies the approach path within 2.5° of the extension of the centreline. Importantly, the localiser can still be used for approaches if the glideslope fails—however, if the localiser fails, the glideslope cannot be used.

Onboard the aircraft, instruments read out the position of the aircraft on the glideslope and localiser to the crew. This is calculated based on the relative signal strength of each lobe, with the ideal lateral and vertical path being when the aircraft senses equal strength on each lobe [66, 70].

Completing an ILS system is DME, used to provide a distance reference whilst on an ILS approach. This is particularly useful in being able to check altitudes at specific fixes, which provides a cross-reference against the glideslope.

Area Navigation

For approaches or en-route flight using the above systems, ground-based equipment is needed. Area Navigation, or RNAV, allows flight on a direct course rather than between ground-based navigation aids [72]. This approach uses GPS signals for guidance, allowing custom routing rather than on existing corridors between navigation aids.

On top of providing en-route guidance, RNAV can be used for GPS-based approaches [73]. This requires higher accuracy GPS fixes than is needed for en-route applications, as well as a database containing instrument procedures. Some more advanced avionics units can offer Vertical Navigation (VNAV) for RNAV approaches, which provide vertical guidance.

Security Research. RNAV itself has not seen security research, but GPS has been considered extensively. Some examples of this include [74] in which the authors investigate the effects of maritime GPS jamming, identifying a range of effects both on and off shore—as well as noting the effect on people, since they are expecting excellent GPS performance. In [75] an analysis of the technical detail of GPS jamming is carried out, alongside adaptations which could be made to improve jamming resistance. Finally, the authors of [76] perform an assessment of the requirements to successfully spoof GPS signals, finding that whilst spoofing is possible, the locations from which it can be done relative to the target are limited.

2.3.3 Surveillance

So far we have covered some systems that an aircraft uses communicate and navigate. In order for ATC to manage the airspace, they must be able to at least establish aircraft location, but ideally their speed, course, and altitude. This is achieved through surveillance, and forms the backbone of airspace management. This can be split into two types: *primary* and *secondary* surveillance.

Primary Surveillance Radar

Primary surveillance is what might be thought of as 'traditional' radar. It is commonly formed of a rotating dish from which an UHF pulse beam is transmitted and reflected by the aircraft [77]. Measuring the time at which a pulse reflection is received at the radar allows the range to the reflecting object to be calculated, which is approximately

$$d = \frac{t_{rtt}c}{2} + t_d \tag{2.3}$$

where d is the distance between the radar and the object (range, in metres), t_{rtt} is the round trip time from pulse transmission to reflection reception (in seconds), t_d is the processing delay (in seconds), and c is the speed of electromagnetic waves through air (in metres per second). Furthermore, the direction to the object can be determined by monitoring the rotation of the radar dish when a reflection is received.

Whilst this can give precise readings for the range and direction to an object, it does not discriminate between objects. This means it is liable to interference from weather, nearby terrain or birds, though signal processing can help reduce this.

Secondary Surveillance Radar

Rather than relying on signal reflections, secondary surveillance radar (SSR) sees the aircraft respond to signal pulses—called interrogations—making it cooperative [78]. An SSR station interrogates a *transponder* onboard an aircraft, to which the transponder replies. This forms a basic data link and so is more data rich than PSR, and has the added benefit of ensuring that only transponders reply. We provide a representation of SSR in Figure 2.7a. A range of modes of operation exist for SSR, each of which can provide far more information than PSR [79].

Mode A. This mode allows identification of aircraft in range of the radar who respond with a locally allocated identifier known as a *squawk*, limited to 4096 values.

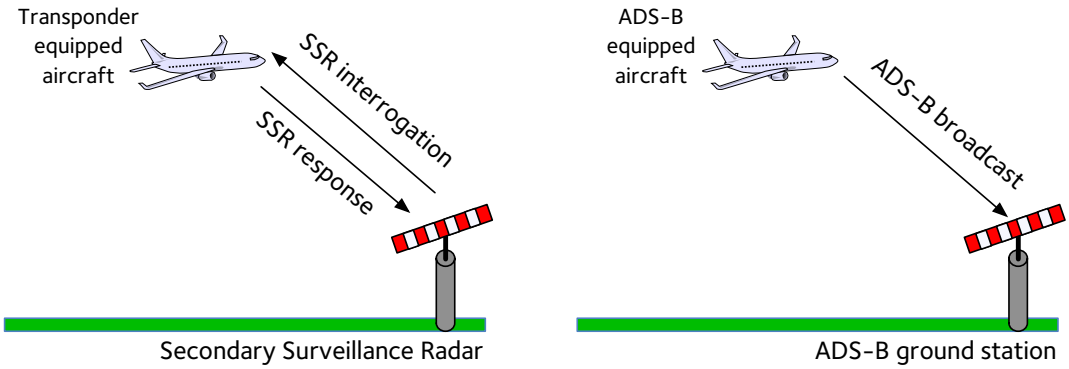

(a) Representation of Mode A/C/S call-response format of interrogation. Aircraft respond to ground-based interrogation from a Secondary Surveillance Radar.

(b) Representation of ADS-B broadcast for collection by ADS-B ground stations. These stations can be a SSR or specialist ADS-B.

Figure 2.7: Illustrations of Mode S and ADS-B operation.

Upon request, they can also identify as a specific aircraft using a transponder setting or can use the squawk code to indicate radio failure and emergency situations.[4]

Mode C. Instead of identification, Mode C interrogations trigger the aircraft to respond with altitude as measured by the aircraft. Typically, Modes A and C will be combined, so that ATC can retrieve both identification and altitude readings using SSR.

Mode S. For this mode, aircraft are allocated fixed identifiers known as ICAO transponder addresses. These addresses are 24 bits long, expressed as six hexadecimal characters and assigned by country of origin. As such, most non-hobbyist aircraft have an ICAO address and are identifiable by it. It is important to note that this is separate to a registration name; transponder addresses are unique to a transponder and attached to all Mode S communications, whereas registrations are not. A Mode S interrogation can be addressed to specific aircraft, and can request other data fields on top of altitude, with responses either be 56 bits or 112 bits in length.

In many parts of Europe, transponders have been mandated for IFR flight for some time, and finds its foundations in the Identification Friend or Foe system

[4]The transponder setting is referred to as `IDENT`, for identify.

used in World War II [80, 81]. Interoperability exists between the modes—when an aircraft or ground station does not support Mode S, the exchange defaults to Mode A/C. Interrogations are sent on 1030 MHz and responses on 1090 MHz.

Since SSR is used extensively by ATC, channel congestion has become a major problem. This has caused Mode S interrogations in particular to be very lossy. One analysis suggests that in quieter airspace, message loss stands at around 20% [82]. However, as airspace gets busier—beyond 60 in-range aircraft in the case of this study—loss exceeds 50%. With more aircraft equipping Mode S, a heavily congested channel will become a bigger problem.

Building on Mode S is Mode S Extended Squitter (also known as 1090 ES), which uses a longer data packet to encapsulate Automatic Dependent Surveillance-Broadcast (ADS-B) data [83].

Security Research. So far, little work has specifically looked at the security of Mode S beyond the analysis in [61]. The authors explore the threat of jamming and spoofing in some detail, considering both to be realistic. Jamming is noted to be arguably easier due to spoofing possibly constraining the attacker more. Some research into SSR usage relevant to security exists, including [84] which discusses the link and its congestion and [85] which considers military usage of SSR and identifies limited equipage. In [86], the authors analyse non-malicious integrity issues, identifying some real world examples including unexpected position changes according to velocity.

Automatic Dependent Surveillance Broadcast

One of the major technological adoptions in the past 15 years, ADS-B OUT is intended to eventually replace radar as the primary surveillance tool for ATC [87]. Aircraft independently broadcast their identity, velocity, position and heading on a regular basis, without the need for SSR interrogation. We provide a diagram of this in Figure 2.7b.

One of the major advantages of ADS-B is that no ground interaction is required to get information about aircraft behaviour. This helps to give a more up-to-date

surveillance picture than SSR in some regions, especially where SSR coverage is limited. Cost reduction is also a benefit; reception of ADS-B messages does not require SSR, meaning ground stations can be deployed more widely, thus increasing ATC coverage. It is currently in the process of being deployed worldwide, with the US and Europe mandating for most aircraft to be equipped by 2020 [88, 89].

Onboard the aircraft, ADS-B can be transmitted either by a variant of a Mode S transponder (1090 ES, as above), or for general aviation, the Universal Access Transceiver [90].[5] UAT is a separate system to Mode S, with messages being transmitted on 978 MHz.

ADS-B OUT is complemented by ADS-B IN, which is a data uplink receiving broadcast messages to improve situational awareness [91].[6] More specifically, these include the following (summarised from [91, 92]):

- Weather and aeronautical reports in the form of the Flight Information Service—Broadcast (FIS-B), currently only available on the UAT link,

- Traffic data from aircraft not using ADS-B, relayed from the ground, as part of the Traffic Information—Broadcast (TIS-B) service,

- Traffic information from users of the other ADS-B transponder types via Automatic Dependent Surveillance—Radar (ADS-R), e.g. if an aircraft is using UAT, then ADS-R can provide returns for aircraft transmitting on 1090 ES,

- Direct traffic returns from other aircraft transmitting ADS-B in the same manner, e.g. an aircraft with a UAT transponder receiving ADS-B transmissions from a nearby aircraft with a UAT transponder.

[5] An extra option exists under VHF data link mode 4, but this is not planned for use going forward. Indeed, it is not an acceptable solution under EU Commission or FAA mandates.

[6] For brevity, when ADS-B is referenced in this thesis it will refer to ADS-B OUT unless otherwise stated. Whilst ADS-B IN is intended for wide usage, current regulator efforts are towards deploying ADS-B OUT to support modern ATC surveillance.

Security Research. Awareness of the security issues in ADS-B stemmed from talks at Blackhat and DEF CON, as well as some early work in journals [15, 16, 22]. Since then, this has seen attention from the academic community, further analysing the problem [20, 81], demonstrating that the system is vulnerable to message injection, modification and deletion.

Since then, a number of solutions have been proposed. Many cryptographic approaches have been suggested, such as [93] which suggests applying a Public Key Infrastructure approach, or [94], attempting to provide authentication through modulation and [95] assessing the suitability of NIST standard cryptography. Since avionic systems are not easily updated or redeployed, such solutions have gained little traction.

Other work has opted to work around the system, providing a second line of defence. The 'gold standard' approach is multilateration, wherein each ADS-B message is received by four or more independent sensors [96, 97]. However, this requires costly infrastructure. As such, other approaches address this and include using crowd-sourced networks with fewer sensors [98–100] or fingerprinting transponder signals [101].

Automatic Dependent Surveillance Contract

Although it has a similar name to ADS-B, Automatic Dependent Surveillance—Contract (ADS-C) is a different technology. ADS-C establishes a one-to-one surveillance contract between an aircraft and the ground. This runs over a data link such as ACARS, and can carry a wide range of information beyond SSR or ADS-B capabilities. It is not broadcasted via a transponder as with ADS-B.

An ADS contract can be negotiated between parties to carry specific data items at a given frequency [102]. On top of current position, it can also provide meteorological information and intended route at different levels of detail. Furthermore, events can be specified by ATC to which the aircraft can report on—for example, changes in altitude or encountering next waypoint.

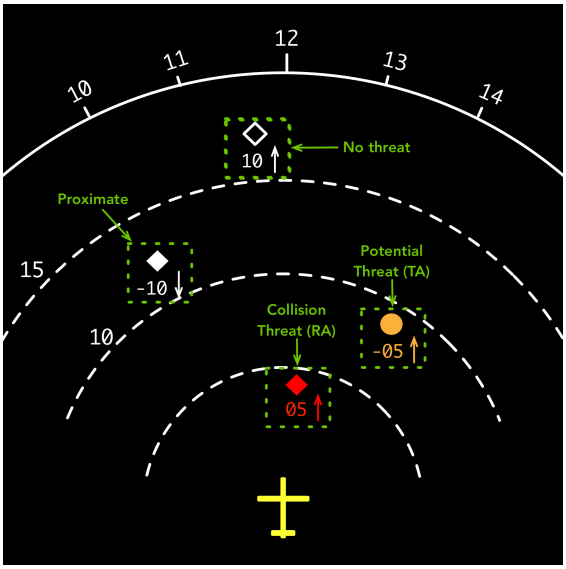

Figure 2.8: A representation of TCAS data as seen by the pilot in the cockpit of an airliner. This is based on the Airbus Navigation Display (ND). Dashed semi-circles represent intervals at a range selected by the pilot, and numbers around the solid semi-circle are heading values.

2.3.4 Safety Nets

Some systems are specifically designed to provide safety protections to crew. Typically, these are designed to give alerts when a situation will become dangerous so that it can be managed. In some cases, the systems will provide guidance or instructions on how to manage the situation. We outline two key systems, the Traffic Collision Avoidance System (TCAS), and Ground Proximity Warning System (GPWS).

Traffic Alert and Collision Avoidance System

Although airspace is tightly controlled by ATC, situations can arise where aircraft become too close to each other. This has resulted in mid-air collisions, such as the 1996 Charkhi Dadri crash, where an aircraft unduly descended and collided with another whilst under ATC control [103]. Incidents such as these have led to regulators requiring aircraft to be equipped with collision avoidance systems, which may take over from ATC control when a dangerous situation arises.

The Traffic Alert and Collision Avoidance System (TCAS) is an implementation of the Airborne Collision Avoidance System (ACAS), designed to help reduce the

chance of a mid-air collision [104, 105]. It has been required in some form on many aircraft since 1993, with TCAS II being introduced in 1998 [106]. In a situation where the risk of a mid-air collision is unacceptable (i.e. two aircraft are on course to collide soon), TCAS on each aircraft will communicate and negotiate deconfliction actions for each aircraft [105].

The system onboard the aircraft uses Mode C and S transmissions to detect and notify nearby aircraft of its existence, responses from which are then processed and displayed to the crew. This will typically be presented as in Figure 2.8, with threats ahead of the aircraft being shown. Other aircraft can be *no threat*, *proximate*, a *potential threat* or a *collision threat*, depending on their distance, rate of closure and altitude difference. If an aircraft is a *potential threat*, a Traffic Advisory (TA) is given to crew, warning them of a potential intruder. If the intruding aircraft gets closer, a Resolution Advisory (RA) alert is given, which is accompanied by instructions that the crew must follow. We cover TCAS in more detail in Chapter 8.

Ground Proximity Warning System

Controlled Flight into Terrain (CFIT) is a major risk for aircraft, especially when operating in difficult weather conditions or around challenging terrain. In an effort to reduce the chance of CFIT incidents, Terrain Avoidance and Warning Systems (TAWS) attempt to identify when situations leading to CFIT begin to emerge and alert the crew.

The Ground Proximity Warning System (GPWS) is a primary example of TAWS which mainly uses the radio altimeter to detect the rate of closure to the ground [107]. GPWS used in larger aircraft has additional features, including the call-out of altitudes on approach, and identification of dangerous weather.

More modern versions of GPWS are beginning to see use. These employ GPS as a location source and include a terrain database against which to compare it. The commercial variant is known as Enhanced GPWS (EGPWS), and is gradually being deployed on commercial aircraft [108]. We cover GPWS further in Chapter 8.

As long as I breathe, I attack.

— Bernard Hinault

3 Threat Model

Contents

Traditionally, avionics were effectively closed-off to outsiders. A combination of expensive multi-document standards, proprietary hardware and software, and high levels of system interdependence meant that it was difficult for third parties to even learn about avionics. This created a culture of security-by-obscurity which carried through to wireless communications; many systems use restricted frequency bands and receiving or transmitting the signals requires specialist hardware.

Recently, this has changed due to the ready availability of commodity software defined radios (SDRs); a significant technological shift making it more accessible and affordable to develop radio transmission and reception tools. Importantly, SDRs have allowed low-skilled users to collect (and sometimes transmit) avionic communications with little technical knowledge.

Because of this, the threat model has changed considerably. Methods and

tools once only in the domain of a well-resourced attacker are now achievable for a relative novice, and advances in the availability of radio technology are enabling new threats [109]. In this chapter, we outline different types of attackers and their possible motivations.

3.1 Causes of the Changing Threat

Multiple factors have contributed to the shift in the aviation threat model, particularly in the past 10–15 years. Whilst the existence of security issues in avionics has not been brought about due to a changing threat, the public awareness of such issues has been heightened. We believe two factors have sped up the process of awareness: wider availability of radio hardware, and a growing interest from enthusiast communities.

3.1.1 Availability of Software Defined Radios

For a long time signals processing was primarily hardware-based since it requires good computational performance on mathematically intensive tasks. As improvements in software-based signals processing were made, coupled with general radio hardware becoming more affordable, SDRs became a popular option for research.

Until recently, these were expensive tools primarily for use by radio experts. One of the first to lower this bar was the Ettus Universal Software Radio Peripherals (USRP), built on open source software and intended to give better accessibility to researchers [110]. Even then it was expensive, with the base unit costing £1000-2000, but was supported by the open source community through the GNU Radio tool [111]. This simplified the development process of radio software by enabling reusable modules for common tasks. Arguably the biggest step-change happened when SDRs became cheaply available and so moved into the reach of hobbyists. Two products which enabled this were the HackRF and RTL-SDR.

RTL-SDR. Often a digital television USB stick (DVB-T) with the RTL2832U chipset, these SDRs the RTL-SDR firmware to repurpose the hardware [14]. Typically available for £10-15, these devices are capable of reception (RX), usually in the range of 20-1500 MHz (dependent on chipset), and can be interfaced with GNU Radio. This makes the development and sharing of software for SDRs relatively straightforward. Recently, flight trackers such as FlightAware have begun to offer RTL-SDR variants tailored to their ADS-B collection service [112].

HackRF. For those wishing to transmit (TX) as well as receive, the HackRF provides an affordable platform to do so. Capable of TX and RX between 1 MHz and 6 GHz it offers a purpose built platform for testing RF systems, costing around £200-300 [113]. As with the RTL-SDR, the HackRF interfaces with GNU Radio, making it easier to create and share radio software.

3.1.2 Enthusiast Community

The effect of an active enthusiast community on the changing threat within aviation cannot be understated. This is not a new phenomenon, with VHF scanner groups collecting ACARS messages since at least early 2000s through `acarsd`, and planespotters logging the coming and goings of aircraft for even longer [114].

The introduction of the RTL-SDR has brought about a wider community interested in aviation signals, prompting many to look at the field who might not have done previously. As such, many people now work on third party tools for aviation communications.

More closely tied to security is the growth of non-academic security conferences, at which the intersection of radio enthusiasts and security researchers can be seen. Venues such as Blackhat, DEF CON or BSides provide researchers the opportunity to present exploratory ideas, often relating to real-world systems. Indeed, some of the early work in aviation security originates from these conferences. Examples include Costin et al. at Blackhat USA discussing ADS-B security and Hoffman et al. at DEF CON analysing the use of BARR blocks [15, 53].

Feeding research from this community into aviation has been slow, arguably due to mismatched incentives. It is important for the aviation industry to not only be safe and secure, but also to be publicly perceived as such. Following the terror attacks on New York on September 11^{th} 2001, US aviation saw the start of a decline in the US domestic market and between 2001 and 2002, recorded losses of $19.6 billion [115, 116]. In contrast, the security community often publicly announce vulnerabilities after a disclosure process, inevitably affecting public perception of industries or organisations either positively or negatively. Furthermore, unknowns around the potential for harm of security vulnerabilities can lead to exaggerated claims. These two approaches are not immediately compatible which has led to high profile disagreements between the communities, in doing so slowing progress in cooperation [19, 117].

Typically, the intention of enthusiast is not malicious but instead focussed on increasing knowledge or gaining recognition within a community. An inevitable consequence is that more light is shed on aviation security, particularly on its shortcomings. Coupled with an increase in open-source tool development, it is possible to access or develop tools which could be used to interfere with avionic communications.

3.2 Mapping the Threat

Given the variety of avionics in use, attackers aiming to leverage security vulnerabilities come with a range of motivations, capabilities and levels of resource. Depending on the type of avionics, attacks can have a many possible outcomes ranging from sensitive information leakage to compromising safety. Some threat actors may be focussed on or capable of one type attack but not others.

In this section we outline threat actors and attacker models. Threat actors summarise the high-level resources and general motivations of an attacker within a group, whilst the attacker models provide three concrete capability sets tied to how they interact with avionics. We also map threat actors to attacker models as a basis for motivations and levels of resource expected for each model. Throughout

the thesis we will develop these models to ground our measurement of security issues in realistic threats.

3.2.1 Threat Actors

Attacks on avionic systems can have range of effects from minor inconvenience to severe safety erosion, depending on the aims of the attacker. We now cover the threat actors relevant to aviation communications.

Hobbyists. This group is different to the others in that they are not aiming to cause harm themselves, but their actions could do so indirectly. They are interested in developing community knowledge through experimentation, modification and data sharing. Although they are passive, they may collect data and store it, act upon it or share it with others. Combined with this, we expect them to at least be capable of setting up a computer with an SDR and antenna for reception, then using existing SDR software to receive aviation signals. In some cases, they could build proof-of-concept tools to do so.

Grey Hat. Rather than being driven by a community effort or for interest in the domain, grey hats instead are security focussed, typically operating alone or in small groups. They are likely to have similar attributes to a hobbyist in this sense. They aim to explore what is possible on systems; because of this, they are not necessarily aiming to cause negative effects but instead may be investigating for community credibility. An example is the case of Chris Roberts, who interfered with an in-flight entertainment system onboard an aircraft [118]. Arguably, his motive was to just explore the system, but also boasted about it online and caused media outrage. In turn, they are more likely to cause disruption than harm or reduction in safety.

Criminal Groups. Seeking to attack systems to make financial gain, criminal groups might target individuals or companies. Attacks could be of a range of complexities but may not necessarily cause disruption—surveillance alone may be sufficient for their task. An example of this saw insider trading supported by paying

someone to manually monitor aircraft movements at an airfield near to many multi-national organisations [119]. In this case, the attacker intended to spot merger and acquisition activity prior to announcement, to direct investment illegally. This actor might deploy hardware to locations to operate remotely or in a distributed fashion.

Activists. This group aim to cause a lot of disruption to highlight their cause—which in turn could have collateral effects—but they are not intending to cause destruction or loss of life. Since the aviation industry consumes a lot of space and has a large impact on many aspects of people's lives, it becomes a natural target for protest. We consider activists to be relatively low-resourced but well-organised since they are aiming to cause an effect with a specific purpose. Attribution is likely to be sought although we limit this group to not aiming to compromise safety. A non-technical example of this is an incident at London Heathrow in 2017, wherein activists attempted to prevent a charter aircraft from deporting people by chaining themselves around it [120].[1]

Terrorists. Probably the most destructive group, terrorists wish to cause harm through a loss of safety, or through causing severe disruption. This might be subtle or overt but will not be concerned by any collateral which might occur due to the attack. They would be aiming for exposure and to cause fear, so would seek attribution. This actor could specifically develop capability to perform an attack. The UK Government considers terrorist groups a threat to civil aviation, albeit one that is currently low risk, as per the Aviation Cyber Security Strategy [109].

Competitors. Originating from the aviation industry itself, malicious competitors can fall into two threat groups; surveillance or attacks. In the former, a competitor may wish to learn how the target operates through observing private communications in order to identify any advantage that they have. The latter category is less likely but must be considered. In this, an attacker wishes to cause reputational damage

[1] It is important to note that the law may treat this actor in the same way as a terrorist actor, but we distinguish them due to their aims.

or financial loss to a competitor though an attack. The motivation for this would be to try to improve their own standing within the market by harming others. Importantly, this actor will probably stay clear of compromising safety, as doing so in safety-critical industries would have a knock-on effect to themselves. The UK Aviation Strategy considers this attacker realistic, especially with regards to causing reputational harm and notes that a competitor threat actor may use insiders to achieve their aims [109].

Nation States. This group is the most advanced adversary, so are very hard to defend against; for the purposes of simplicity, we include state-backed entities. Their focus would be specific, possibly limiting the potential for collateral effects. Again, aims can be split into two categories, according to whether the effect would be attributable. They may wish to perform surveillance on other nations, using their resource to widely deploy hardware. On the other hand, their actions may have a more visible outcome, such as disruption, denial of service or reduction in safety. An example of this kind of action was seen in 2016, where an apparently state-backed Russian advanced persistent threat (APT) group caused Swedish ATC to lose capability for a period of time [121].

3.2.2 Attacker Models

Having considered the capabilities and the threat actors, we now tie them to three attacker models used in this thesis. We summarise the mapping of potential threat actors to attacker models in Table 3.1.

Passive Listener

This attacker model concerns actors which simply collect data from one or more locations using commodity hardware such as the RTL-SDR, and so are passive with respect to the medium. It is coupled with freely available collection software, thus making the baseline resource requirement low. Their minimum skill level is low, capable of setting up and storing collected data, which may be shared with a community. This aspect presents the main cause of threat from the attacker, as

Table 3.1: Mapping of attacker models to potential threat actors.

	Threat Actors						
Attacker Model	Hobbyist	Grey Hat	Criminal Groups	Activist	Terrorist	Competitor	Nation State
Passive Listener	✓	✓	✓				✓
Passive Codebreaker		✓	✓	✓		✓	✓
Active & Determined			✓	✓	✓	✓	✓

data aggregated over time and across a large geographic area gains power. Data collection might be for honest-but-curious activities such as plane spotting, or with malicious intent such as uncovering personal data. Actors which might fit within this model include hobbyists, grey hats or criminal groups.

Passive Codebreaker

Extending upon the *passive listener*, the passive codebreaker has the same collection intention as the above attacker but rather than simply collect and observe data, they attempt to break any encryption or protected messages seen. This is a slight distinction in that they are still passive with regard to the medium but are intending to break protections in place. To achieve this, their capabilities might be slightly increased—both in terms of resources and skill. This would allow the collection of more data and the skill needed to break encryption where possible. Threat actors fitting into this category might be grey hats, activists, criminal groups or competitors.

Active & Determined

Our final attacker model is our strongest, wherein the threat actor is aiming to cause disruption, financial and reputational loss, and possibly a reduction in safety. To achieve this, we presume they have a medium to high level of skill and resource, capable of researching and implementing new attacks. The hardware used will

be of a higher specification than our attackers above, capable of transmission as well as reception. To do this, they will need amplifiers and directional antennae capable of higher transmission powers. They can deploy hardware and operate it remotely if needed. We consider the threat actors in this group to be criminal groups, competitors, activists, and possibly terrorists or nation states if they are aiming for lower impact attacks.

If you go with a break, you can either win or not win.
If you don't go for it, you definitely won't win.

— Jens Voigt

4

ACARS Data Sources & Collection

Contents

In this chapter we outline the data collection process and third-party sources used for the work in the first part of this thesis, namely the analysis of security and privacy on the ACARS data link. We also cover the survey used to establish some

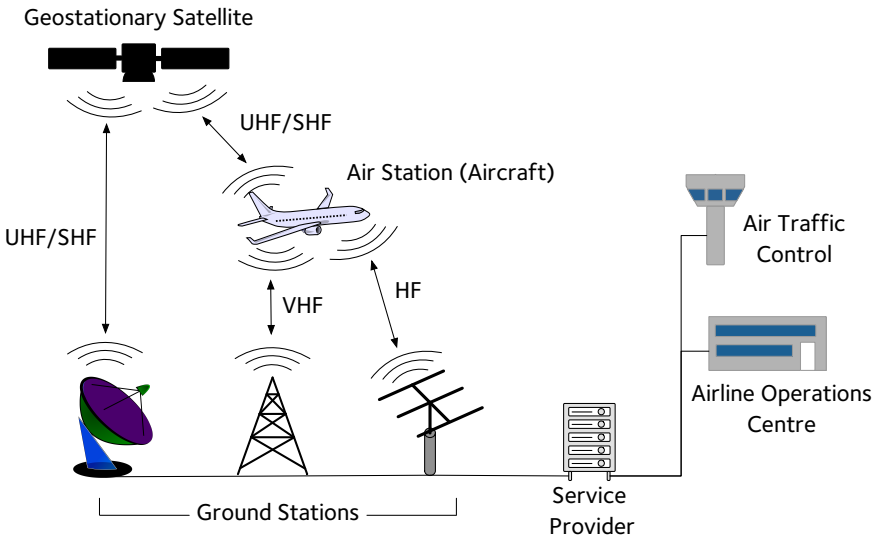

Figure 4.1: Diagram of ACARS subnetworks over Very High Frequency (VHF, Plain Old ACARS and VHF Data Link Mode 2 systems), High Frequency (HF, High Frequency ACARS), and Ultra/Super High Frequency (UHF/SHF, satellite communications based ACARS).

ground truth on usage of security within avionic data links. We will refer to the data in this section throughout the thesis and begin by describing ACARS in detail.

4.1 Aircraft Communications Addressing and Reporting System

ACARS is an avionic communications system used worldwide. Deployed in 1978, it provides support for airlines and ATC to communicate with the vast majority of commercial aircraft [122]. Although originally deployed as a coupled data link and application-layer communication system, ACARS is now also used as application which runs on a range of data links.

System Overview. Fundamentally, ACARS operates in a very similar manner to text messaging over mobile cellular networks. The network is operated and maintained by a service provider. This is a limited market with only two providers—SITA and ARINC. On the ground, organisations can be connected to the ACARS network via a service provider, so that they can send and receive messages. On board

Table 4.1: Key feature comparison of ACARS delivery subnetworks.

Subnetwork	Coverage	Frequency	Link Speed
HF	Worldwide	2.9-22.0 MHz[1]	Up to 5.4 kbps[2]
'Plain Old' ACARS	Continental, over land	∼131 MHz	∼2.4 kbps
VHF Data Link Mode 2	Continental, over land, limited deployment	∼136 MHz	∼31.5 kbps
SATCOM	Worldwide, except polar regions	L-Band (1-2 GHz) up C-Band (6-8 GHz) down	Up to ∼400 kbps[3]

an aircraft, systems interface with a Communication Management Unit (CMU) through which they can send and receive ACARS messages to organisations on the ground. We now go through each part in detail, and provide a diagrammatic representation in Figure 4.1.

4.1.1 Data Link

Although primarily carried on airband VHF (known as Plain Old ACARS, POA), ACARS can also be served in other ways. As shown in Table 4.1, it can be carried on the High Frequency Data Link (HFDL), the VHF Data Link Mode 2 (VDL Mode 2) or over satellite communications (SATCOM) [55].[4] In this section, we provide some technical detail on the subnetworks used to deliver ACARS.

VHF Subnetworks

Due to VHF being the original ACARS subnetwork, usage is widespread. As of 2012, ARINC network coverage comprises over 1100 POA stations and 400 VDL2 stations, with much of the populated regions of Europe, North America, South America, Asia and South Pacific served [123]. As such, VHF links are extensively used for travel over land or near the coast and are usually the primary data link option.

[1] Depending on atmospheric conditions, HF frequencies are reassigned regularly.
[2] This depends on the baud rate and keying used.
[3] Exact speeds vary depending on service, here 10 kbps is provided by the Inmarsat ClassicAero service, with the higher rate provided by their SwiftBroadband service.
[4] For brevity, in this thesis we will refer to VDL Mode 2 as VDL2.

Plain Old ACARS. In some form, POA has been available for over 40 years and is defined in ARINC 618 [124].[5] Channels have a bandwidth of 25 kHz and use Minimum Shift Keying (MSK) to encode data onto the signal [125]. This gives a notional data rate of 2.4 kbps. To reduce transmission collisions, POA uses Carrier-Sense Multiple Access (CSMA). In Europe, POA is carried on frequencies around 131 MHz. Although this version of ACARS is still used, it is considered the legacy option since it is slower than the newer VDL2.

VDL Mode 2. A more modern link, VDL2 also operates in airband VHF like POA but on frequencies around 136 MHz. It is a more general link than POA, primarily intended to carry CPDLC, but also able to carry ACARS messages via technology called Aviation VHF Link Control (AVLC)—serving ACARS over this is called ACARS over AVLC (AOA) [126]. By using a newer encoding technique in Differential 8-level Phase Shift Keying (D8PSK), the link has a higher bitrate of 31.5 kbps.

SATCOM

Satellite-carried ACARS is offered via the Iridium and Inmarsat satellite constellations, each with slightly different options and service levels.

Inmarsat. This cluster offers two ways to serve ACARS: via a data service called SwiftBroadband, and through Classic Aero [127]. SwiftBroadband is a more modern, data-focussed service which uses two channels to offer data prioritisation for safety-related information [128]. Classic Aero has been on offer for over 25 years, delivering voice and data services [129]. Both services are offered on the Inmarsat-4 cluster, made of geostationary satellites, which covers most of the world except polar regions [130]. These satellites provide communications on the L-band (1-2 GHz), with the choice of service dictating the data rates [131].

[5] Although less common, POA can also be referred to as VDL Mode 1.

Iridium. Offering a much less expansive service, Iridium instead can serve ACARS over its existing network with a service called Short Burst Data (SBD). Unlike Inmarsat, Iridium uses low earth orbit satellites instead, requiring lower transmission power to communicate with than for geostationary satellites [132].

High Frequency Data Link

Although the lowest priority of the subnetworks, HFDL provides almost worldwide coverage (excluding the south polar region) using 32 of 167 available frequencies at any time [133]. It is defined by ARINC 753 and operates on frequencies between 2.9–22.0 MHz, using HF voice channels for data [134]. Said frequencies are rotated depending on atmospheric conditions, offering a range of data rates (300–1800 bps) and Phase Shift Keying (PSK) encoding methods (2, 4 and 8 PSK) [135]. Rates and encoding methods are also chosen depending on the atmospheric conditions through which the signal must be sent. At a high level, a lower the data rate allows more room in the data packet for error correcting methods, ideal if conditions for transmission are poor.

4.1.2 Message Format

All ACARS subnetworks use the same message structure, illustrated in Figure 4.2, but some subnetworks allow longer text portions of the message. Characters are from the ASCII character set, as per ARINC 618 [124]. However, this is a more recent addition—some older parts of the ground network are only compatible with a reduced ASCII set [124]. Depending on which parts of the network are used, the even further reduced Baudot character set may need to be used, effectively limiting the set to `A-Z`, `0-9`, `,-./`, and some control characters. This would be due to messages travelling through legacy systems which cannot handle the full character set. Thus, if maximum compatibility is required, a smaller character set must be used.

Most of the information in an ACARS message is carried a free text element, which forms the largest part of the message (see Fig. 4.2). Other fields of note are the `ICAO registration` and `Flight ID` fields which identify the aircraft involved

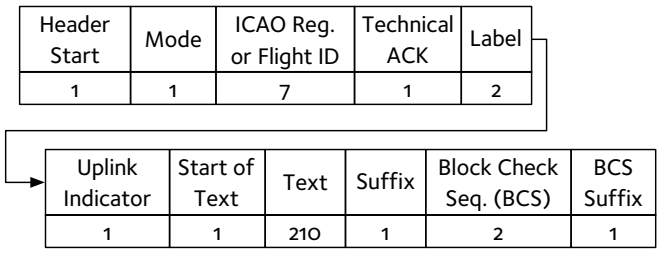

(a) Uplink message format

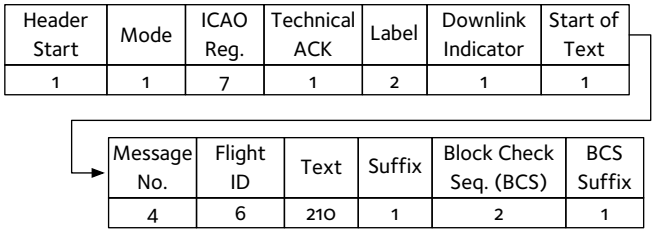

(b) Downlink message format

Figure 4.2: ACARS message structures for uplink (air-to-ground) and downlink (ground-to-air) based on ARINC 618 [124]. Field sizes are in ASCII characters (i.e. bytes).

in the messages. ICAO registrations are unique to the aircraft rather than the avionics, as with a transponder code. Both allowing identification across flights, though registrations are usually displayed on the outside of the aircraft. In contrast, flight IDs are tied to a single flight and only rigorously used by commercial aircraft.

4.1.3 Message Handling

ACARS messages originate from and are directed towards specific systems. The original mechanism to do this is the `label` field within the message, as indicated in Figure 4.2, with labels being standardised in ARINC 620 [136]. Labels act as routing fields for the Communications Management Unit (CMU) onboard the aircraft, which routes messages to the correct endpoint system [55]. Specific labels for a message type can vary depending on whether they are used on the up or down links. Some example downlink labels are:

- `00` – Hijack situation report,
- `5U` – Weather information request,
- `Q2` – Estimated Time of Arrival (ETA) report.

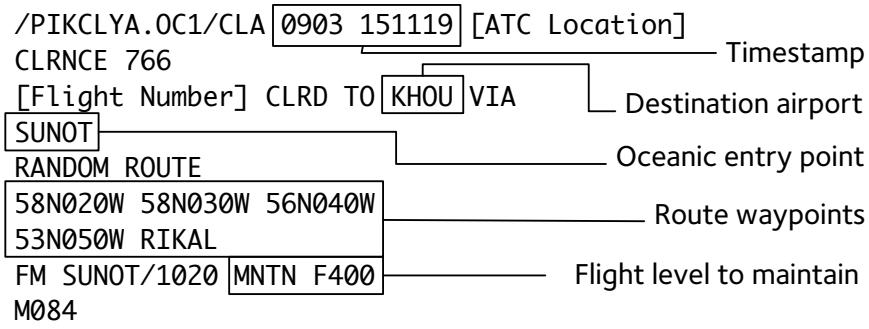

Figure 4.3: Example ACARS oceanic clearance message as transmitted from air traffic control to aircraft.

Although many labels are defined, ACARS usage has grown considerably and is used for purposes beyond its original intention. To handle this, further message handling data is stored in the `Text` field, or if needed, an organisation can use free-text to accommodate their own structure.

4.1.4 Uses of ACARS

As indicated above, the system can be used for many purposes, broadly divided into ATC and AOC/AAC groupings. We now outline some primary uses within each grouping for ACARS.

Air Traffic Control

Air traffic control (ATC) messages are used to ensure that the aircraft can fly on its route safely. This takes the form of pre-departure and oceanic clearances, as well as aerodrome information, served as the Digital–Automatic Terminal Information Service (D-ATIS, see Chapter 6). An example of an oceanic clearance message is given in Figure 4.3. It is particularly useful in situations where voice is not available or too busy, such as over the ocean or on congested frequencies near airports. Although ACARS does see some ATC usage, this is unlikely to be a long-term solution. Instead, as CPDLC becomes more widely adopted, that will be used instead of ACARS for ATC purposes [137].

Aeronautical Operational Control & Airline Administrative Control

Non-ATC communication forms a significant part of ACARS traffic, since it performs many important tasks in operating larger aircraft. These can be roughly divided into operational (AOC) and administrative (AAC) messages. Both use a combination of structured as per standard and free-text messages.

AOC messages include:

- Maintenance reports for broken or malfunctioning items,
- Out, Off, On, In reports for timekeeping,
- ETA messages,
- Weather reporting,
- Transfer of takeoff parameters,
- Downlinking engine health management data,
- Position reporting.[6]

AAC messages instead focus on passenger and crew needs so include:

- Passenger connection information, such as sending connecting flight information to the aircraft,
- Requests for information on events such as elections or sports fixtures,
- Medical reports on unwell passengers,
- Requests for wheelchairs to meet passengers,

These messages use the free-text nature of ACARS, with messages ranging from automated, structured reporting to text messaging between crew and ground operators. Lists of passengers transferring to other flights, maintenance issues and requests for aid of disabled passengers are common sights, though exact usage varies between airlines. It is also common for flight plans to be served over ACARS, which a pilot will then input into the flight computer.

[6]This is usually in a custom format unlike the well-structured ATC surveillance reports.

4.1.5 Security in ACARS

ACARS as an application has no security mechanisms by default. Whilst some subnetworks (notably some Inmarsat options) may provide some mechanisms at the link layer, most do not. Furthermore, awareness or concern with ACARS security is not common throughout the industry. We now look at the existing work in this area.

ACARS Message Security

A standardised security solution for ACARS does exist—it is known as ACARS Message Security (AMS) and is standardised as ARINC 823 [138]. It was partially developed by Honeywell, Inc. and the US Air Force, the former of which have an implementation named Secure ACARS [139]. It provides a range of provisions to sign, authenticate and encrypt message content, and some configurations can be used to protect aircraft identity.

Two methods for secure communication are defined; using public/private keypairs or using a shared secret. Furthermore, Secure ACARS uses well-established cryptographic mechanisms at the time of development to enact signing, authentication and encryption [140]. These match to NSA Suite B, which was a US Governmental information security standard but is now deprecated.[7] As such, Secure ACARS would likely need to adopt a more up-to-date cryptographic suite to be effective going forward.

In terms of the protocol for itself, Blanchet undertook verification on parts of AMS and proposes some fixes for the problems found [142]. Namely, they identify potential message freshness issues, a possible fall back to in-the-clear communication if session negotiation fails, and vulnerability to replay attack. Each of these can be addressed with adjustments to the standard.

Another implementation of AMS is in development by ARINC, and is known as Protected ACARS (PACARS) [143]. This is intended as an 'open' option—rather than being reliant on the service provider to deliver a secure solution, it aims to

[7] In recent years, the NSA has moved away from Suite B, marking the related IETF RFCs to historic [141]. Furthermore, the NSA appears to have removed Suite B related references from its website, indicating that this standard is no longer active.

run on top of ACARS as an application. This is a more recent system and uses similar cryptographic mechanisms to Secure ACARS.

State of the Art

Despite the existence of AMS, usage appears limited. As indicated in [142], at least one airline does use it, though many other ACARS users—some of whom are privacy sensitive—do not. We cover this in relation to privacy and security in Chapters 6 and 7, but first look at the state of the art and existing literature in this area.

Unlike some parts of the aviation community, militaries are engaged in the security challenges posed by ACARS. For example, [144] identifies that the clear-text nature of ACARS is considered an important weakness especially for military aircraft, due to operational privacy requirements. Furthermore, in [145] the authors demonstrate efforts to manage the lack of security through encryption, highlighting the requirement for privacy in the military context. In both, ACARS defaulting to sending in the clear is the motivation for users to require some measure of security.

The role of ACARS security has occasionally been discussed outside of academic research. In [146], the authors note the challenges of deploying Secure ACARS, as well as its development process with the US military. In the hacking community, [17] claims to use ACARS to upload malware onto a flight management computer. Despite the strong claims in this paper, the results were refuted by Honeywell, whose systems were used in the attack.

In [143], the author considers issues caused by the lack of security on the link. They highlight that crews rely on information sent via ACARS, which could have safety implications since the source of that information cannot be confidently determined on an unprotected link. The paper also has concerns with third-party message interception and transmission, identifying that this could cause trust issues for pilots using ACARS.

Some other security solutions apart from AMS have been proposed. One such example is [147], in which a solution for 256 bit AES is presented but lacks further implementation detail. Since it primarily focuses on providing message security

through AES-256 but does not address how to fit this within the operation of ACARS, it has not seen implementation.

User perceptions are also notable; in [148] we get an insight into such views. A survey was carried out on pilots, air traffic controllers and other aviation industry professions, investigating their perceptions of security on a range of avionic systems using wireless communication. Within this, they were asked about the integrity and authenticity of ACARS with most believing that the protocol offered some level of protection by default.

Although not directly covering ACARS, [60] investigates security for CPDLC, which is carried on links shared by ACARS. The authors explain that the lack of security mechanisms means that no meaningful security guarantees can be made. It also suggests that PACARS could be used to protect CPDLC, though implementation details on this are unclear. Furthermore, it suggests adopting Host Identity Protocol, a TCP/IP addition which uses asymmetric cryptography to verify sender identity. However, this would require adaptation to work with CPDLC and would need a public key infrastructure, something not yet in place.

Clearly, the concept of security for ACARS has been considered for some time. Even though a standardised approach to security on the link exists, the previously referenced airline which does use AMS appears to be the exception rather than the rule. One explanation for this might be a lack of external incentives to use security—a situation which could change as awareness of how ACARS is used for sensitive data evolves. These incentives could be customer requirement, avoidance of financial loss through data breach or compliance to regulation. However, such motivating factors rely on a public understanding of how the link is used which, currently, does not appear to be the case. We discuss how to improve this understanding in the rest of this chapter.

4.2 Data Collection

To understand how the security provision on ACARS affects those who use it, we collected real-world usage data from the system. We were particularly interested in

representing a passive listener as described in Chapter 3, since this is the lowest bar for an attacker to meet both in terms of requirements and skill.

Data collection was carried out from a single location with commercially available hardware and software, on SATCOM and both VHF subnetworks, intermittently over the course of two years. All data was collected from a single location in Thun, Switzerland, with the data being summarised in Table 4.2 and a full list of collection frequencies given in Appendix A. We collected data in this location in order to comply to laws and regulations regarding data collection of aviation data.

4.2.1 VHF

Collection for both POA and VDL2 is very similar and only differs at the software level. A standard computer was connected to an RTL-SDR, which was fed by an airband ground plane antenna (e.g. the Sirio GPA 108–136 [149]). This replicated a representative set up of a passive threat as described in Chapter 3.

Our POA setup used the previously described hardware with the ACARSDec software [150]. We collected for 141 days (May–October 2016) followed by a further 220 days (March–October 2017) on the three European channels: 131.525 MHz, 131.725 MHz and 131.850 MHz.

Collecting VDL2 messages requires the same hardware set up but with VDL2 decoder; in this case, it was `dumpvdl2` [151]. We collected VDL2 ACARS messages in a separate collection period for 211 days (March–October 2017) over five European frequencies: 136.725 MHz, 136.775 MHz, 136.875 MHz and 136.975 MHz.

For both links we collected the downlink, i.e. from the aircraft to ground stations. Downlink VHF is collectable from any location where aircraft fly overhead, whereas uplink VHF would require either an airborne platform or locations near the uplink transmitter. Since we do not know the locations of the latter, and the former would be very expensive and potentially illegal, we collected downlink only. Again, this is also representative of a passive threat with limited resources.

Table 4.2: Summary of data collected for ACARS analysis on the VHF (POA and VDL2) and SATCOM subcarriers.

Subnetwork	Time Period	# Days	# Messages	Link Direction
POA	May–October 2016 March–October 2017	141 220	372,996 114,308	Down
VDL2	Mar.–Oct. 2017	211	157,568	Down
SATCOM	November 2016–January 2017	68	1,170,745	Up

4.2.2 SATCOM

For SATCOM ACARS, we collected from Inmarsat satellites. For this constellation, the SATCOM uplink is located in the L-band around 1.5 GHz. As with VHF subnetworks, reception uses a computer with an RTL-SDR, this time fed by patch antenna through a low-noise amplifier. Decoding software is in the form of JAERO [152]. We recorded all 11 uplink channels of INMARSAT satellite 3F2 for 68 days between November 2016 and January 2017.

As with VHF ACARS, collecting one direction of the link is easier than the other. Uplink messages are transmitted from the satellite to the aircraft, which is a moving platform. To ensure that messages are delivered, a higher power and lower wavelength signal is used, since there is limited space on the aircraft for antenna. Furthermore, the beam is large to provide good coverage and account for high aircraft mobility. This allows it to be collected with a patch antenna.

Downlink, located in the C-band around 3.5 GHz, has much shorter wavelengths and increased path loss. Ground stations receiving downlink can use bigger receivers than aircraft and are stationary, allowing satellites to transmit at said shorter wavelength with a smaller, more targeted beam. This makes intercepting downlink messages very difficult without requiring third parties to use a large satellite dishes or the ability to listen from nearby the receiving stations. Since this was out of scope for our collection capabilities, we only collect uplink for SATCOM.

4.2.3 High Frequency Data Link

HF ACARS uses long wavelength signals with much lower frequencies than VHF or SATCOM. Using the representative range of 2–22 MHz, this gives wavelengths of ~149 m to ~13 m respectively. Collecting signals with such long wavelengths, especially where transmission power can be low, requires large antenna. We did not have the capability to install such antenna at our collection location, so could not reliably collect HF ACARS. On top of this, our collection location is far from oceans and the polar region, thus meaning that the potential users of HF ACARS would be very far away.

4.3 Third Party Data Sources

Data collected through wireless avionic communications will contain a limited amount of information about the sender or recipient. Usually, this will just be the ICAO transponder address or the aircraft registration. We can use these identifiers to gather information about the aircraft from third-party data sources. In this section, we cover the data sources used to complement the collected ACARS data.

4.3.1 Aircraft Positional Data

In some instances, we observe aircraft sending ACARS messages but without location. We can cross-reference this with data from the OpenSky Network provides high-quality historical ADS-B data, and overlaps with our ACARS collection period both geographically and temporally [153]. We can use ADS-B data collected from across the world to check whether an aircraft was transmitting ADS-B signals at the time, and if so, find its position. This becomes particularly relevant when assessing how much additional privacy impact ACARS has on top of ADS-B, as we analyse in Chapter 6.

4.3.2 Aircraft Metadata Sources

To assess the extent of sensitive data transmission, we need gather context on the aircraft. Throughout our work on ACARS we compared our ACARS data with several publicly accessible sources to establish this context. The sources usually provide the aircraft type (e.g., Airbus A320) and the owner/operator (e.g., British Airways, but the owner and operator may not always be the same).

We exploited the following sources:

- FlightAware and Flightradar24 allows a rudimentary check of whether an aircraft is attempting to obscure itself from public flight trackers [23, 154].
- Junzi Sun maintains a database of aircraft from Flightradar24, which allows us to establish whether aircraft have attempted to hide from flight trackers in the past [155].
- Airframes.org offers background knowledge such as pictures and historical ownership data, maintained by a community of hobbyists [156].
- National registries usually provide non-sensitive owner records. One of the largest is for US-registered aircraft, containing over 320,000 records [157].

Most of these sources are important in establishing if an aircraft is *blocked* from public view. Since this is a significant concept in this thesis we expand upon it in Chapter 5. It is also worth noting that these sources are naturally noisy, since they rely on compiling many smaller databases, and aircraft around the world are registered, de-registered and transferred regularly.

4.3.3 Third Party ACARS Collection

At the time of collection there were no active public efforts to collect ACARS messages. More recently, AvDelphi has made a database of ACARS messages available to the public, though details of how it is constructed are limited [158].

Table 4.3: Observed aircraft on the ACARS data link. Breakdown by stakeholder of identifiable aircraft using any subnetwork.

Aircraft Category	Observed Aircraft			
	# Aircraft	% Total	# ADS-B equipped	% Category
Commercial	6899	69.6	6832	99.0
Business	2366	23.9	2323	98.2
Military	438	4.4	423	96.6
Government	183	1.8	178	97.3
Unknown	22	0.2	0	0.0
Total	9908	100	9756	98.5

Before this, ACARS data collection was popular with the aviation scanner community. A wide range of ACARS decoders existed in the early 2000s, though appear to no longer be maintained.[8]

4.3.4 Categorising Aircraft

To reason about privacy and security, we also need to categorise aircraft into business, military or state. Some of this can be automated using the World Aircraft Database as above [155]. For a portion of the observed aircraft, the database contains information about the operator such as whether it is private, military, state or commercial.

Since some aircraft do not appear in this list, part of the process involves manual gathering. This is due to a limited number of third party data sources on aircraft ownership being publicly available. At the time of writing, the most complete database is Airframes.org [156]. This provides comprehensive and historical records on many aircraft but does not provide API access.

4.4 Collection Statistics

Over the course of collection, we obtained 1,815,617 messages across both links, with 1,170,745 (64.5%) being from SATCOM uplink, 487,304 (26.8%) on POA and

[8]The most prominent of these was *acarsd*, though this was last updated in 2007 [114]. Lists of feeder stations illustrate that decoding was quite widespread.

the remaining 157,568 (8.7%) being POA. We identified 9908 individual aircraft, of which 5802 were seen over POA, 4834 over VDL2 and 4540 on the SATCOM channels. A total of 534 aircraft were seen over all three links and 2883 were transmitting on both terrestrial technologies.

In Table 4.3 we show the number of aircraft belonging to each stakeholder group described in Chapter 2, along with their level of ADS-B equipage. We assess ADS-B equipage from OpenSky Network data, cross-checking ICAO numbers with existence according to OpenSky as per Section 4.3.1.

Commercial aircraft make up the majority of all ACARS users at 69.6%, with those qualifying as business jets comprising the other significant portion at 23.9%. Military and state-based aircraft groups were observed to be much smaller at 4.5% and 2.0% respectively—unsurprising considering their exclusivity. Note the high level of ADS-B equipage across the board; ADS-B poses its own security and privacy challenges and so any system which offers other avenues of compromise or leakage compounds the problem [20, 27].

4.5 Survey on Security and Privacy in ACARS

Whilst there is growing awareness of the general security and privacy challenges in aviation, application to specific technologies is limited. We carried out a survey of industry professionals, including pilots, avionics engineers and air traffic controllers to establish a ground truth on the opinions on security and privacy in ACARS. We will refer to this throughout the survey but cover it in this section. This work was approved by our local ethics committee with reference R53464/001.

4.5.1 Background

Our primary aim in running the survey was to get a better understanding of whether members of the aviation industry understood that ACARS lacked meaningful security mechanisms, and so may not be suitable for sensitive data. To do this, we constructed questions to solicit opinions on privacy and security in ACARS,

especially with regard to how it is currently used. Our full set of questions can be seen in Appendix B.

The survey ran for a period of 56 days during October and November 2017. We ran parallel versions of the survey. One was sent to closed pilots' communities and emailed to a range of international ANSPs and regulatory bodies. The other was shared on a range of aviation professional forums, primarily an ACARS user group. This allowed us to filter the data should the 'open' survey collect spam messages, however this did not occur, so we merged the results.

4.5.2 Response Analysis

We received 41 responses to the survey, with the majority coming from flight crew and engineers. Our profession demographics are shown in Figure 4.4a, with organisational in Figure 4.4b. Some 21 (51.2%) of respondents were flight crew, either in a First Officer or Captain role. The remaining majority group was from Avionics Engineering, at just under 25%. This is also reflected in the organisation demographic, with 24 (58.6%) of respondents employed by an airline.

Although many of the participants are unlikely to have the ability to unilaterally choose to either secure or not use ACARS, their responses are still valuable as we aimed to understand awareness of the security of the ACARS link. This is important since the link has grown to be used for a wide variety of purposes, some of which may not be as originally intended. As such, whilst this sample is self-selecting and relatively small, we feel that it provides a useful cross-section of participants. Future work would increase the sample size and target other demographics such as aircraft manufacturers, in order to draw out more views.

Looking to the general questions on ACARS, we can establish a baseline view of the system from industry, with regard to usage for sensitive messages, AMS deployment and the existence of anomalous messages. Charts for this data are presented in Figure 4.5, with the respective questions as below:

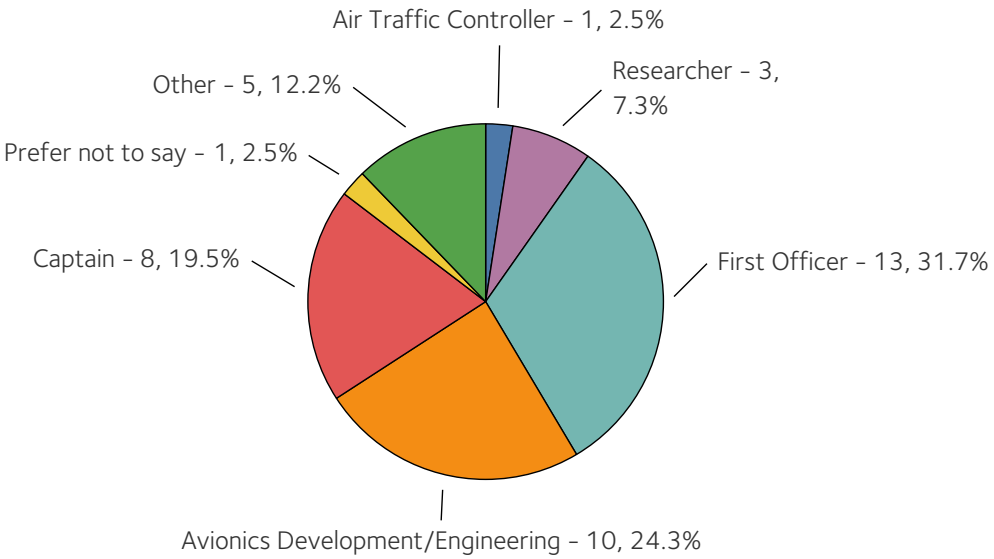

(a) Respondents by Profession

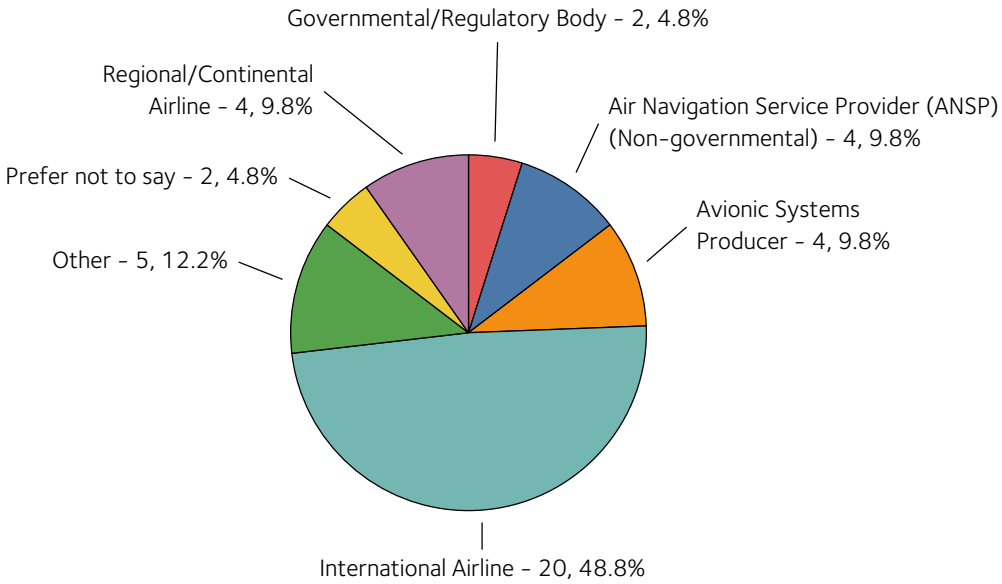

(b) Respondents by Organisation

Figure 4.4: Data link opinion survey demographics, by profession and organisation.

- Question 6: Do you have any experience of ACARS being used to share sensitive or private information? This might include personal data (e.g. names, addresses) or commercially sensitive data.

- Question 10: ACARS Message Security (AMS, also known as ARINC 823P1) is a standard for providing secure messaging over ACARS. Do you have any experiences with, or knowledge of, it being used in practice?

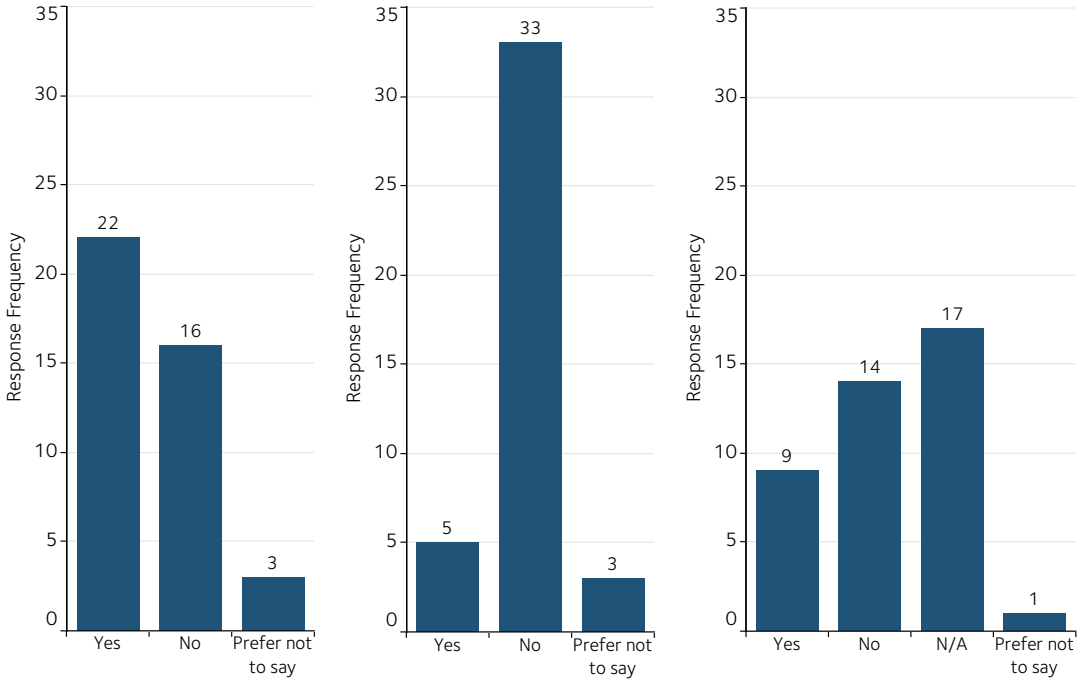

(a) Q6: Sensitive data **(b)** Q10: ACARS Message Security **(c)** Q15: Anomalous messages.

Figure 4.5: Survey responses relating to message sensitivity, security usage and anomalies. N/A reflects no response given.

- Question 15: Have you ever had experience with reception of anomalous ACARS messages?

Furthermore, in Figure 4.6 we provide responses to privacy and safety assessments to ACARS according to the following questions:

- Question 8: How suitable would you consider standard ACARS (unencrypted) to be from a safety point-of-view? In other words, to what extent do you think that ACARS is secure enough for safety-related data?

- Question 9: How suitable would you consider standard ACARS (unencrypted) to be from a privacy point-of-view? For example, for transmitting sensitive data such as names or addresses.

In Figure 4.5a, we can see that over half of respondents have experience of ACARS being used to transmit sensitive data. This is in comparison to Figure 4.5b, which shows that the vast majority, at 80%, did not have knowledge of AMS

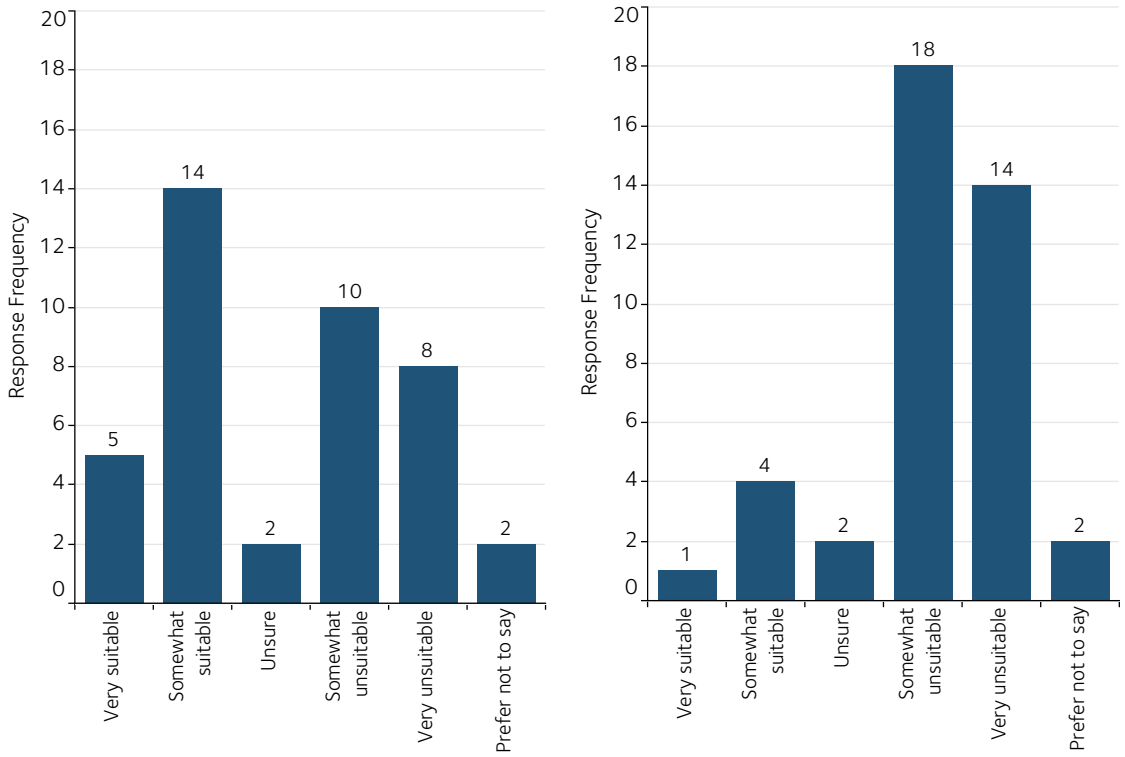

(a) Q8: Safety **(b)** Q9: Privacy

Figure 4.6: Responses to privacy and safety assessments of ACARS.

usage. Although a small sample, this indicates that although the system is known to be used for potentially sensitive data, there is limited appetite to use the standardised security solutions.

Looking at the attached free-text responses for Q6, a range of potentially sensitive information types were included, such as:

- Passenger lists, emergency health information or passport data,
- Crew information, such as roster changes,
- Maintenance and engineering information, including engine data,
- A proprietary information sharing system,
- Security warnings.

This is interesting when considered with Figure 4.6b, in which ask about ACARS suitability for private data. The vast majority of respondents (32, 78.0%)

felt the system was either *somewhat unsuitable* or *very unsuitable* for sending sensitive data over.

For safety, Figure 4.6a shows responses with regard to ACARS suitability for safety-related messages. Responses to this are more mixed, with 19 (46.3%) feeling that it is suitable for safety-related messages, and 18 (43.9%) feeling it to be unsuitable. This divide indicates that the respondents do not necessarily relate communications integrity to security mechanisms, since such mechanisms would help protect message integrity and authenticate their source—vital for safety-related tasks. However, in [148], participants rate the safety impact of ACARS to be moderate with the likelihood of attack being low. Our response distribution approximately matches to this.

In a similar vein, experience of anomalous messages (and so possible efforts to interfere with the system) appear limited. In Figure 4.5b, we see responses regarding anomalous messages received via ACARS.[9] However, in the related free-text responses, these messages are all cases of scrambled content, possibly due to link failure. Indeed, successful attempts to interfere with the system may not be spotted—we analyse this further for other systems in Chapter 8.

We use the data from this survey throughout the rest of the work, and so present it where relevant.

[9] A flaw with the survey implementation allowed respondents to ignore this question if they wished. We presume that a no response is analogous to a negative response, but mark them as N/A.

I don't mind what people say about style. If it's the quickest way to get from A to B, that's how we're going to do it.

— Chris Froome

5 Privacy in Aviation

Contents

Outside of commercial aviation, privacy is not a new concern; business aircraft owners in the United States have had a long-running campaign to protect their privacy through the National Business Aviation Association [159]. Military and state aircraft adopt a different approach, relying on their ability to use regulation to protect themselves. In this chapter, we cover the precedent for privacy in non-commercial aviation and outline the framework for privacy requirements on which our ACARS work is based.

5.1 Why is Privacy a Challenge in Aviation?

Due to the way that aviation operations have evolved, privacy does not necessarily come as naturally as safety. The ATM system does not accommodate privacy well, with lots of communications taking place on open channels (e.g. voice ATC) and identification being clearly displayed upon the aircraft. Furthermore, commercial aircraft publish extensive schedules detailing their movements which are available through many sources, such as airport websites or flight trackers. However, not all stakeholders are satisfied with such little privacy. Business, military and state aircraft operators extensively use aviation in their activities, and some desire privacy.

This has become a popular topic in recent years, with SDRs enabling many users to collect aircraft communications data easily. Especially in the case of ADS-B, operators have begun to become aware of the challenges that in-the-clear links face. Flight tracking websites such as Flightradar24, FlightAware or OpenSky Network are using SDRs at scale to collect aircraft location and identity data and making it publicly available [24, 25, 160].

A consequence of this distributed collection is that there is no single entity or method which can be used to protect privacy. For example, whilst an operator can request that their aircraft is not shown on flight trackers, they will have to do this for each different tracker.[1] Furthermore, in many cases there is no provision civil aircraft to use encrypted channels to hide.

This is not unique to ADS-B. One pre-ADS-B example of tracking at scale was presented at DEF CON 20, in which the authors used speech recognition on ATC voice feeds to extract clearances for certain privacy sensitive aircraft [53]. However, carrying this out over a wide area is resource-intensive, and some regions forbid retransmission of their ATC frequencies.[2] This is not always the case with data communications.

[1] This is known as *blocking* and is covered in Section 5.2.1.

[2] In the UK, for example, listening in to addressed communications, as voice ATC is, is illegal under the Wireless Telegraphy Act 2006 [161]. This also applies to passively observing ACARS messages.

Numerous examples of the impact of aircraft tracking have been seen in the media, some using ADS-B. A prominent case saw the authors train a machine learning classifier to identify surveillance aircraft based on their flight path [28]. For state aircraft, the Geneva Dictator Tracker has generated much attention due to its use of ADS-B to identify when aircraft from countries with authoritarian leadership—as defined by the US Central Intelligence Agency—arrive into Geneva Airport [162, 163]. In some cases, ADS-B tracking has even revealed military aircraft movements such as deploying equipment to war zones [164].

Indeed, ADS-B is only one link which a sensitive operator might wish to protect. Being a general-purpose data link, ACARS is potentially problematic since many types of data flow over it, not just location. Identifying which flows might be sensitive and preventing their use is non-trivial, and whilst simply not using it is an option, it does not appear a popular one.[3]

Whilst this is clearly a motivating example for why links carrying potentially privacy-sensitive data should be encrypted, this is not realistic in the short term due to the inability to quickly redeploy avionics. As such, operators must look to other approaches instead.

5.1.1 Stakeholder Privacy Requirements

To reason about different privacy approaches, we must first establish the privacy requirements of business, military and government stakeholders. We draw upon work in [27] and [30] to do this, and it is worth noting that a government and military's stance on privacy heavily depends on their stance on transparency. For example, some countries (e.g., Switzerland) publish military and governmental aircraft records in full [166].

Business

Privacy for business aircraft has been, and currently is, a point of contention. This is a multifaceted issue which sees shareholders and the public wanting more

[3] We know of one airline choosing not to use ACARS due to cost reasons but this makes operations more difficult [165].

Table 5.1: Summary of privacy requirements and concerns for stakeholders.

Stakeholder	Example reasons for privacy	Aircraft operator
Business	Protecting business activities e.g. mergers, acquisitions	Charter or owned by business
Military	Operational or tactical secrecy for transport or logistics	Owned by military
State	Sensitive diplomatic missions	Operated by military or national flag-carrier airline

transparency, whilst companies firmly believing that having no aviation privacy would seriously impede business. Some examples of this are provided in [27, 167], in which aircraft movements are claimed to be able to reveal mergers, acquisitions and personal executive actions which would unduly affect share price. In one case, tracking of business aircraft by manual means has been used by law enforcement as evidence of insider trading information [119].

Based on this, situations where business aircraft movements can be tracked are seen as sensitive with respect to business activity, and so not only as undesirable but potentially damaging. Since businesses are not required to make all of their operations public, the use of avionic data links for sensitive data such as location information is problematic.

Military

In many cases, military operators aim for secrecy in their movements—this is perhaps best demonstrated in the fact that operational military aircraft (e.g. fighter aircraft) use separate standards and systems for designed for operational purposes, of which little public information exists. As such, we reason that they have a strong desire for privacy in order to avoid revealing operational information to other nations, making situations which reveal sensitive information about military aircraft, such as location, are problematic.

However, armed forces do use civilian technologies such as ACARS and ADS-B for some purposes such as transport or logistics. From this we can infer that they strongly value privacy. Work in the area of ACARS was driven by US military,

specifically with the intention of providing confidentially on the data link [139]. This work led to Secure ACARS and the later AMS standard as discussed in Chapter 4. Their concerns are further supported by their stance to ADS-B. In [168], US military figures highlight that the ability to receive unencrypted signals with cheap hardware is a major concern, since it allows aircraft to be tracked.

State

State, and as a subset government, aircraft have a slight different relationship with privacy; the way in which the aircraft are operated depends on the country, with some being military run, others by the flag-carrier commercial airline and some a mix of both. For example, in the UK, Royal and UK government flights are operated by the Royal Air Force [41, 169]. Many countries have specific aircraft for official government business (e.g. Air Force One in the US) but use charter or commercial flights for smaller groups of officials.

Unlike business or military aircraft, citizens often want more transparency and accountability from state officials, particularly when travelling to foreign nations or with regard to spending taxpayer money. This could create an expectation where state and government aircraft on regular activities should not expect to have privacy by default. Summits such as Bilderberg Meetings and the World Economic Forum make attendee lists available on the Internet, thus making concealing travel to them ineffective [170, 171].

Even with such an expectation of accountability, situations exist where state aircraft might desire privacy, such as on sensitive diplomatic missions where exposing attendance could unduly reveal intentions. Although situations exist where it can be argued that state aircraft should be able to operate privately, doing so may be contentious. An example of state attempts to obscure aircraft is seen in the 2016 attempted coup d'état in Turkey. President Erdoğan allegedly attempted to evade coup leaders by hiding as a civilian aircraft in Turkish airspace [172]. Even so, we note that our data indicates high levels of aircraft blocks on state or government aircraft.

5.2 Current Approaches to Privacy

Since many civilian avionic communications do not have security measures which can be used to protect message confidentiality, other methods can be used for privacy. The most established approach is that of *aircraft blocking*, which involves restricting the public sharing of data about an aircraft. We use the concept of blocking as an indication of privacy in Chapters 6 and 7, so consider it in this section, alongside other privacy approaches.

5.2.1 Aircraft Blocking

A widespread approach to privacy for aircraft is through blocking schemes. Fundamentally, these seek to limit the sharing of data about a particular aircraft such that they cannot be easily tracked. These schemes have been in existence since at least 2000, but have become much more prevalent since ADS-B and flight trackers such as Flightradar24 or FlightAware expanded [159]. Business aircraft operators in particular have expressed concern about how easy it is to receive ADS-B signals with cheap hardware, and how this allows them to be tracked [26].

Blocking in the United States

The most established blocking mechanism concerns and FAA data feed called the Aircraft Situation Display to Industry (ASDI) which provides information on the aircraft observed by the FAA through their ATC capability. Until recently, the National Business Aviation Association (NBAA) ran a programme called Blocked Aircraft Registration Request (BARR) which allowed aircraft owners to obscure themselves from this feed [173]. Submitting a BARR request would allow an aircraft owner to either prevent their data reaching a subscriber (source-level block) or allow it to reach a subscriber but not be shared publicly (subscriber-level block). The latter would allow operators to still use flight trackers, for example, but not have their data shared publicly. ASDI has been decommissioned but the feed and BARR program are instead part of the new Traffic Flow Management Systems (TFMS) [174, 175].

Aircraft owners can submit directly to the FAA for a block, who administrate the list for the new feed. Until recently, the comprehensive list of blocked aircraft was only available to TFMS subscribers and thus private—a Freedom of Information request made public a version of this list from March 2017 [176]. At the time of writing, BARR is the most prominent public blocking scheme with government cooperation. Outside of the US, little evidence of blocking for business aircraft at a national scale exists.

Blocking outside of the United States

For aircraft registered outside of the US, blocking is instead reliant on contacting individual flight trackers. FlightAware block non-commercial non-US aircraft by default, and offer a fee-paid service through which said aircraft owners can choose to privately view tracking data [177]. They explicitly state that they do not track military aircraft, using blocking of US presidential fleet movements as an example, but been subject to pressure from military and law enforcement/government aircraft owners [154].

Flightradar24 has some public information about how to block aircraft and offer a fee-paid express blocking service [178]. They also note that some classes of aircraft are blocked by default. Specifically, they state:

> Fourth, the aircraft might be blocked. By default we block sensitive aircraft such as those associated with military or governmental operations. For private aircraft, we can block flight tracking information at the owners request.
>
> —*Flightradar24 on blocking* [179].

Some flight trackers take a different approach to this, most notably ADS-B Exchange [180]. This tracker explicitly does not comply with block requests and offers unfiltered data to the public.

Effectiveness of Blocking

Blocking is ultimately a form of obfuscation; whilst it does not prevent the surveillance signals from being transmitted from the aircraft, it places the burden

on a potential threat to identify them. It also makes it harder for such a threat to track an aircraft outside of its radio horizon without the threat actor running their own multiple-location collection operation. This is because surveillance messages reveal where the aircraft currently is, so once the aircraft flies beyond the collection range of a given sensor it cannot be tracked further unless another sensor is in range.

This approach has seen attention from the security community in the past. As described previously, at DEF CON 20 researchers used voice processing on ATC feeds from US airports to attempt to identify where blocked aircraft were flying from and to [53].

Indeed, blocking a legal right only in the case of FAA data feeds; if data is collected through other networks then FAA block lists will not apply. In this case, operators will have to contact individual networks to request blocks, who can chose to comply. This approach was far more effective when few streams of aviation data were available. Now that many flight trackers exist—as well as it being trivial for hobbyists to collect surveillance data on their own—blocking is less comprehensive.

Stakeholder Usage

Blocking is applicable to business, military and state aircraft. The previously discussed BARR programme was developed in conjunction with the US National Business Aviation Association, and so was originally intended for business aircraft [159].

Within most of the described blocking mechanisms, military and state aircraft appear to be granted a presumed right to privacy. As highlighted in [173, 181], governments interact with aviation agencies to filter out sensitive aircraft from official streams before the data is made available to third-parties. This demonstrates a clear desire for operational privacy from both military and state aircraft operators.

The NATS Airspace Explorer provides an example of collaboration on flight blocking with governmental agencies. The tool does not show a number of sensitive aircraft agreed with the UK Centre for Protection of National Infrastructure (CPNI) [181]. Similar statements have been made with regard to FAA feeds [182].

Blocking as an Indicator for Privacy Requirement

It could be argued that blocks are not privacy measures since they are low effort, or since the aircraft still transmit sensitive information regardless. However, the way in which blocking was devised and is used today suggests that is indeed an indicator that an operator is seeking privacy.

In [173], the authors conduct a legislative assessment of the BARR program and its history. They favour the existence of such schemes, arguing that whilst the list of those using the BARR program is too coarse to be harmful, detailed aircraft movements and history as provided by flight trackers does have the potential for harm. Furthermore, they argue that the United States Congress has supported privacy in aviation by repeatedly enacting laws which allow blocking to take place [173].

Furthermore, the use of fee-paid schemes to block aircraft suggests that owners are willing to pay to protect their privacy as much as possible. One such service is `FltPlan.com`, through which operators can pay to use a callsign rather than their own registration [183]. These callsigns have FAA blocks in place already so are effectively prepared for private operation [184].

5.2.2 Other Approaches to Privacy

Whilst blocking is an established approach, other methods exist. We briefly cover these, and they are detailed in [27].

Use of Commercial Aircraft

One approach to protect privacy is to 'hide amongst the crowd' and use flights on commercial aircraft to avoid being identifiable by private aircraft usage. Although commercial aircraft have public schedules, passenger lists are not publicly available. Furthermore, the volume of commercial flights lends itself to creating enough noise that tracking a person of interest becomes very difficult. This approach is likely to be cheaper than operating a private jet, though can be less convenient since flights will be between larger aircraft. Some governments choose to use commercial flights as a matter of policy, due to the high cost of maintaining a governmental fleet [185].

Use of Shared/Charter Aircraft

Instead of using a private aircraft tied to an organisation, shared or charter aircraft could be used. This is a lesser extent of the above approach of hiding in a crowd, since it may be difficult to identify when specific groups are using the aircraft. Maintaining some element of privacy over time might require use of different charter companies. The cost-effectiveness of this approach depends on the frequency of use—regular charter flights might eventually cost more than owning a private aircraft. Furthermore, the users of the aircraft will have less control over how avionic systems are used. However, new schemes such as 'all-you-can-eat' subscription-based charter models are emerging, which may offer more flexibility to users [186].

Avoiding Exposing Technology

In certain cases, some stakeholders can choose to switch off or not use technologies which might expose their location or other sensitive information whilst in civilian airspace. Not all aircraft are permitted to do this under IFR flight or at higher altitudes, usually limited to state or military aircraft [40, 187]. For technologies which do not have provision for confidentiality (e.g. ADS-B and ACARS), this is the surest way to avoid data being collected by potential threats. Some transponders such as those manufactured by Becker Avionics are being developed with the option to disable Mode S, ADS-B or both transmissions during flight [26].

This does not come without downsides. As ATM uses new technologies such as ADS-B more heavily, switching off these systems increases the workload for ATC. It can also reduce situational awareness since other aircraft may rely on these signals to locate the aircraft, for example in the case of TCAS.

Should this approach be taken, the operator of the aircraft must ensure that they limit other sources of data coming from the aircraft too. If not, they may negate the benefit of switching off systems in the first place.

Pseudonym Schemes

One technical approach to privacy which is standardised is that of transponder pseudonyms, though this is only possible under VFR conditions and lower than 18,000 ft in the United States. If fitted with a UAT-based Mode S transponder, these aircraft can enable 'anonymous mode' which generates a randomised, temporary ICAO address. This is transmitted instead of the real address of the aircraft [188, 189]. Work in the area suggests that these random identifiers can be decorrelated simply by tracking the aircraft before it switches into anonymous mode [188].

As with opting to not use surveillance signals, using the UAT anonymous mode has drawbacks. Whilst in this mode, the aircraft cannot be flying on a filed flight plan or use ATC services [190], somewhat limiting the applicability of this approach.

5.3 Aircraft Privacy in Practice

Having considered a range of privacy measures, we now look at one as deployed in practice. Aircraft blocking is a good metric to use since it requires some positive interaction from the aircraft operator, and is measurable. In this section we provide an overview of the extent of blocking for aircraft observed in our ACARS data collection.

5.3.1 Measuring Blocking

In order to measure how many aircraft are blocked and their category, we rely on third-party data sources and flight tracking websites themselves.

U.S. Registered Aircraft

Since these can request an FAA block, they may appear on the FAA-maintained blocked aircraft list, as described above [176]. If an aircraft in our ACARS data sets appears on these lists, it has an FAA block. Note that military and state aircraft are automatically blocked so do not appear.

Non-U.S. Registered Aircraft

Outside of the U.S., governments and flight trackers do not have publicised block lists. In order to check blocks for these aircraft we instead infer from flight tracking websites. As the most prominent flight trackers, we focus on Flightradar24 and FlightAware.

FlightAware. In this case, the website will either explicitly state that an aircraft is blocked, or it will state that it does not have information on an aircraft. Since we know that the aircraft exists and FlightAware has good coverage in our collection regions, we can infer that it is blocked.

Flightradar24. Similar to FlightAware, the website will show either limited or no information for blocked aircraft, allowing us to infer when an aircraft is blocked. On top of this, Junzi Sun maintains the World Aircraft Database which is gathered from Flightradar24 [155]. This collection has been running since 2015 so if an aircraft does not appear in the database, it is likely to be blocked.

5.3.2 Blocked Aircraft Statistics

In Table 5.2 we provide a breakdown by aircraft category and the type of block implemented, as well as the ACARS subnetwork they were observed on. Unique blocks are where an aircraft is blocked on one of the three links. Note that FAA blocks are significantly less frequent than Flightradar24 or FlightAware blocks; this is highly likely to be due to the European collection location, thus observing fewer US-registered aircraft. We include the figures as a reference, due to FAA being a well-established governmental scheme.

Business

Considering Table 5.2a, we see heavy usage across all three subnetworks of ACARS, though few aircraft transmitted on all three links. Flightradar24 saw the highest level of blocks, with 90.7% of aircraft having some level of blocking in place, and $90 \pm 2\%$ of aircraft blocked across all three subnetworks. Just over 40% were

Table 5.2: Number of blocked aircraft (AC) from each category observed on each subnetwork, by type of block. Note that the 'unique' row and column is the number of unique aircraft with each type of block and the number of unique blocked aircraft on each subnetwork respectively, e.g. an aircraft might appear on POA and VDL2 but is counted as one unique aircraft. Percentages for each figure are of 'all aircraft' for that row.

(a) Business aircraft

Subnwk.	All AC	Flightradar24		FlightAware		FAA		Unique Blocks	
		# AC	%	# AC	%	# AC	%	# AC	%
POA	1168	1033	88.4	408	34.9	128	11.0	1050	90.1
VDL2	1035	942	91.0	386	37.3	241	23.4	952	92.3
SATCOM	970	882	90.9	426	43.9	299	30.8	890	91.8
Unique AC	2366	2147	90.7	951	40.2	514	21.7	2173	91.8

(b) Military aircraft

Subnwk.	All AC	Flightradar24		FlightAware		FAA		Unique Blocks	
		# AC	%	# AC	%	# AC	%	# AC	%
POA	29	20	69.0	6	20.7	0	0.0	22	75.9
VDL2	25	19	76.0	10	40.0	0	0.0	20	80.0
SATCOM	418	402	96.2	357	85.4	1	0.2	405	96.9
Unique AC	438	416	95.0	365	83.3	1	0.2	420	95.9

(c) State aircraft

Subnwk.	All AC	Flightradar24		FlightAware		FAA		Unique Blocks	
		# AC	%	# AC	%	# AC	%	# AC	%
POA	84	43	51.2	28	33.3	11	13.1	49	58.3
VDL2	72	33	45.8	22	30.6	8	11.1	40	55.6
SATCOM	119	66	55.5	43	36.1	11	9.2	80	67.2
Unique AC	183	97	53.0	63	34.5	19	10.4	116	63.4

blocked on FlightAware, with a wider range at 9 percentage points. The cause of the gap between these two flight trackers is unclear, but reasons could include relative popularity or that Flightradar24 is a European company thus causing European operators to seek blocking with it first.

Looking at the FAA blocks, we see that 21.7% of business aircraft have one in place; the lower percentage likely due to our European collection location, thus seeing few US registered aircraft.

Military

Aircraft operated by the military primarily were seen on the SATCOM subnetwork, with little activity on POA or VDL2 ACARS. As we will see in Chapter 6, many military aircraft we observed were perform oceanic crossings, so may be more likely to use satellite communications.

In terms of blocks, Table 5.2b shows that Flightradar24 and FlightAware blocks are very common for military. For the former, 95.0% of aircraft have a block in place, the majority of which come from the satellite subnetwork. On FlightAware, the lower levels of blocking for POA and VDL2 are heavily outweighed by 85.4% of aircraft using SATCOM having blocks. Very few aircraft in this category have FAA blocks, which is likely due to some of these aircraft being redacted by default, as implied by the FAA and Flightradar24 statements (see pg. 71) [179, 182].

State

Although sometimes sharing characteristics and operators with their respective military, state aircraft have a more even split of subnetworks, with only a slight bias to SATCOM. However, as Table 5.2c shows, rates of blocking are much lower than business and military aircraft, with Flightradar24 having the highest rate at 53.0%. Across the subnetworks this varied by 9.7 percentage points. State aircraft blocking on FlightAware was lower at 34.4%, with a range of 5.5 percentage points. Unlike military, 10.4% of this category had an FAA block in place, which was biased in the favour of the aircraft seen using POA.

Lower levels of blocking across the measures could be attributed to governments taking an approach of being more transparent, thus not seeking to hide their movements to the same extent. Another reason might be that as described above, some state aircraft are operated by the military so could fall into that category instead.

5.4 Summary

In this chapter we have focused on the privacy issues arising from a lack of confidentiality protection on surveillance technologies, namely ADS-B, and how operators can try to mitigate them. Having considered the different approaches to privacy for surveillance technologies, we have established that operators do try to obscure their activity through blocking. When analysing the status of aircraft seen in our ACARS data collection, we can see that business and military aircraft extensively block themselves from flight trackers, with some trackers having over 90% block rates. State aircraft are less prolific in this regard, but still see at least a third of the observed aircraft blocked on public trackers.

There are too many factors you have to take into account that you have no control over... The most important factor you can keep in your own hands is yourself. I always placed the greatest emphasis on that.

— Eddy Merckx

6

Implications for Privacy of Using Unsecured Data Links

Contents

Having considered how aircraft can seek privacy against surveillance signals, we now look at how ACARS provides an additional channel to leak sensitive location

data. The wide variety of ways in which ACARS is used lends itself to misuse, with sent messages undermining attempts to protect privacy through blocking. In this chapter, we perform a measurement study on the extent to which ACARS can leak in this way for aircraft which are otherwise seeking privacy through blocks.

Concerns about the clear-text nature of ACARS have been highlighted by individuals within the aviation community as far back as 1998, but primarily from a security angle [191]. Recently, an ex-pilot discussed his view of ACARS usage from the cockpit, acknowledging the eavesdropping threat and providing anecdotes of messages circulated widely despite being intended for a narrow group of people [192]. Assessments from the Airline Pilot's Association and the US Air Force assert that message injection with false information is a realistic threat, due to the lack of security [193, 194].

Our findings show that there is a significant leakage of privacy-sensitive data on the ACARS channel. Messages which initially appear innocuous can leak significant amounts of information, especially for stakeholders who otherwise demonstrate a desire for privacy. Usually, these actors try to hide flight information—ACARS messages regularly undermine this effort. This is in spite of the results from our survey of aviation industry professionals which found that 77.5% of respondents do not find ACARS suitable for private data.

6.1 Threat Model

For this work, we take the stance of the passive listener as defined in Chapter 3; they are collecting data from one or more locations on POA, VDL2 and SATCOM using commodity equipment with the intent of analysing it for sensitive information. As per our collection setup in Chapter 4, this involves equipment costing £100-200 per subnetwork. Since we are focussing on non-commercial aircraft, our threat is aiming to uncover location information which might give them an advantage. This might be inferring confidential meetings or movement of personnel.

6.2 Method

To investigate this, we use our datasets collected across three links as described in Chapter 4. Having established the extent of blocking in Chapter 5—namely considering an aircraft *blocked* if it has a block on one of the three sources—we now look to understand whether ACARS provides an extra channel which reveals location data on top of existing surveillance signals.

The wide range of ways in which ACARS can be used gives rise to message types which could reveal location. We focussed our search by looking for messages which fitted into one of the following categories:

- **Inferred location**, revealing the area an aircraft is in or where it is travelling to, which is usually tied to an airport,
- **Current location**, providing the location of an aircraft at a given time, usually specifically,
- **Route information**, which includes detailed descriptions of where the aircraft has flown from and where it will fly to.

Typically, this data is shared with an aircraft operator or company over data link since surveillance signals are available to ATC but not necessarily other parties. Furthermore, route information is often transferred to share flight plans without paper copies or to update the route.

Crucially, if an aircraft has flight tracker blocks in place yet transmits messages in the above categories via unencrypted ACARS, they stand to leak not only current location but past and future locations too in a single message. We now consider how to identify these messages.

Inferred Location. The least specific type of location information is tied to an area rather than a specific location. This will come in the form of messages containing airport codes or waypoints. For airport codes we look for ICAO codes which are four characters long; the first one or two of which will identify the

country. Waypoints define some geographical location, so can be referred to by a coordinate, distance and radial from a VOR/DME station (see [195]) or a five letter identifier assigned by a local aviation authority.

Current Location. Information in this category is usually in the form of coordinates. Format varies depending on the sender, since some will explicitly specify the longitude and latitude components whereas others will rely on the format to reveal this. In this area, we are particularly interested in either messages which contain coordinates in plaintext or via encoded ADS-C messages. The former can be retrieved with regular expressions though format differences make full coverage difficult. For the latter, the ADS-C decoder from JAERO was repurposed to check for positions both in the form of waypoints and coordinates.[1] These messages can be identified by `ADS` occurring before an encoded blob.

Route Information. One of the most comprehensive forms of information sent over ACARS is the route of a flight. These vary by operator but typically consist of waypoints, airport codes, coordinates and navigational fixes. Flight plans are the main source of this data, which can be retrieved by searching for variants of `FPLAN`, the mnemonic used for flight plans.

6.2.1 Comparison with ADS-B

Many of the aircraft seen in our dataset also transmit ADS-B, as observed by the OpenSky Network and outlined in Table 4.3 (pg. 58). Where relevant, we compare the data given by ACARS to the ADS-B or Mode S data available from OpenSky. We explicitly state where we use this method.

6.2.2 Scale of Collection

Whilst we collect in a single location it is important to note the relative ease with which this can be scaled up. Since ACARS can be collected with similar—or sometimes the same—sensors used for networks such as OpenSky or Flightradar24,

[1] The code for this can be found at `https://github.com/tangohead/ADS_CDec`.

adapting these for use with ACARS would be relatively inexpensive. As such, we believe there are no major technical barriers to wide-scale ACARS collection networks emerging. Some initial work towards this can be seen by AvDelphi who offer a public ACARS stream and appear to be moving towards establishing an ACARS collection network [158]. Should ACARS coverage widen in this way, the privacy consequence of sensitive messages transmitted over the link is magnified as the chance of a passive observer contributing the message to a public store is much greater.

6.3 Measuring Sensitive Information

Using the described categories of data, we now explore the extent of this information being used by blocked aircraft as defined in Chapter 5. We further assess the potential loss of privacy breaches caused by using unsecured communications. During our data collection, we observed 2987 non-commercial, identifiable aircraft, 2709 (90.7%) of which were blocked.

6.3.1 Position Reporting

One of the many purposes of ACARS includes position reporting. Although ADS-B is slowly becoming compulsory worldwide, ACARS-based position reporting is also widely used by business, commercial, military and state aircraft alike. This allows airlines or third-parties to provide services based on location data which they would otherwise not necessarily have access to. Despite the ARINC 620 standard defining some methods to do this, a range of ways beyond standardised approaches are also used to share position [136].

Position reporting via ACARS presents a challenge to aircraft which are trying to avoid being detected due to transmissions of other systems, such as ADS-B. The lack of confidentiality protection on the ACARS link means that messages of this type sent over it create a location privacy issue. We look at two types of report: Automatic Dependent Surveillance-Contract (ADS-C) and text-based.

Table 6.1: Statistics of positional report transmission by blocked aircraft (AC) on POA. Percentages are of all blocked aircraft of that type.

Aircraft Category	All Aircraft		Blocked Aircraft Only			
	All	Blocked	Position Reports	# Aircraft Reporting	% All Blocked	Reports per AC
Business	1168	1050	4828	516	49.1	9.4
Military	29	22	89	5	22.7	17.8
State	84	49	83	18	36.7	4.6
Total	1281	1121	5000	539	48.1	9.3

ADS-C

As described in Chapter 2, ADS-C is a method for surveillance over data link in which an ATC centre will establish a surveillance contract with an aircraft. This contract defines the regularity of reporting and the data which should be contained within [102]. These messages are encoded inside a data link message and so may contain more information than other location report, such as those by ADS-B, due to the larger message size available to ACARS. Contents might include imminent waypoints, notification of route changes and emergency events. Transmitting ADS-C messages on an unprotected link can reveal a lot of routing data; for aircraft which wish to hide, it becomes a privacy issue.

Sensitive aircraft usage of ADS-C was mostly confined to the military on SATCOM, with 104 blocked aircraft (25.7% of all military blocked on SATCOM) sending 1062 messages. As we only observe SATCOM uplink, we can only see the ground-to-air side of ADS-C, e.g. requesting reports. However, with this we can determine which aircraft use the system, and so are revealing their location through the data link. On average, an aircraft using this system sent 10 messages, with one US Air Force aircraft receiving 171 messages alone, indicating heavy use of ADS-C on downlink. On top of this, ADS-C can be transmitted outside of ACARS, directly on the SATCOM and VDL2. Since we focus on ACARS, we do not consider this here—however, our numbers provide at least a baseline for usage on other parts of the links.

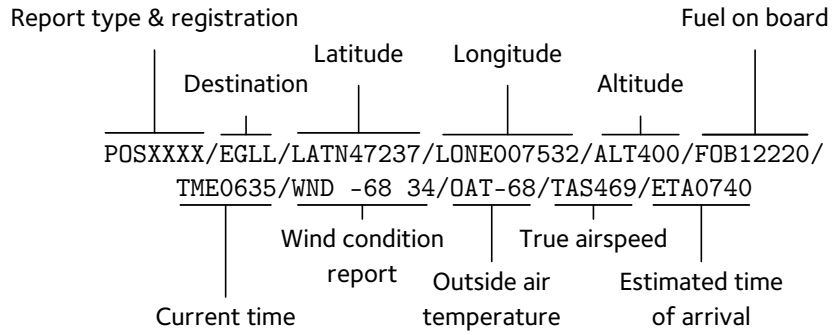

Figure 6.1: Labelled example of a text-based position report sent over VDL2. Identifying information has been omitted.

Text-Based Position Reports

Whilst ADS-B and ADS-C provide well-defined systems for transmitting location, our data shows that many aircraft send positional reports in custom text-based formats. Often, they simply send coordinates as part of a message, along with a timestamp—a more complex example can be seen in Figure 6.1.

This approach is heavily used on the POA link above all others, as illustrated in Table 6.1. Slightly over 49% of blocked business, 22% military and 36% state aircraft sent text-based position reports during our data collection. Furthermore, aircraft within these groups on average sent 9.4, 17.8 and 4.6 messages each, respectively. The situation for some business aircraft is worse on SATCOM; although only 13.9% of blocked aircraft send this type of message, they send 12.8 messages each on average. A plot of these reports can be seen in Figure 6.2. We observe a nominal POA reception range of 400 km, with three outliers at 430, 440 and 460 km.

Clearly, this is a problem for business aircraft. Any notion of obscuring oneself from a flight tracker is undermined by sending ACARS position reports in the clear. Although VHF-based ACARS might have a similar range to ADS-B, it is yet another clear text transmission of sensitive data for blocked aircraft and in some cases contains more than position, but also information on location intention or history. SATCOM presents an additional challenge to VHF; its effective reception range is larger than VHF based links. Since a given satellite serves a large geographic area—larger than the radio horizon—intercepting part of the link gives access to

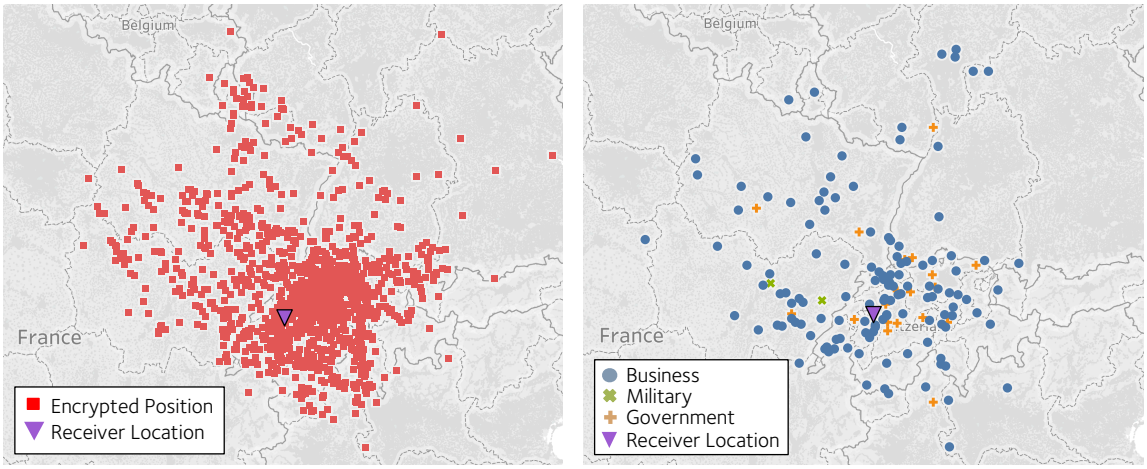

(a) Encrypted position reports. **(b)** Unencrypted position reports by aircraft type.

Figure 6.2: Plot of encrypted and in-the-clear position reports across POA and VDL2 links, with in-the-clear messages split by business, government and military aircraft. Encrypted positional reports are explained in detail in Chapter 7.

data sent to this wide area. As such, we see data in our SATCOM collection considerably outside of our VHF range.

6.3.2 Use of Information Services

Information services, such as the Digital-Aerodrome Terminal Information Service (D-ATIS, or simply ATIS), reveal a great deal about aircraft intention. ATIS reports are used by pilots to reduce ATC loading and VHF congestion [196], containing information about an airport such as runway condition and activity, weather, Notices to Airmen (NOTAMs) amongst other safety-related data. Thus, requesting ATIS information is an intention to land at a given airport, which implies the destination and approach timing of an aircraft. Table 6.2 demonstrates that ATIS is used by blocked aircraft across all links. Whilst usage is confined to a small section of the blocked aircraft for each category, the level is consistent across the categories and those using it typically to send several messages.

SATCOM saw heavier use, with close to 19% of all blocked business aircraft and 20% of both blocked military and state aircraft using ATIS. Business and military aircraft sent three messages each on average, whereas government sent over five. Although a normal part of aviation operation—and an important frequency

Table 6.2: Number of blocked business, military and government aircraft (AC) using ATIS requests/responses, with aggregated across all ACARS subnetworks.

Aircraft Category	All Aircraft		Blocked Aircraft Only			
	All	Blocked	ATIS Messages	# Aircraft Reporting	% All Blocked	Messages per AC
Business	2366	2173	1210	322	14.8	3.8
Military	438	420	314	84	20.0	3.7
State	183	116	115	22	19.0	5.2
All	2987	2709	1639	428	15.8	3.8

Table 6.3: Number of blocked business, military and government aircraft (AC) using messages encrypted with a proprietary cipher across all ACARS subnetworks.

Aircraft Category	All Aircraft		Blocked Aircraft Only			
	All	Blocked	Encrypted Messages	# Aircraft Reporting	% All Blocked	Messages per AC
Business	2366	2173	3288	294	13.5	11.2
Military	438	420	75	11	2.6	6.8
State	183	116	44	14	12.1	3.1
All	2987	2709	3407	319	11.8	10.7

congestion reducer—it allows a listener to determine where the aircraft is landing. Since this is mostly via SATCOM, the listener need not be geographically close. Because of this, it reveals intention information which would otherwise be hard to gather at this scale with a single sensor. Again, this is a problem for those hiding from public flight tracking.

6.3.3 Use of Proprietary Encryption

Proprietary encryption is regularly used by business-type aircraft, with our data demonstrating that these messages are sent over all three observed links. As a separate investigation into ACARS, we analysed this encryption of which details are provided in Chapter 7. We will cover the relevance to privacy below. One of the common forms of these messages, when decrypted, reveals much about the aircraft's route. On top of current location, they also reveal origin, destination

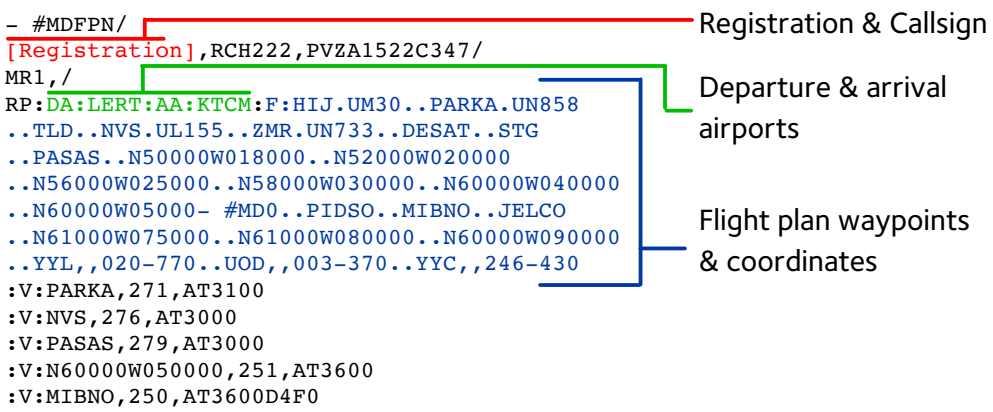

Figure 6.3: Labelled example flight plan transmitted over SATCOM ACARS. Sensitive information has been omitted.

and estimated time of arrival. Furthermore, the link is used for other message types which can be problematic for privacy.

We observed 319 blocked aircraft (11.8% of all blocked) using this cipher across all links, sending 3407 messages, as shown in Table 6.3. Whilst this is a low proportion of blocked aircraft, the few which use it do so heavily. Government aircraft had the lowest average at 3.1 messages per aircraft, with military at 6.8 and business at 11.2.

Of the VHF encrypted messages 80.3% originated from blocked aircraft, and 44.6% of all messages (including unblocked aircraft) were positional reports as above. On SATCOM, aircraft leaked in a different way, with flight plans, information services and contact information being used extensively. Once again, the blocked aircraft otherwise obscure their actions, and use a cipher to further obscure the data. The fact that these messages are trivial to decrypt and contain aircraft intention information constitutes a clear privacy issue.

6.3.4 Flight Plans

SATCOM transmission of flight plans was used by 69% of blocked military aircraft; with 405 of the 418 military aircraft on the link are blocked, the majority of this group of aircraft engage in this. The average observed aircraft sent around four flight plans, with the format of these messages varying. We provide a labelled military flight plan in Figure 6.3.

Typically, the message content is used to transfer flight plan data to the aircraft. Like clearances, this involves the departure and arrival airports and in some cases routes and call signs. Call signs can indicate the type of flight—for example, `RCH` is a 'reach' flight, which is a troop transport. Since military aircraft have a degree of operational sensitivity, the fact that these messages are sent in clear text at all is a problem. Additionally, a significant proportion of these aircraft have both a Flightradar24 and FlightAware block in place, indicating that there is some active, rather than just presumed, attempt at privacy. We look at military usage of flight plans further in Section 6.4.2.

Although not directly related, some military aircraft on SATCOM make use of free text messages via ACARS for operations or logistics. We observed 115 blocked military aircraft sending 630 messages in this way. Most of these include flight operational content relating to cargo or estimated arrival times, with some revealing destinations or route adjustments.

6.3.5 Undermining Aircraft Blocks

Clearly, use of ACARS by business, military and state aircraft poses a significant threat to privacy. By listening to the ACARS subnetworks in a single location, messages worldwide can be collected, used to track movements and reveal intentions. Although some aircraft might not consider this a problem, a significant proportion of blocked aircraft (which 90.7% of the non-commercial aircraft observed were) use clear text messages—or an easily breakable cipher—to transfer information which they otherwise try to obscure from public knowledge.

It is important to compare data obtained from ACARS with similar data gathered by collecting airport landing logs or listening to ATC voice channels to understand the additional impact. Gathering data from landing logs or flight plans involves collecting copies from ANSPs and processing them to a structured format. To cover numerous countries would incur a lot of manual work; on top of this, some countries or airports may not make all of their logs public. ACARS provides more structured, consistent data stream than this—having a formal message structure

makes it easier to process at scale. Furthermore, the structure is the same, or very similar, worldwide, meaning low-cost sensors could be geographically distributed to significantly scale up collection.

For ATC communications, language processing could be used with multiple receivers tuned to different frequencies, as shown in [53]. This can provide granular information including turn headings and airspeeds of aircraft. However, aircraft regularly change frequency as they move between different areas, thus adding risk of tracking loss. On top of this, voice channel quality is variable and congested, so not all phrases can be recovered. Although ACARS may provide less detail, it avoids these challenges. It also offers operational information which would not appear on ATC voice channels.

For these reasons, we feel that ACARS usage by blocked aircraft poses an additional and unique risk to said aircraft privacy beyond the accepted—or known—risk.

6.4 Case Studies

With the various ways in which ACARS usage can cause problems explained, we now expand upon this with three case studies. For each aircraft category we show how sending one or more of the message types above can reveal a significant information about a flight; in each case, this is far more than ADS-B alone, and than what would be collectable manually.

6.4.1 Business Aircraft Case Study

Many business aircraft reveal their locations or destinations with positional reports. However, some aircraft reveal even more through flight plans and ATIS reports. In this case study, we look at a business jet owned by a Saudi Arabian company and operated by a British private aviation firm. This aircraft regularly transmits on SATCOM, sending 118 messages of which 50 were encrypted with a monoalphabetic cipher (see Chapter 7), 10 were ATIS and 10 were flight plans.[2] This aircraft

[2]The encrypted messages contained weather reports, some free text discussion about ETA and requests for flight plans.

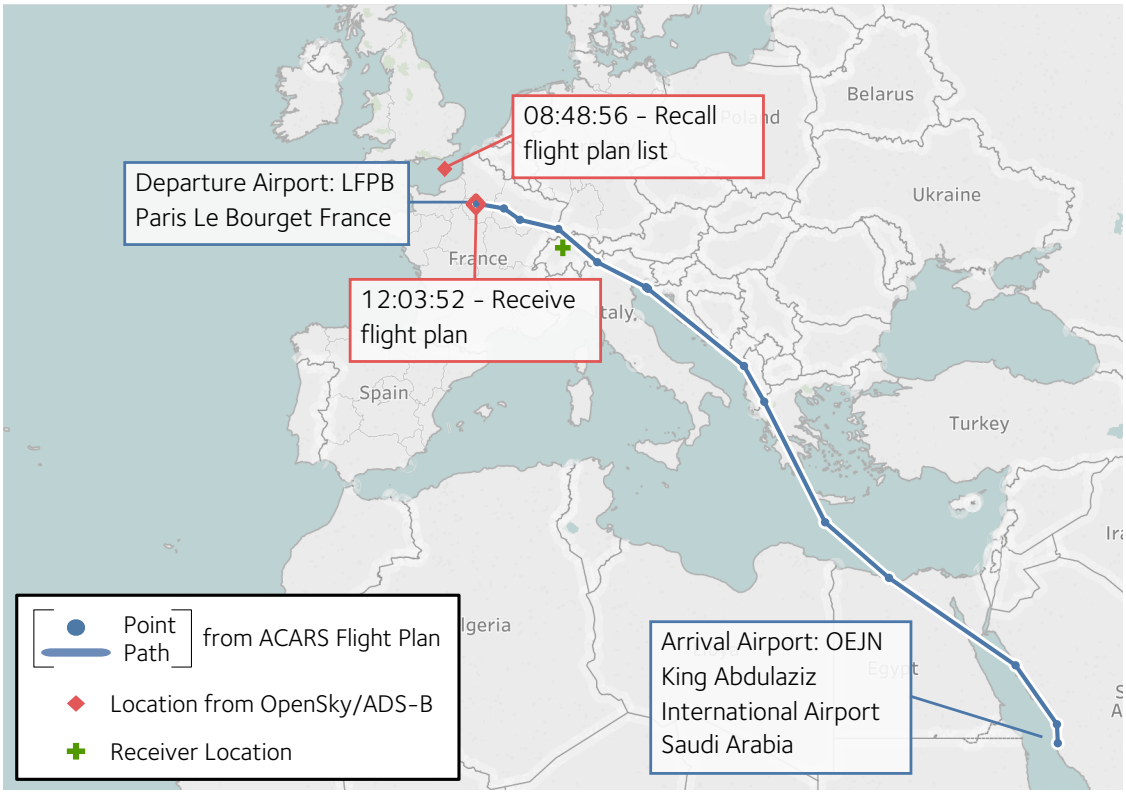

Figure 6.4: Flight plan of a blocked business aircraft from SATCOM ACARS messages. Red shapes indicate position obtained from OpenSky ADS-B messages. Blue shapes constructed from the ACARS flight plan.

has a BARR source block and is blocked on both flight trackers, Flightradar24 and FlightAware.

In Figure 6.4 we see the reconstruction of one day of ACARS traffic from this aircraft, consisting of 20 messages. This data is cross-referenced with OpenSky, a collaborative sensor network collecting ADS-B and Mode S data for use by researchers [153]. Whilst flying to the departure airport, it is sent the list of available flight plans, then whilst on approach to Paris, is sent the flight plan for the flight Paris to Jeddah. Using the flight plan we can construct the main waypoints along the route. It follows this message with some considerably more detailed route information over the course of three messages. Two hours later, the aircraft resumes sending messages, though this time no ADS-B data could be retrieved. During this period, it sends 10 encrypted messages and uses ATIS to check the landing conditions at the arrival airport.

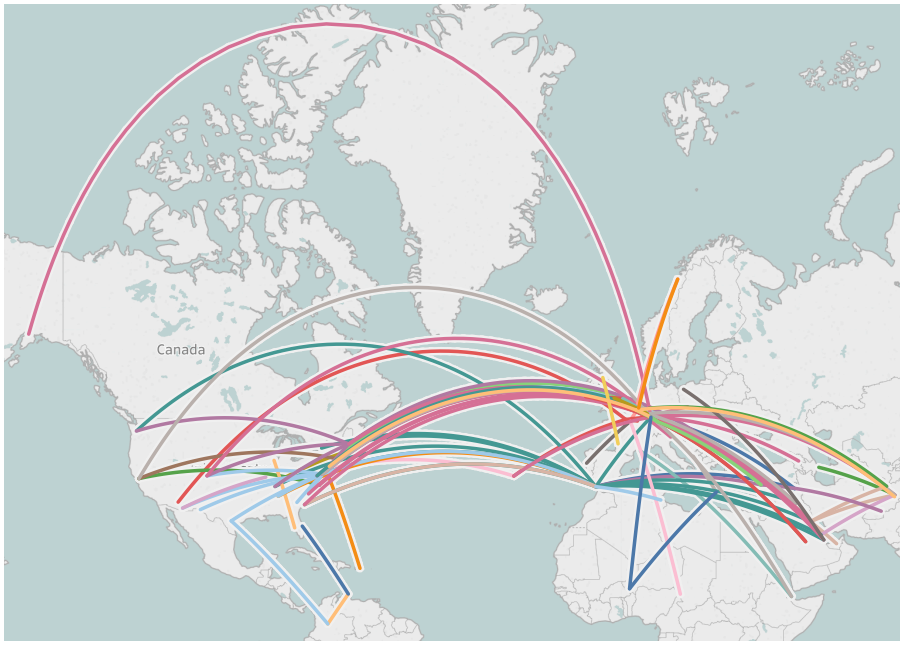

Figure 6.5: Interpolated flight tracks of blocked military aircraft from departure/arrival airport pairings recovered from flight plans sent over the SATCOM ACARS subnetwork. Colours are assigned by departure airport.

Clearly, this is a significant amount of information recovered from relatively few messages, especially for a blocked aircraft. A similar process could be followed for the other flight plans and exchanges; considering that the aircraft is apparently blocked on all public flight trackers, this is a significant leak.

6.4.2 Military Aircraft Case Study

Military aircraft are permitted to turn ADS-B off in order to provide privacy with respect to their location [40]. As shown in Table 5.2 (pg. 77), 416 (95.0%) of these aircraft had a Flightradar24 block, with 365 (83.3%) having a FlightAware block. We can see there is a real desire for privacy, as many of these aircraft were 'unknown' to the trackers, i.e. the websites claimed not to recognise them.

This category made heavy use of ACARS, with the average aircraft sending three messages in the collection range. Some reveal much more than others, however. Over our collection phase we received over 1206 flight plans transmitted by 280 blocked military aircraft via SATCOM. We were able to reconstruct the flight path for most of them, as shown in Figure 6.5.

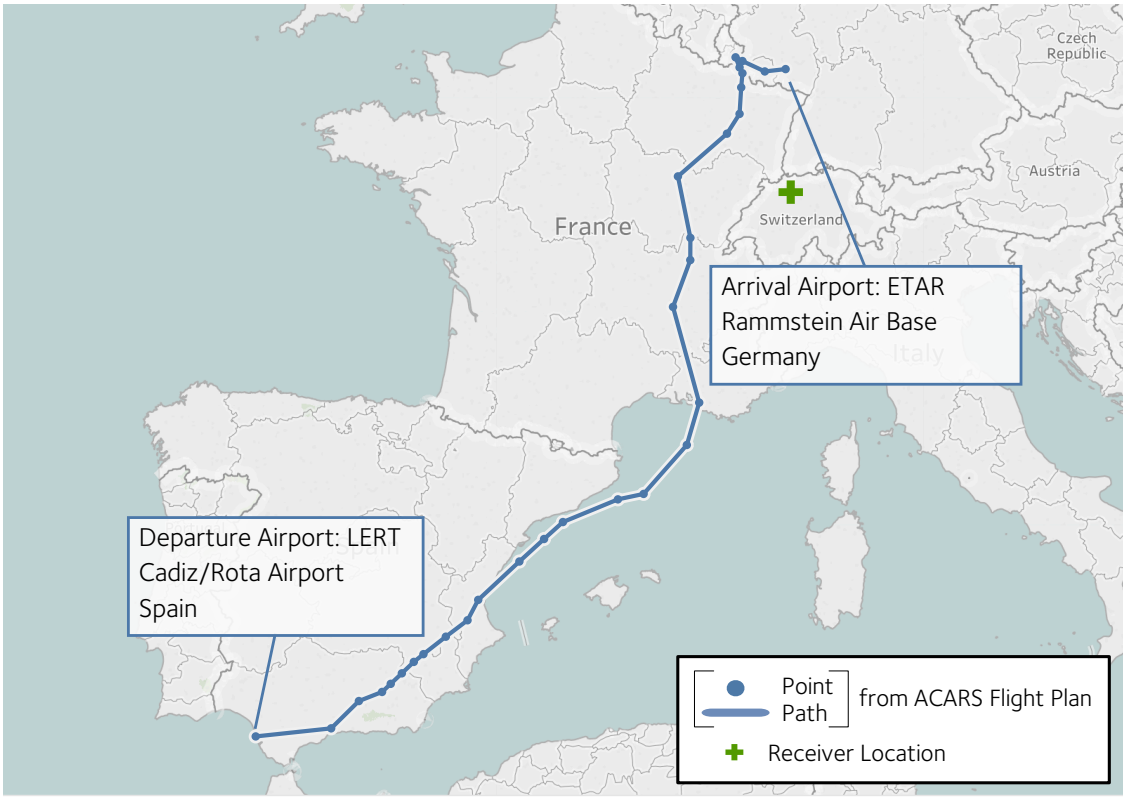

Figure 6.6: Plotted flight update from a US military aircraft in January 2017. This was collected on the SATCOM link.

We present the case of a specific military aircraft to demonstrate the impact of ACARS further. One American military aircraft only observed on the SATCOM link transmitted 513 messages between November 2016 and January 2017. Message content varied between free text, flight plans and weather reports.

When looking specifically at the content of flight plans from this aircraft we can see a lot of data; despite being blocked, it received 16 flight plans and 32 ATIS messages. On top of this, it transmitted messages providing far more detailed routes, apparently updates on the route taken so far. In one instance the aircraft transmitted a flight plan, then an update 40 minutes later. We have plotted this update in Figure 6.6. Despite the aircraft being blocked on flight tracking websites, an attacker can plot the movements of this aircraft based on either of these messages, collected far away from the departure airport. Importantly, by using both the flight plan and update, they can see where the aircraft is going without relying on more surveillance data.

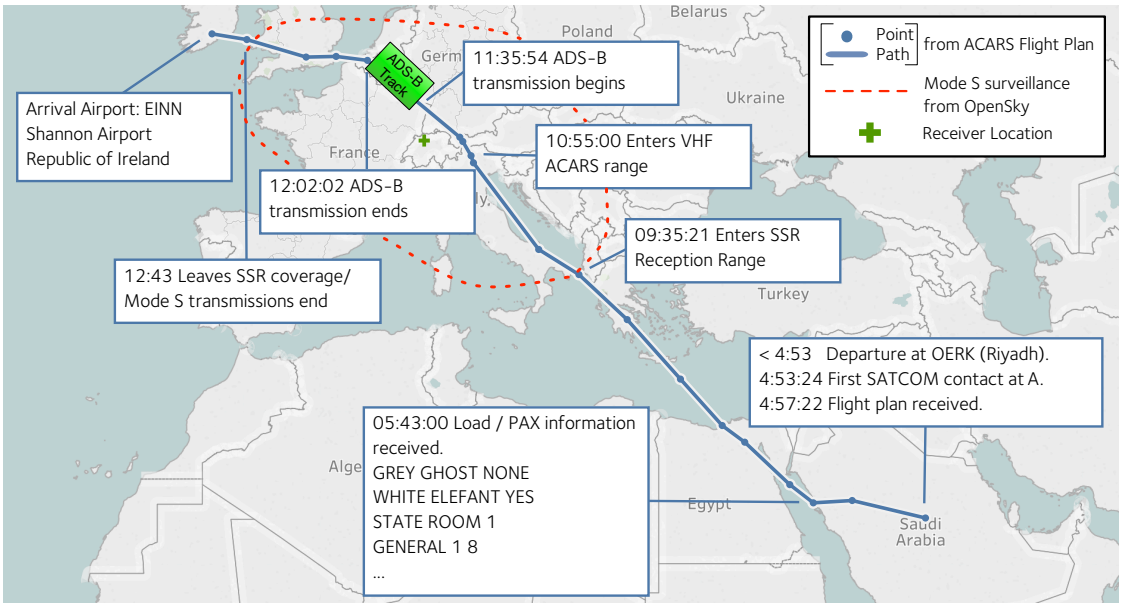

Figure 6.7: Flight track of a US diplomatic fleet aircraft in December 2016, as reconstructed from ACARS messages. Note that aircraft switches on ADS-B only for cruising, never for departure or landing.

6.4.3 State Aircraft Case Study

One of the strongest examples of the power of ACARS message interception is in data collected from an aircraft of the US diplomatic fleet. This aircraft appeared to have a source-level block, since it is considered US military. As such, no information about it can be seen in FlightRadar24 or FlightAware at the time of collection. Using SATCOM ACARS, we were able to not only receive messages and track the aircraft as it flew over our collection location, but also were able to gather load sheets and a flight plan. Constructed from five SATCOM ACARS messages collected within one hour, Figure 6.7 reveals the full route, despite the collection location being in middle of the track. These SATCOM messages were received before the aircraft entered VHF range, out of the line of sight.

Using Mode S and ADS-B data from OpenSky, we see the additional level of sensitive positional data leaked by ACARS. Firstly, the aircraft rarely turns ADS-B on apart from a period of 27 minutes at cruise altitude over France. This appears to be some attempt at hiding, since ADS-B is mandated for use by all civilian aircraft in US and European airspace by 2020—military aircraft do have the ability and

permission to turn this off though [40, 197, 198]. Alongside the flight plans revealing the origin, destination and waypoints along the route, other messages include a load sheet for items on-board. In terms of the latter, most is standard information, e.g. passenger count, fuel and take-off parameters. This has the potential to be sensitive since it indicates how many passengers are on board and as such, the type of mission underway. However, additional items are seemingly coded entries such as `WHITE ELEFANT` or `GREY GHOST`. No public information on these items exist, yet they appear in many load sheets sent by US diplomatic fleet aircraft.

Beyond this flight, the aircraft conducted other flights revealing route information, ADS-C updates, further inventories and free-text messages.

6.5 Industry Opinions

Clearly, ACARS is used to transmit location information regardless of how other systems are used; as shown in this chapter, this can often reveal information which an operator is trying to hide, thus making it sensitive. When compared to our survey of industry professionals as covered in Chapter 4, this highlights some contradictions.

In particular, we look at Figure 4.6b (pg. 63) which shows that of the 41 respondents, 18 felt ACARS is '*somewhat unsuitable*' for private data, and a further 14 felt it is '*very unsuitable*'. Even though this sample has 31 (78.0%) respondents describing ACARS as unsuitable for private data, many operators transmit sensitive location data over the link. Furthermore, we also asked participants if they had experience of sensitive data being transmitted over the link; 22 (53.6%) participants noted that they had.

Although our survey demographic (Figure 4.4b, pg. 61) shows a sample bias towards commercial aviation with 24 (58.5%) participants working in this domain, they highlight a similar contradiction. As covered in Chapter 5, these operators have different privacy requirements to business, military or state aircraft since a lot of their actions are public. Even so, their free-text responses indicated similar privacy issues to the those demonstrated in this chapter—13 participants noted that

passenger information was shared, 10 considered operations or maintenance data sent via ACARS to be sensitive and 5 stated that crew information was transmitted.

This highlights an interesting contrast—the vast majority of respondents felt the link was unsuitable for sensitive data, but we observed it used for this extensively, which aligns with their responses. Indeed, slightly over half of the respondents noted that they had seen sensitive data transmitted via ACARS.

One explanation for this is that ACARS is so pervasive that the convenience it brings overpowers considerations of privacy. Many of the services it provides save time or offer a service which might otherwise be difficult to provide. In the example of location reporting, ACARS may provide a cheaper or already-implemented system compared to accessing ADS-B data. For flight plans, sending them via ACARS allows easier modification and avoids having to copy them from printouts.

Another reason might be that participants do not have control over which systems are used, and the procedures relating to them. Even though a system may be clearly unsuitable for private data, their company procedures require them to use it. In turn, this could be due to a separation in decision-making between the technical and business aspects of operating an aircraft. This would explain why, even though effort goes into blocking aircraft, ACARS is still used for purposes which reveal a range of types of location information.

Ultimately there is a discrepancy in system understanding, especially with respect to the privacy limits of avionic systems. Even though some sections of industry know the shortcomings of the data links, this does not seem to map to how the shortcomings might adversely affect privacy.

6.6 Mitigations

Many of the leaks described in Section 6.3 are as a result of the type of message content; without it, the most a single receiver can determine is aircraft existence. Because of this, protecting message content is paramount. We cover two types of approach to this: technical and policy and provide a high level summary in Table 6.4.

Table 6.4: Summary of mitigations available for reducing sensitive information leakage on ACARS.

Mitigation	Difficulty	Cost	Level of Protection	Main Drawback
Standardised security	Low-Med.	Med.-High	High	Few deployment options, requires new hardware
Proprietary security	High	High	Low-Med.	Many pitfalls if design is not robust and well tested
Disabling/ firewall	Low-Med.	Low-Med.	Med.	Removes some capability from the system with no replacement
Data link policy	Low	Low	Low-Med.	Does not prevent accidental/deliberate misuse
Data legislation	Med.	Med.	Med.	Only applicable to charter

6.6.1 Technical Measures

There are some existing technical measures (in use or standardised) by which ACARS message content can be protected. Whilst they can provide improved security, they also have drawbacks.

ACARS Message Security

As discussed in Chapter 4, ACARS has provision for standardised security as part of ACARS Message Security, though deployment of this appears very limited. For example, in our data set we did not observe AMS being used. This also aligns with our survey, in which we found that only two (5%) professionals surveyed had any knowledge of AMS usage in practice, with 36 (90%) having no knowledge of deployment.

Using AMS would likely solve the issue of leaks to a passive attacker and would do so using standardised cryptographic approaches. It does, however, come with some challenges. As with any distributed security solution, implementing a public-key infrastructure is costly and requires thoughtful, security-conscious design. Especially in the case of aircraft, which must be able to communicate with

unexpected ground stations, keeping up-to-date credentials for all communication partners is a key challenge.

Furthermore, it requires specific software and hardware updates which take a lot of time and money to produce and deploy. It appears that the added cost of AMS has proven to be a major hindrance and the main reason for its almost non-existent deployment, even though the investment may be offset by the potential reputational damage and legal costs.

Non-standardised Security

In lieu of the standardised AMS, non-standardised and potentially proprietary cryptography is used by some operators. Naturally, the effectiveness of such measures depends entirely on the quality of the cryptography. Thus far, all attempts at this that we have been able to identify in the wild have provided no meaningful level of security but rely on insecure mono-alphabetic substitution ciphers instead.

Whilst the temptation of such cheap, proprietary cryptographic solutions is great, weak encryption is to be avoided at all costs. This has been observed in wide use, for example by business jets, and is covered further in Chapter 7

It could be deployed faster than standardised efforts and still designed within the restrictions of ACARS, particularly in the case of business aircraft or for small-scale organisations. To be effective it must be thoroughly tested by cryptography experts. However, as key management is still likely to be an issue, this solution is best used within a company or organisation rather than being a general solution. Even so, it might still require some change to hardware if implemented as an ACARS peripheral, which would be expensive.

Disabling ACARS Messages

As ACARS is not a technology mandated by any civil aviation authority, it is possible to forgo its use partly or completely. Some commercial airlines such as Ryanair do not use ACARS, reportedly due to cost reasons, and rely instead on mobile phone networks while they are close to the ground [165]. Thus, it is conceivable though operationally complicated and costly to abstain from using

ACARS for most aircraft. As an aside, it is interesting to note that economic pressure led to an airline to avoid using ACARS, but similar pressures do not lead to efforts to secure the link, as discussed in Chapter 4.

Instead of this extreme option, a monitoring system could be deployed at the network level to identify only potentially privacy-sensitive messages. Aircraft which wish to partake in this could request that certain message types be blocked from transmission if they are sent unencrypted. Since business, military and government aircraft which wish to hide from public data sources already try to do so, they could register a set of ACARS message restrictions for their aircraft. Should those messages then be sent without their knowledge or automatically, they would be filtered at the network level.

Of course, this would still restrict functionality for some aircraft and their operators. It would not work for ATC clearances, for example, which would cause the flight crew to have to fall back to voice communications. However, in the case of blocked aircraft transmitting position reports, many are using ADS-B thus have that as a source of tracking. Since ACARS is not designed for tracking it is arguably better—unless there is a specific reason otherwise—to use ADS-B which is designed for the purpose. Indeed, it would be one less privacy risk to manage.

6.6.2 Policy Measures

Another option lies in the creation of better policies to improve data security over ACARS. Of course, these are not mutually exclusive to technical measures, but may be more feasible to deploy in a timely fashion considering the often decade-long development cycles found in aviation [199]. As analysed in the previous sections, unless an effective security measure is in place none of the data links should be used to transfer sensitive data. A strong sensitive data policy could stem the issues described in this paper without heavily reducing functionality.

Data Protection Laws and Regulations

In many parts of the world, data protection legislation is a key measure in enabling citizens to protect their privacy. This is particularly relevant to the aviation scenario; those on board the aircraft are unlikely to control how their data is treated, and the primary method of transferring data is by default not secured. For charter companies carrying passengers but using ACARS, this could create legal difficulties.

With recent General Data Protection Regulation (GDPR) from the European Union, the potential consequences for failing to adequately protect data have increased to the greater of 4% of company turnover or €20 million [200]. Instances of sensitive data transmission in-the-clear ACARS appear to be a breach of this regulation since no effort at protection is made, despite the data being sensitive. This is especially the case if other information on top of location data is sent via the link.

Compliance is not necessarily cheap. It would require a review of the ways and means of ACARS usage with respect to data protection laws, which is likely to be costly. To then address these, another mitigation would be needed or ACARS usage for sensitive data would need to be avoided entirely—when in place, this solution would be comprehensive. This mainly applies to charter companies who are handling using ACARS for data relating to client location.

Internal Procedures and Education Policies

External regulatory pressure must be complemented by internal measures to be effective. As part of this, it is crucial to educate users both on the ground and in the air on the fact that all communication sent via ACARS is in the clear and effectively public, so should be treated as such. Our survey indicates that this is not currently the case. Where possible, codified processes should be adapted to reflect this mindset. While this is not offering a complete solution to the security issues of ACARS, it can at least mitigate them.

An example for how such policies can be effective was provided during our measurement campaign conducted for this paper. Although not a consistent data breach, some air transport providers use ACARS to validate credit cards used for

substantial purchases on board of aircraft [201]. We notified several airlines of their misuse of ACARS and provided proof of intercepted data. At the time of writing, at least one airline responded to us and changed their procedures to close this data protection issue by using tokens instead.

Whilst this would be one of the easier approaches to mitigating sensitive data leakage, it would not entirely solve the problem. End-users could still transmit sensitive data over the link without consequence, be it accidentally or deliberately. However, it would be relatively straightforward to develop information security policy into this area, as it can work around existing procedure.

6.6.3 Future Steps

In the longer term, steps should be taken to move away from the current ACARS technology completely, or at least to data links with network-level security. Since it was designed with a significantly weaker threat model in mind—i.e. one of no malicious activity—it is not equipped to deal with cybersecurity threats. According to the newest Global Air Navigation Plan by ICAO, which sets out the technology roadmap through to 2030, ACARS is intended for use in some form until at least 2030, if not longer [202].

Due to uptake on available security solutions being limited, a newly developed data link with security as the default may be the better option. However, given typical technological cycles in aviation, this would take decades to deploy fully [199]. In the meantime, our recommendation is that aircraft which insist on using ACARS fully but require security seek out AMS and absorb the cost as a necessary investment. Where possible, duplication of systems (e.g. position reporting with ADS-B) should be stopped such that sensitive data is not sent in the clear on two channels. Beyond this, the fastest way to achieve change in ACARS security and privacy will be to educate users such that they demand better systems.

6.7 Summary

In this chapter we demonstrated that ACARS usage poses a notable privacy risk to business, military and government across all three data links. Basing privacy on the notion of blocking aircraft and the legal, political and governmental pressures relating to it, we showed that for a modest attacker with a single set of sensors, much can be learned. This is particularly true of the SATCOM link which can collect far beyond line-of-sight to the aircraft due to the nature of the link. With further investment, an attacker could—for a relatively low cost—expand collection to a significant area and capture a great deal of privacy sensitive data.

After highlighting the message types which cause the most significant location privacy issues we illustrate the problem with case studies. This emphasises the importance of ACARS privacy for the military and government stakeholders, as both revealed a lot of data that they otherwise try to conceal. We contextualized this with industry opinions, showing that there is little awareness of existing security in ACARS, even if those surveyed believe that it is not suitable for private data. We concluded by providing mitigations and recommendations, given that this is not a trivially patchable problem, ultimately highlighting the fact that aviation must account for security at the point of design in the future.

Baroudeurs aren't always liked. 'Again!' the peloton says to itself when they attack. It is the baroudeurs who hand out leg ache.

— Paul Fournel

7

Deployed Security on the ACARS Data Link

Contents

Despite having a standardised security solution in AMS, deployment and usage of

security on ACARS is inconsistent at best. However, as demonstrated in Chapter 6 and in the literature, there is awareness of the issues caused by ACARS being unencrypted [139, 144]. As users of the link became aware of its practical insecurity, some required improvements to the confidentiality of their data.

This has been alluded to in existing work but not covered in detail. First mentioned by Roy, proprietary ciphers are identified as character substitution thus offering limited protection [139].[1] Later, monoalphabetic ciphers in use are referred to by Teso in [17] whilst covering ACARS as a system, though no detail is provided. These works suggest awareness of the limited ability of proprietary substitution ciphers to protect data. Since we have no evidence to suggest that the stronger AMS is used, it appears that weak ciphers have provided a quick fix but are now used as a long-term solution.

Our data provides an example of these ciphers, and indeed the only one seen used at scale; a monoalphabetic substitution cipher with nine static, shared keys. Furthermore, it is primarily used by aircraft who are attempting to otherwise protect their privacy, as in Chapters 5 and 6. In this chapter, we look at this cipher in detail and argue that this type of solution provides a false sense of security for ACARS users. Consequently, it does more harm for their reasonable expectations of privacy than no solution at all. More specifically, we:

- Show that the current most commonly used security solution for ACARS is highly insecure and can be broken on the fly,
- Quantify the impact on different aviation stakeholders and users,
- Provide lessons for the development of security solutions for existing legacy technologies, particularly in slow-moving, safety-focused critical infrastructure sectors.

[1] Roy uses this as a motivation for the design of Secure ACARS, which led to the standardised ACARS Message Security.

7.1 Threat Model

We extend our threat model used in the analysis of privacy leakage on ACARS in Chapter 6 to the *passive codebreaker* model. This focuses on an honest-but-curious threat who is passive with respect to the medium but actively decrypts messages: they collect ACARS messages and aim to break the cipher and decrypt messages that use it.

Since we use the same datasets as described in Chapter 4, our threat actor is acting from a single location but collecting on POA, VDL2 and SATCOM ACARS using commodity hardware and freely available software. A more capable attacker would be able to deploy multiple collection units across a larger geographic area in order to increase the message collection rate and the number of unique aircraft observed. As demonstrated below, this will increase the rate at which the analysed cipher can be broken.

7.2 Cipher Usage

Before covering our steps to break the cipher, we look at the extent of its usage, and the type of aircraft doing so. Relative to our full dataset, the number of aircraft using this cipher is low; the 3745 message forms 0.21% of collected messages and is transmitted from 337 aircraft, 3.40% of all observed aircraft.[2] The fact that these aircraft use encryption at all is anomalous in ACARS, thus warranting investigation.

7.2.1 Identifying Encrypted Messages

Structurally, ciphertext from this system is similar across all three links, with some slight differences in a prefix. This prefix indicates the key used, and in some cases the kind of messages. From our observations, we can identify encrypted messages as follows:

[2]Due to the small proportion of all messages, and significant variations in cipher usage between aircraft categories, we provide values to two decimal places in this chapter.

Table 7.1: Breakdown of encrypted messages (Msg.) across stakeholder groups and subnetwork (SN). Percentages are of all encrypted messages.

	Stakeholder									
	Business		Military		State		Commercial		Total	
Subnwk.	Msg.	%	Msg.	%	Msg.	%	Msg.	%	Msg.	%
POA	1097	29.29	1	0.03	2	0.05	7	0.19	1107	29.56
VDL2	830	22.16	14	0.37	10	0.27	0	0.00	854	22.80
SAT.	1689	45.10	60	1.60	35	0.93	0	0.00	1784	47.64
Total	3616	96.56	75	2.00	47	1.25	7	0.19	3745	100

- For VHF-based links **POA** and **VDL2** a string beginning with two numbers between `01` and `09` followed by other ASCII characters, with message labels `42` to `44`; e.g. `09a(s80(s0a...`
- For **SATCOM**, a string beginning with `TWX` or `MSG` followed a first set of two numbers (observed as `01` to `03`), then a second pair of numbers between `01` and `09`, followed by other ASCII characters; e.g. `MSG0104asd8*((0sa....`

When comparing messages beyond the above prefixes, it becomes clear that groups of messages see repeating characters in the same position across separate characters which implies the use of a monoalphabetic substitution cipher. This is especially clear when messages are grouped by the numeric prefix on VHF and the second pair of numbers in SATCOM messages; we refer to these as *key identifiers*.

7.2.2 Subnetwork Usage

In Table 7.1 we can see the number of encrypted messages in our dataset. Messages were sent via all three observed links, with SATCOM seeing the highest proportion at 47.64%. POA and VDL saw lower usage at 29.56% and 22.80% respectively, possibly due to the collection range being smaller than for SATCOM.

Looking to Table 7.2, we provide by-aircraft usage of encryption on each subnetwork. In contrast to Table 7.1, POA had the lowest number of aircraft using the link at 61 (18.10%) instances of aircraft, with VDL2 significantly higher at 209 (62.02%).

Table 7.2: Breakdown of aircraft (AC) using encrypted messages for each link. Note that aircraft can transmit on multiple links. Unique provides the number of individual aircraft across all subnetworks, hence an aircraft may be counted in multiple stakeholder groups. Percentages are of all aircraft sending encrypted messages.

	Stakeholder									
	Business		Military		State		Commercial		Total	
Subnwk.	AC	%	AC	%	AC	%	AC	%	AC	%
POA	57	16.91	1	0.30	2	0.59	1	0.30	61	18.10
VDL2	196	58.16	5	1.48	7	2.08	1	0.30	209	62.02
SATCOM	129	38.28	6	1.78	6	1.78	0	0.00	141	41.84
Unique	309	91.69	11	3.26	15	4.45	2	0.60	337	100

Compared to message count, POA has the highest average number of messages per aircraft at 18.1; VDL2 has an average of 4.1 and SATCOM has 12.7. One reason for this disparity might be that VDL2 still sees lower coverage and usage than POA, hence many aircraft continue to primarily use POA.

7.2.3 Stakeholder Usage

Looking to stakeholder splits, business aircraft sent the vast majority of messages across all three subnetworks, totalling 3616 (96.56%). This is matched to their usage of each subnetwork; over 90% of the aircraft observed on each link are business; as shown in Table 7.2, 309 (91.69%) unique business aircraft used it in total.

Very few state and military aircraft appeared to use this cipher at just 15 and 11 respectively, sending 1.25% and 2.00% of all encrypted messages. Although the numbers are small, these stakeholders appear to have a bias away from using POA, with the majority of messages for both of these stakeholders being on SATCOM. However, a similar number of these aircraft using VDL2 to SATCOM indicate a willingness or sufficient equipage to use VDL2 over POA.

An anomalous pair of commercial aircraft were observed using this encryption scheme so included in our results, despite lying outside of our non-commercial scope. The reasons for this aircraft using the cipher are unclear, but one reason

Table 7.3: Summary of aircraft manufacturers and models using ACARS encryption, with names omitted. Percentages are of all aircraft using this cipher. This table is an extract from the full table in Appendix D, focussing on manufacturers with more than three aircraft observed.

Manufacturer			Model		
Name	# Aircraft	%	Name	# Aircraft	%
D	11	3.26%	D-1	2	0.59
			D-2	1	0.30
			D-3	2	0.59
			D-4	1	0.30
			D-5	3	0.89
			D-6	2	0.59
E	20	5.93	E-1	2	0.59
			E-2	3	0.89
			E-3	1	0.30
			E-4	14	4.15
F	296	87.83	F-1	87	25.82
			F-2	154	45.70
			F-3	55	16.32

could be that they appear as a commercial aircraft in all but purpose, instead being a flag carrier-operated state aircraft.

7.2.4 Airframe Usage

Since such a high proportion of the aircraft were operated by business stakeholders, all except two could be categorised as corporate jets. To further analyse this, we identified the model of aircraft for each transmitting aircraft. An excerpt of these results are shown in Table 7.3, with the full table given in Appendix D. Clearly, manufacturer F has significantly more aircraft using this cipher than any other manufacturer, with model F-2 being the most prevalent of these.

We found that the vast majority of aircraft (333, 98.8%) using this cipher are equipped with the Primus suite avionics equipment from Honeywell, Inc. [203]. Since avionics handle ACARS messages, this implies that the cipher may be a feature of this avionics suite. With the different models of aircraft using the cipher, in turn using different variants of Honeywell Primus, we believe it is a feature of

this avionics suite. Therefore, we believe that any aircraft choosing this suite will be affected by the weak cipher, should they opt to use it.

7.3 Cryptanalysis of the ACARS Cipher

Having considered the type of aircraft using this cipher, we now cover the steps taken to recover keys from the cipher. To do so, we follow several classic cryptanalytic steps. We first describe how character substitutions can be recovered before moving to analyse the properties of the cipher.

7.3.1 Recovering Character Substitutions

As covered above, inspecting the available ciphertext, we note above that all messages ciphered under this label are prefixed in some form by two digits, from `01` to `09`. We refer to this as the *key identifier*. When messages are grouped by these digits, repeating characters in the same position across messages can be seen. From the similar set of characters used between messages of the same key identifier, this implies the use of a substitution cipher as well as an underlying common structure between messages.

Next, frequency analysis was used to compare the per-character distribution for each key identifier against all messages in our dataset. Since the encrypted messages are a small portion of our overall message set, we expected the character distribution of the underlying plaintext to be similar to the overall ACARS character distribution.

An example of this can be seen through frequency comparisons for key `01` in Figure 7.1, which is an excerpt from the full chart in Appendix C. We can see that in the cleartext distribution, `0` is clearly more prevalent than other characters, which matches to similar peaks for `l` or `b`. On top of this, `A` provides another plaintext peak, which is likely to be mapped to one of `l` or `b` in the ciphertext.

Due to the relatively small amount of messages compared to the total number of ACARS messages collected, this can only provide a starting point for decryption. To recover further characters, we can use known plaintext. Some aircraft send cleartext messages of the same length and ACARS label but without encryption.

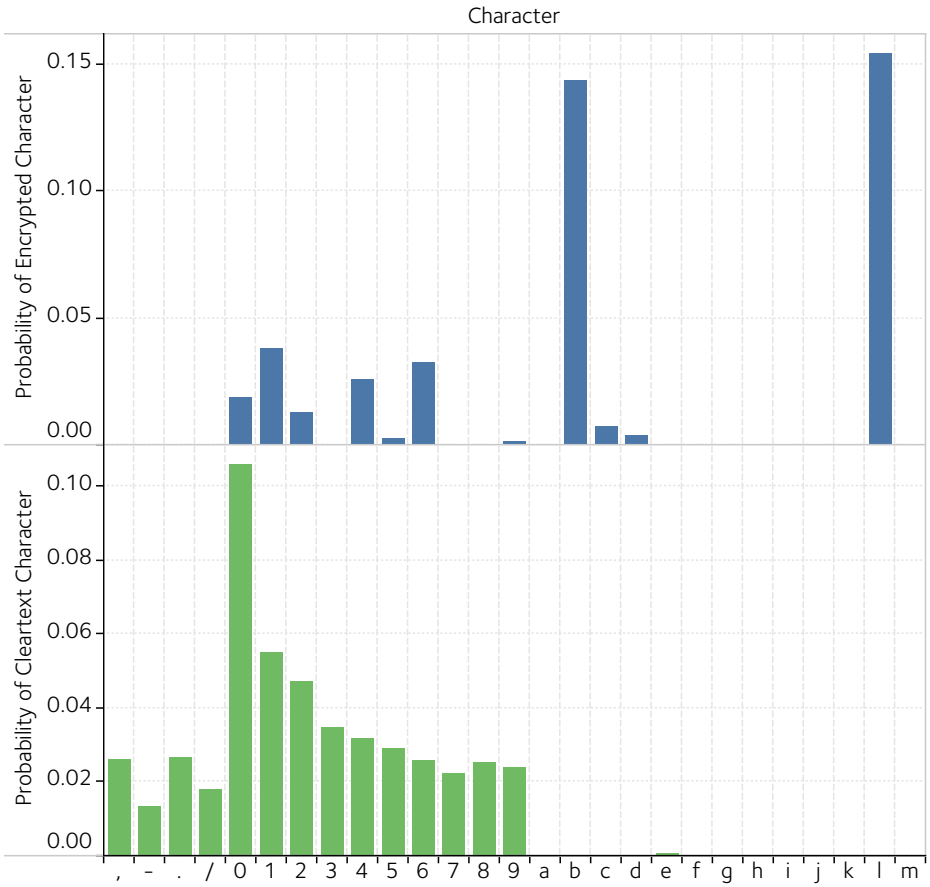

Figure 7.1: Excerpt of character frequency distribution of cleartext and ciphertext, from Appendix C. The selected areas in this chart are the characters of high probability in the character set.

Simple comparison of repeating characters in each, coupled with substitutions gained from frequency analysis, provides the next step. Furthermore, repeating characters in the ciphertext helps identify spacing or separating characters between portions of the message. A labelled cleartext status report message can be seen in Figure 7.2, in which we identified the fields based on meta-information and structure. Using this technique, we recovered other substituting characters using domain knowledge and heuristics.

7.3.2 Character Recovery Heuristics

Since we have a limited set of ciphertexts but now possess knowledge about the underlying structure of one message type and content of the fields, we can use heuristics to recover the remaining characters.

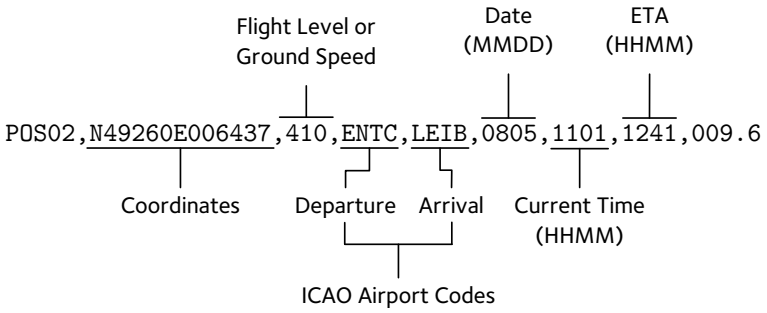

Figure 7.2: A labelled cleartext status report message sent under label '44', of the same format as VHF-based encrypted messages.

Recovering Coordinates.

As the second field in plaintext messages is a coordinate field, we use this to retrieve many substitution characters by exploiting the position of the receiver. Since the collected messages are received on VHF, we know the transmitter was within reception range limited by the radio horizon. Assuming the aircraft is reporting current position, we can then calculate that the coordinates in-message are limited to ±2-4 degrees longitude and latitude from the receiver, i.e. within reception range. This means that the options for the first two digits and direction letter (i.e. `N` for north) are also restricted. However, this approach becomes less reliable if the collection location lies on a point of 0° longitude or latitude.

Message Prefixes.

For some message types, the first field follows the structure of a three-letter code followed by two digits which we refer to as a message prefix; in the plaintext example of Fig. 7.2, this is `POS`. Looking at all plaintext messages received, some three-letter codes are significantly more common than others. Combined with already known letters, this reveals more substitution characters across key identifiers; we look at this further in Section 7.4.1.

Airport Codes.

As indicated in Figure 7.2, two of the fields are ICAO airport codes. Typically, countries will be assigned a prefix; in the case of the UK, this is `EG`. Based again

on the collection location, we can determine that local airport codes are more likely and use this as a heuristic for recovering substitutions—for example, if the collection range solely covers a part of the United States, one of the airport codes is likely to begin with the country prefix `K`.

We can also exploit partially decrypted messages containing ICAO airport codes by considering the set of possible codes from the possible decryption. Using known country prefixes and a list of airport codes, this allows comparison against various possible airport codes with a common encrypted character and so limiting the possible plaintext mappings.

SATCOM Meteorological Messages.

Not all character substitutions can be recovered from the reporting messages due to their limited set of used characters. However, aircraft receive periodic meteorological data over the SATCOM uplink to inform the pilots about the weather on their destination airport. Such messages take the form of Pilot Weather Reports (PIREP), Notice to Airmen (NOTAM), Meteorological Aerodrome Reports (METAR) and Terminal Aerodrome Forecasts (TAF). NOTAM, METAR and TAF messages originate from ground stations whereas PIREPs are reports of the conditions that the pilots encounter. Each has a consistent structure and contains regularly occurring phrases, which allows for character recovery when compared with plaintext obtained from other aircraft.

7.3.3 Key Analysis

Based on our observations, many of these messages use a limited set of ASCII characters, namely the printable characters between positions 32 and 126; this includes digits `0-9`, characters `A-Z` and symbols `,.*-:/?` and whitespace. This broadly aligns with the limited ASCII sets defined in ARINC 620 [136].

Using the 3745 messages, we recovered 661 of 855 (77.3%) substitutions across the nine keys. Further collection both in range or period, would allow the recovery of the remaining characters.

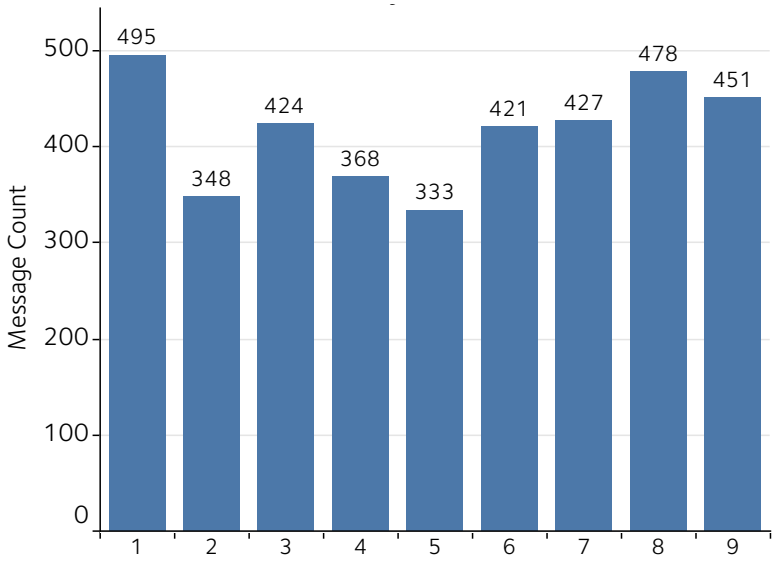

Figure 7.3: Distribution of encrypted messages across key identifiers. Data is aggregated across all ACARS subnetworks.

Theoretically, the ACARS alphabet size of 127 (i.e. the full ASCII range) offers a potential space of 127! keys. However, only 9 of these 3×10^{213} possible keys appear to used. Alone this makes the deployed cipher very weak. However, these keys are shared across all aircraft using this cipher with no evidence of re-keying. Key usage for a given aircraft does vary, though not in an apparently predictable way. In Figure 7.3 we display the distribution of key usage across encrypted messages, with $\lfloor \bar{x} \rfloor = 416$ and a range of 162. Clearly, this significantly reduces the difficulty of recovery across all keys by quickly providing enough ciphertext for which we can infer plaintext for each key.

Given the small set of global keys, users of many aircraft models might have the illusion of privacy when in fact this security solution is breakable. Furthermore, we have seen no attempts at key distribution or rekeying over the course of several months; the keys recovered from the data at the start of our sample work on our most recent data, too.

7.4 Message Content

We now consider the type of information contained within messages, especially with regard to potentially sensitive content. Message format across is consistent across

Table 7.4: Summary of encrypted message (Msg.) prefixes count for POA and VDL2 content. Percentages are of all POA and VDL2 messages.

	POA		VDL2		Total	
Prefix	# Msg.	%	# Msg.	%	# Msg.	%
ATI	2	0.10	1	0.05	3	0.15
ETA	62	3.16	48	2.45	110	5.61
POS	875	44.62	680	34.68	1555	79.30
TOC	80	4.08	61	3.11	141	7.19
TWX	68	3.47	60	3.06	128	6.53
OFF	10	0.51	0	0.00	10	0.51
ONO	5	0.25	0	0.00	5	0.25
OUT	1	0.05	0	0.00	1	0.05
INO	2	0.10	0	0.00	2	0.10
MWX	2	0.10	1	0.05	3	0.15
FLT	0	0.00	3	0.15	3	0.15
Total	1107	56.45	854	43.55	1961	100

two groups, the VHF links, POA and VDL2, and SATCOM.

7.4.1 VHF Messages

From the encrypted messages collected, 1961 (52.36% of all encrypted messages) are some form of status reports similar to Fig. 7.2. Although we have no official documentation on these messages, we have observed a range of prefixes within the message content. We summarise the frequencies of these prefixes in Table 7.4. Some match to flight phases such as Out-Off-On-In reports, whilst others indicate position, estimated time of arrival (ETA) or top of climb (TOC). The intention of `TWX`, `MWX`, `ATI` and `FLT` are currently unknown, as is the effect of a prefix on the message content.

From the message format we can deduce field content in these message types such coordinates, ICAO airport codes (implying departure and arrival airports), date, current time and ETA. Importantly, a single instance of these messages can be used to infer a great deal about a flight; current location, where it is landing and when it will do so.

Of the observed messages, 987 (50.33%) were transmitted by blocked aircraft, with POA comprising 839 (75.79% of messages on that link) and VDL2 the remaining

148 (17.33%, as with POA).

7.4.2 SATCOM Messages

Data sent via satellite takes a different form but is a significant portion of the dataset at 1784 messages, or 47.63% of all messages. Since SATCOM messages are not necessarily constrained to short text-field lengths by subnetwork limitations, like VHF-based systems, we see messages containing multiple features such as weather reporting, flight plans or notices to airmen. We also see evidence of an email platform operating over ACARS.

Flight Plan Metadata

Whilst we do not see detailed flight plans themselves protected with this cipher, we do see evidence of them on the SATCOM link. We observed 185 instances of flight plan requests or metadata, of which 159 (85.95%) were from blocked aircraft. Typically, these messages take the form of a response to a flight plan request, in doing so revealing the aircraft registration, departure and arrival airports as a minimum. Due to our collection of SATCOM uplink only, the downlink may contain more detailed data such as the plan itself.

Information Services

The vast majority of encrypted messages observed from SATCOM are information services, namely Meteorological Aerodrome Reports (METAR), Terminal Aerodrome Forecasts (TAF) or Notice to Airmen (NOTAM). As explained in Chapter 6, whilst these messages do not contain aircraft location, they provide information relevant to the flight such as nearby locations and aerodromes.

We observed 1171 METAR/TAF reports, of which 1097 messages (93.68% of METAR/TAF) were transmitted by blocked aircraft. For NOTAMs, we saw 548 instances, of which 481 (87.77% of NOTAMs) were transmitted by blocked aircraft. Often these message types will be combined into one message—within these message counts, 534 messages contained METAR/TAF data and NOTAMs, with 467 (87.45% of combined) being blocked.

Table 7.5: Summary of privacy status of aircraft(AC) stakeholders using the encryption. Commercial is included due to two instances of apparently commercial aircraft using the encryption. Percentages are of all aircraft observed using the cipher.

	Stakeholder									
	Business		Military		State		Commercial		Total	
Status	AC	%	AC	%	AC	%	AC	%	AC	%
Blocked	294	87.24	11	3.26	14	4.15	1	0.30	320	94.96
Public	15	4.45	0	0.00	1	0.30	1	0.30	17	5.04
Total	309	91.69	11	3.26	15	4.45	2	0.60	337	100

Contact Information

In some instances, we saw references to contact details. Primarily these were emails, however some instances had phone numbers attached. Emails were particularly easy to identify due to a consistent format. We found 174 instances of emails, some of which appeared to be for charter operations teams but others seemed to be personal contacts. Of these, 158 (90.80%) were from blocked aircraft.

7.5 Privacy Analysis

Clearly, a range of message types exist which can contain sensitive data; when considered against the number of blocked aircraft using this link, this cipher is undermining other efforts to protect privacy. Detailed above, messages revealing coordinates, arrival and departure airports and points of interest related to the flight are common.

In the case of the VHF-based links, over half of the messages sent were from blocked aircraft, thus revealing location. For SATCOM, we see fewer direct locations but instead data allowing us to infer routes or destinations; in some cases, we saw contact details. Given other efforts to protect privacy through blocking, and the use of this cipher, it appears that this information is considered sensitive by the aircraft operators.

At the stakeholder level, Table 7.5 summarises the extent of blocked aircraft usage across all links. Business aircraft form the majority of aircraft using the

cipher, as well as requesting blocks, with 294 (87.24%) in place. This category also accounts for most of the encrypted messages sent by blocked aircraft at 3288 (87.80%, shown in Appendix E), almost half of which are sent via SATCOM.

Whilst a smaller number of state and military aircraft use the cipher, all except one aircraft have blocks in place. This is further reinforced by message counts, with just three messages being sent by an unblocked state aircraft. The remaining 47 state and 75 military messages were all from blocked sources.

With 94.96% of aircraft and 91.00% of messages using this cipher having flight tracker blocks, it is clear that this group of aircraft have a privacy requirement. Indeed, the fact that they are using an optional cipher on their ACARS messages supports this. However, this cipher provides no meaningful protection—instead, it arguably provides a false sense of security. This creates a situation where private information can be retrieved by attackers with relative ease, despite efforts to prevent it.

7.6 Discussion

As protocols are in use for many decades and are surpassed by technical progress and new user requirements, the temptation for quick fixes is great. In aviation, data links evolved to serve applications for which they were not initially intended (e.g., ACARS for ATC as in [122]) and requirements changed to include confidentiality to enable privacy for its users. Unfortunately, the presently deployed attempt to protect ACARS does not meet these requirements as we have shown.

It is thus critical to take away several lessons from this study. Research suggests that similar cases can be found not only in the wider aviation scenario (such as in surveillance technology [21]) but in many safety-focused critical infrastructures using communication systems such as SCADA, which typically do not support encryption at all [204, 205].

7.6.1 Time in Deployment

As the discussed solution has been greatly obscured, we could not obtain the exact time when it was first deployed but the range of aircraft using it indicates that this has been the case since the mid-2000s. This in turn means this solution has been in use for at least 10 years without proper independent analysis, and once installed is hard to remove at scale. Integrating the security community early on could have avoided the deployment of inferior solutions.

7.6.2 Proprietary Cryptosystems

The described attack serves to illustrate the dangers of attempting to produce cryptosystems without due peer-review or use of well-known secure primitives—indeed in this case, without any reasonable primitives at all. This is especially the case in this situation where the nature of ACARS limits the cryptographic solution due to character set, message size and bit rate. Indeed, proposals such as Secure ACARS use AES, which is standardized and widely tested [206]. To draw parallels outside of the aviation scenario, WEP encryption suffered a similar fate in that an attempt to devise a security solution was critically impaired simply by misusing cryptographic primitives [207]. However, in the case of WEP, the primitives themselves were sound—in the system discussed in this chapter, even the primitives were not sound.

7.6.3 False Sense of Security

Developing—and deploying—solutions without such expertise can indeed be harmful. A solution that provides no effective protection has two distinct negative effects. First, it undermines the development and use of better solutions. In the case of ACARS, a demonstrably secure solution based on ACARS Message Security would be standardized and use reasonable primitives, but users who want data link confidentiality have opted exclusively for the discussed cipher be it for cost or marketing reasons. Secondly, it provides its users with a false sense of security.

Believing in the hardness of the encryption may lead operators to rely on the confidentiality they seek and potentially even modify their behaviour.

This case presents a robust example of how, coupled with other information typically transmitted over ACARS as shown in Chapter 6, encryption systems such as the one considered in this chapter can amplify, rather than mitigate, privacy issues. By taking steps to protect a data link, sensitive users may choose to transmit data over it which they otherwise would not. This causes exposure even for those trying to secure their data on the ACARS link.

7.6.4 Future Cipher Usage

Based on these lessons, we recommend that this security solution should not be used further. With little cryptographic knowledge or resources, message content can be recovered in real time. At the very least, manufacturers should discontinue the inclusion of it in future systems. Ideally, it would be patched out or replaced with a more secure option on existing aircraft and avionics. For users relying on this cipher and seeking better protection, we propose that they demand an established solution such as Secure ACARS which is a more complete security suite.

7.7 Legal and Ethical Considerations

Due to the sensitive nature of this work, we have ensured that it has been conducted in a manner which upholds good ethical and legal practice. At the start of the work we obtained ethical approval process to sensitive messages, and we followed a responsible disclosure process with Honeywell, Inc. We adhered to all relevant local laws and regulations.

We have further chosen not to name the aircraft manufacturers and models affected, as this could unduly impact the users of the affected aircraft before there is a chance to address the problem. Furthermore, we have outlined the steps taken to break the cipher but decided to omit further details and unredacted messages to avoid making such an attack straightforward to replicate. Overall, we believe it is crucial that all aviation users are aware of weak security solutions protecting

their communications so that they do not fall prey to a false sense of security but instead can take the necessary steps to protect themselves.

7.8 Summary

In this chapter we have demonstrated the shortcomings of a proprietary encryption technique used to protect sensitive information relating to privacy-aware aircraft operators. More specifically, we have shown that it cannot meet any security objective. We recommend its users are made fully aware that it does not provide actual protection; users should either seek a more robust security solution or avoid using ACARS for sensitive material.

Also, we highlighted the privacy issues arising due to this, since the cipher is primarily used to transmit locations and destinations by aviation users attempting to hide their existence and intentions. We show the cipher's weakness consistently undermines the users' efforts to hide their positional reporting, or protect message content which might be valuable to an attacker.

Consequently, we claim that when such solutions are deployed in practice it does more harm than good for users who require confidentiality from their data link. It is crucial that the aviation industry takes the lessons learned from this case study and addresses these problems before they are widely exploited in real-world attacks.

Training can be monotonous, and it is hard work, but you never lose sight of why you are doing it. Every single effort of every single session counts in the months and years leading up to a big event.

— Chris Hoy

8

Analysing Pilot Response to Attacks on Avionics

Contents

Having considered security and privacy on a widely used communications link, we now turn our attention to navigation and surveillance systems. As with ACARS, many avionics using wireless communication do not have security mechanisms and so are theoretically vulnerable to a range of attacks. However, understanding the effect of such attacks is difficult.

One of the reasons for this is the inclusion of humans in the loop in the form of pilots. When attacks affect what the crew observe through their aircraft systems, this may influence how they react and fly the aircraft. Unexpected aircraft or system behaviour due to attacks can lead to grey areas in procedure and as a result, unpredictable actions from flight crew. This is interesting when applied to safety systems that pilots are taught to trust, but we believe can be compromised. For these systems, it is useful to learn whether existing pilot training leads to consistent responses to theoretical attacks, how this might affect the safety of the aircraft, and whether pilot responses mitigate attack effects in any way.

Furthermore, developing and testing attacks on real hardware is expensive. Performing preliminary research to identify where countermeasures are most needed would allow researchers to target their efforts. Whilst developing practical proofs of concept for attacks would be a first step, doing so for avionics is challenging due to their specialised nature. Many systems will only operate as part of a wider aircraft electronics system, which can be replicated with a harness but is expensive and requires domain expertise. For example, a general aviation Mode S transponder might cost in the region of £2000-4000 but also would need antenna, wiring looms and supporting instruments configured as in a real aircraft to function properly.

In order to direct mitigation research to attacks of the highest impact, this chapter proposes the use of flight simulation to assess attack feasibility and impact. To do this, we take the following steps:

- We identify theoretical attacks based on current aviation security research for three systems: collision avoidance, instrument landing and ground proximity.

- We construct such attacks in a flight simulator and run experiments for 30 Airbus A320 pilots.
- Using in-simulator and interview debrief results from the experiments, we demonstrate that these attacks have disruptive impact.
- We suggest countermeasures or mitigations, and preliminarily assess the simulation approach as a training tool.

Assessing Humans in the Loop

Pilots, as humans in the loop, are trained extensively. Gaining a licence takes 3-4 months for private pilots, with a year or more full-time training needed for commercial aircraft [208]. As part of flight crew education, pilots are taught the specifics of the aircraft they will fly. In the commercial aviation setting, this is known as *type rating*, with a pilot only able to fly an aircraft they are type rated for. One of the aims of type rating is to provide a detailed education on the systems and features of a particular airframe [209]. Furthermore, to ensure that training is current, pilots undergo reassessment over time, often in a flight simulator [210]. Having such a structured approach to training helps the flight crew to be current and take consistent steps in response to many situations, from routine to emergency.

As part of their training, pilots are taught about the many systems used in an aircraft to reduce workload, increase efficiency and improve safety. We have covered some of these in Chapter 2, noting that each is designed to address specific problems and threats. Pilots are taught to trust the systems and their output to keep the aircraft safe.

At the same time, SDRs are enabling new security research to look at these systems. As well enabling research, this is creating new threats. Such a shift raises the question of the how much impact an SDR-equipped attacker could have if they compromised key avionics, especially those used for navigation or surveillance. Attacks on these may lead to systems acting outside of 'normal' behaviours, possibly inducing situations outside of existing crew training. We wish to understand what

flight crew responses to these attacks might be, from mitigating it successfully, to failing to spot an attack entirely, or switching an under-attack system off.

So far, opinion is split on whether attacks would effectively be contained by normal crew procedure. In a review of pilot perspective on cyber attacks, the Air Line Pilot's Association highlight that views on the new threat are split: some believe it is addressed already in aviation's extensive development and safety certification processes whilst others believe attackers could create 'unsafe flight conditions' [193]. Existing literature as covered in Chapter 2 suggests that current technical approaches to security are not effective.

Numerous cases demonstrate that consequences can be serious when avionic systems malfunction or are not used as intended. In 2006, two aircraft collided in Brazil, partly due to a transponder failing and not providing collision avoidance messages [211]. Similarly, Turkish Airlines Flight 6491 flew into terrain on approach due to failure to acquire instrument landing system signals [212]. In both instances, procedures and systems now exist to identify and avoid the underlying causes.

System 'noise' can also be a problem; if malfunctions or glitches happen on a regular basis, they can seed mistrust and create potential for severe harm. For example, instances of controlled flight into terrain have occurred due to a mismatch between where crew believed the aircraft was and what the terrain warning systems were telling them, leading to crew disregarding alarms. The 2012 crash of a Sukhoi Superjet saw crew treat terrain warnings as spurious before colliding with a mountain [213].

Procedural grey areas also have the potential for severe harm. Perhaps the most notable example is the Überlingen crash in 2002 [214]. Two aircraft collided mid-air due to mismatched commands from ATC and TCAS. At the time, procedure did not set out command priority between ATC and TCAS, resulting in one crew obeying ATC commands whilst the other crew followed the TCAS command. This led to the aircraft moving towards—rather than away—from each other. Such instances of ambiguity are almost inevitable in complex systems and can go undetected until an incident occurs.

Whilst these cases were not a result of cyber attacks, they demonstrate the worst-case scenario which can arise when key systems are disregarded, are producing suspicious information or have ambiguous procedure. Indeed, even if attacks on the systems are not destructive, they may well cause impact to both the targeted aircraft, or to nearby traffic.

Simulation for Flight Training

One way to assess the feasibility and potential effect of attacks on an aircraft is to use flight simulation, a well-established means of training and assessing pilots. To the best of our knowledge, this study is one of the first to use flight simulation to assess the impact of cyber attacks on aircraft. Because of this, it is worth considering the related work with regard to the effectiveness of flight simulator training for unexpected events.

Time spent in the simulator is a vital part of professional pilot training. A body of research analyses the configuration of simulator scenarios such that they transfer most easily to flying the real aircraft. Early research indicated that it provides notable benefit over aircraft-only training [215]. However, it is not a given that high-fidelity simulation transfer skills well, and the literature suggests that well-designed scenarios are vital in equipping pilots effectively [216, 217]. As such, using cyber attack scenarios for training may be beneficial.

One of the key factors in cyber attacks is that there may be no forewarning, leading to surprise and loss of capacity. In [218], a survey of aviation incident reports highlights that 'normal' events can be surprising to pilots when they occur out of context, e.g. alerts when the conditions do not warrant it. The authors in [219, 220] simulate this for stall recovery manoeuvres, a regularly tested skill for pilots. Both papers find that pilots struggled to follow even well-known procedures when the stall occurred in unexpected conditions.

Addressing this, the authors of [221] argue that unpredictability and variability in simulator training improves performance when encountering surprise scenarios.

Whilst their work uses failure scenarios instead of malicious interference, this is promising for use with cyber attacks.

Simulating Cyber Attacks

Some work addressing simulation for cyber security has begun to emerge. In [222] the authors conduct a human factors focussed study to assess how pilots respond to an attack on ground-based navigation systems. They find that pilots under attack lose some monitoring capacity, and that warnings can help mitigate this. The authors of [223] (and the extended [224]) conduct a more avionics-focussed set of attacks, looking at six variants of navigation and flight management system threats. Multiple attacks inserted over the course of one flight with the intention being to assess if pilots notice the attacks. They found that most attacks were identified during flight, however some happened without detection.

The work in this chapter differs in that it focusses specifically on safety-related systems using some form of wireless transmission, through which the attacker is instead aiming to disrupt normal flight. In doing so, we explore a different set of systems and cover the principles of these attacks in technical detail.

8.1 Threat Model

For this work our attacker is at its strongest as the *active and determined* model according to Chapter 3; we now have an active attacker who is more technically capable and equipped.

We focus on those aiming to induce reputational and financial harm by reducing safety levels and causing inconvenience. This could be through forcing diversions or missed approaches to land, known as a go-around, or pushing crew to switch off distracting systems. Although these procedures are for safety thus are not a negative or problematic response, they inevitably result in inconvenience.

An example of the kind of disruption possible was seen during the Gatwick Airport drone incident of December 2018, wherein unidentified drones flew in restricted airspace around and over the airport [225]. On top of passenger travel

plans being heavily affected, airlines fell off schedule and had displaced aircraft worldwide. Extensive defence efforts were deployed to try to stop the drones. Easyjet, one of the major operators at Gatwick Airport, saw costs reach £15 million, with policing costs reaching £400,000 [226, 227].

For our attacks, the attacker's budget is presumed to be in the region of tens of thousands of pounds to buy commercially off-the-shelf antennae, amplifiers and SDRs which are able to transmit at sufficient power to communicate with airborne aircraft. They have the capability to develop software which can be used to generate both digital and analogue signals or do this using existing open source tools. Attackers can deploy their systems remotely or create a mobile platform from which to do this. A more capable attacker such as a nation state could scale this up considerably.

We isolate the impact of attacks to one aircraft to manage the simulation complexity. Some attacks described in this paper have variants which could affect multiple aircraft or cause more serious safety incidents but require a stronger attacker. Handling these situations well and being able to accurately measure responses would require a full crew in a familiar simulation environment. Since this is an initial study, we leave this for future work.

We note that we constrain our threat model to not consider intentionally destructive attacks such as causing mid-air collisions or controlled flight into terrain. Whilst non-destructive attacks appear to be less important than attempts to crash an aircraft, we believe that a non-destructive threat model is more realistic at this stage of research and according to real-world threat actor capability. Whilst some variants of the attacks presented may be destructive, these are likely to be more overt and lead to the pilot defaulting to flying the aircraft 'by hand' rather than relying on the aircraft systems. At this stage, we instead want to understand more subtle cases where attacks may not be immediately obvious.

Experimental limitations should also be considered; destructive attacks may induce high-workload situations dependent on a full flight crew with familiar controls, and should also factor in air traffic control interaction. Since we cannot replicate

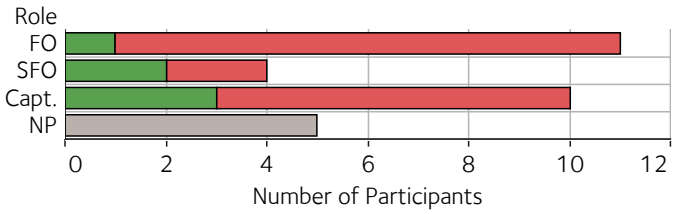

Figure 8.1: Participant role demographics in flight crew: Captain (Capt.), First Officer (FO) and Senior FO (SFO). NP is where participants chose not to provide data. Green bars indicate a training role, red for those without and grey is 'not provided'.

this yet, we felt we could not assess such attacks accurately. Limitations are discussed further in Section 8.5.6.

8.2 Experimental Method

Since our attacks were designed to fall in procedural grey areas and have unpredictable responses, we needed to adopt a methodology which allowed participants to react in real time. We chose to use a flight simulator to recreate the environment of the cockpit. This section describes our method in more detail. The work was approved by our local ethics committee with reference number R54139/001.

8.2.1 Participants

We recruited 30 pilots who had current A320 type-rating or had held it in the past few years but since had moved to larger Airbus aircraft. Our sample was recruited through pilot forums, and open to pilots of any level of experience, First Officer or Captain. This is appropriate since pilots are trained and kept current with a homogeneous skill set across a given type of aircraft. Thus, all pilots are similarly skills-equipped to handle the scenarios we presented to them.

We collected demographics from participants, with an option not to provide information if desired. In Figure 8.1 we show participants by both role and whether they hold a training capacity. We split into the three key crew roles in order of decreasing seniority: Captain, Senior First Officer and First Officer. Furthermore, the colours indicate whether the participant trains other pilots as part of their job.

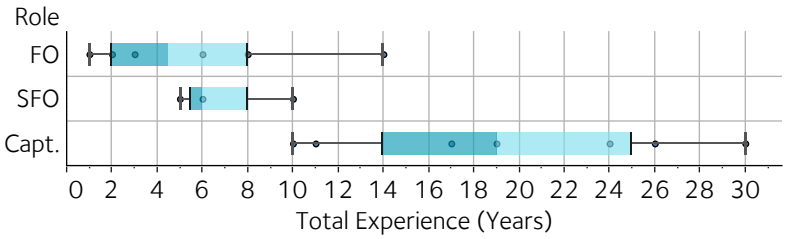

Figure 8.2: Plot of participant commercial flying experience by role: Captain (Capt.), First Officer (FO) and Senior FO (SFO).

In Figure 8.2 we provide a chart of participant commercial flying experience, grouped by role. Note that captains have a wide range of years of flying experience due to the requirements for taking a captain role varying between company and location. The median total years of commercial flying experience for a Captain was 19, for an SFO was 6 and for an FO was 4.5.

8.2.2 Protocol

For the purposes of control, we used the same weather conditions, traffic, and route for four runs. Pilots were asked to fly between runway 09R at London Heathrow to 33 at Birmingham International, cruising at 12,000 ft, for a total flight time of around 30 minutes. The route followed the BUZAD4J standard departure, and the GROVE1A standard arrival on direct arrival, rather than via the hold. The route specifics are provided in Appendix G, with a map of the route given in Figure 8.3.

We used a single-pilot setup with the experimenter taking the role of a limited co-pilot, and in doing so partly an air traffic controller. For consistency, the experimenter was capable of providing support in flying the aircraft (e.g. enabling modes or pressing buttons on request) in a way that a pilot not flying would but did not give decision input.

Each pilot was given the first run as a familiarisation flight, in which they could get used to the controls of the simulator. The following three runs included one attack, each followed by a short debrief interview about the attack; we include the experimental procedure in Appendix F. The interview assessed the pilot response to each attack, focussing on perception of impact, trust, workload and safety. We recorded data from the simulator to correlate with interview responses.

Figure 8.3: Plot of route used for simulation scenarios. The route runs between London Heathrow runway 09R and Birmingham International runway 33.

Participants were told that they were taking part in a study looking at cyber attacks on avionics but did not know about the timing or type of attack. Whilst this does prime them slightly, a limited amount of information helped us to recruit participants from a range of places in the aviation industry. This was important as we had no guaranteed access to participants, many of whom travelled in their free time to take part. The details of the attacks were explained by the experimenter in a debrief interview.

8.2.3 Equipment

Our hardware setup consisted of two high-end gaming PCs, running X-Plane 11 and an aftermarket Airbus A330 model, as no reliable A320 models were available [228]. We checked the model's fidelity with type-rated Airbus A320 pilots to ensure similarity to an A320. We provided single pilot non-type-specific hardware controls, which allowed flexibility in the design stage but does limit the pilot interaction slightly. However, as the majority of flying on such an airliner involves manipulating automatic flight, rather than directly flying with manual controls, these controls were still suitable. A picture of the setup is show in Figure 8.4.

Figure 8.4: Image of the flight simulator setup as used for experiments.

As this setup is naturally somewhat different to an A320 specific simulator, we assessed the level of limitation felt by pilots from '*heavily*', '*somewhat*' to '*not limited*'. Of these, 20 (69%) felt '*somewhat*' limited, with 7 (23%) not feeling limited. Most participants put the limit down to the effectively single-pilot setup, and non-type specific controls.

Related work suggests that lower simulator fidelity does not necessarily pose problems and may encourage participants to engage in the tasks deeply [217]. Whilst our limits were resource-based, we found that on the whole participants did indeed engage with the tasks well; all found the scenarios to be useful, with 28 (93.3%) feeling that this type of scenario could be useful for further training.

8.3 Systems and Attacks

We now describe the systems and theory of attacks used in the experiment, providing details of the capabilities an attacker would need to carry out said attack and the expected participant response.

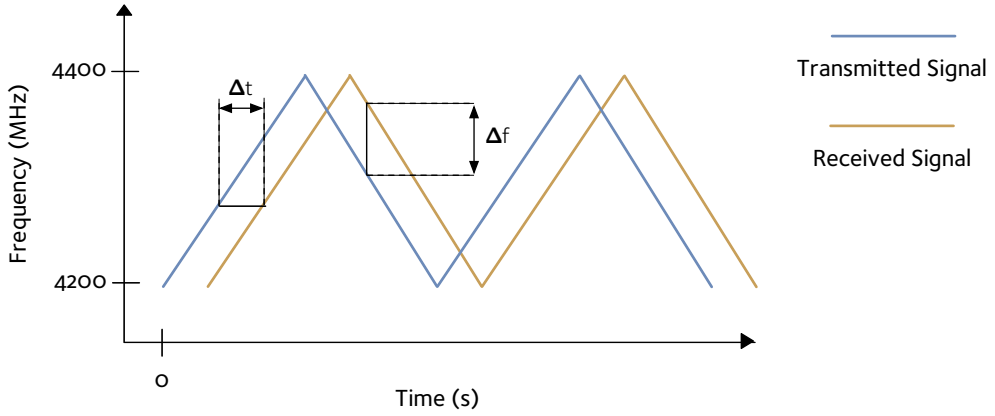

Figure 8.5: Representation of signals sent and received by an FMCW radar used for GPWS radio altimetry at a constant altitude, based on diagrams in [230].

8.3.1 Ground Proximity Warning System

A fundamental part of an aircraft's 'safety net', the Ground Proximity Warning System (GPWS) provides early warning of the aircraft becoming too close to terrain [229]. Failure to respond to this warning can result in a controlled flight into terrain incident.

System Description

As covered in Chapter 2, two versions of this system exist under the more general Terrain Avoidance and Warning System (TAWS) label: GPWS and the newer Enhanced GPWS (EGPWS). The part of the system used in this study is the same in both. GPWS takes input from a range of sensors including the radio altimeter, barometric altimeter and a vertical speed sensor [107]. EGPWS additionally uses GPS and a terrain database to assist in determining the position relative to terrain. The systems have a range of modes, each of which cater to a specific type of alert and phase of flight. We are interested in Mode 2, which relates to excessive closure on terrain [229].

Mode 2 GPWS uses the radio altimeter to determine the rate of closure on nearby terrain. A radio altimeter typically uses a Frequency-Modulated Continuous Wave (FMCW) radar. FMCW radars perform frequency sweeps in the 4200–4400 MHz range, using the frequency shift in the received signal to calculate absolute height

above terrain. This height is also referred to as *above ground level* (AGL). A representation is shown in Figure 8.5.

Mode 2 has two sub-modes, A and B: A is active during climb, cruise and the first part of the approach, whilst B is used in landing configuration. We provide a representation of mode B on approach in Figure 8.6a. Mode 2 GPWS is primarily used on approach; for example, on the Airbus A320 the instrument becomes active below 2500 ft, at which point the height readout is displayed on the primary flight display. In this mode, excessive terrain closure will be met with audio alerts, the most serious of which is *'Terrain Terrain, Pull Up.'* GPWS is considered one of the highest priority alerts to which crews must respond however the exact response is determined by airline procedures.

Attacker Aim

For this system, our attacker aims to create a spurious and unexpected terrain warning when an aircraft is on final approach. They intend to use this alert to negatively impact situational awareness and causing an unwanted go-around. As a result, aircraft will then have to perform a second approach or divert to a different airport. During this time, the aircraft will be using extra fuel, incurring delay, as well as adding workload for the pilots. In some cases, it may delay other aircraft waiting to land.

Attack Description

Mode 2 GPWS relies on a FMCW radio altimeter to measure height above the ground. By transmitting false radar pulses on final approach, the attacker causes the GPWS to believe that the terrain closure rate is significantly higher than it is. This would require the attacker to calculate the timing and frequency shift of the wave or take flooding approach to saturate the frequency range with pulses as indicated in Figure 8.6b. This will cause it to trigger a *'Terrain Terrain, Pull Up'* whilst the aircraft is close to the ground yet within the 'safe' range.

From a technical point of view, two approaches exist which could achieve this attack. In a targeted scenario, the attacker would aim to replicate the rapid

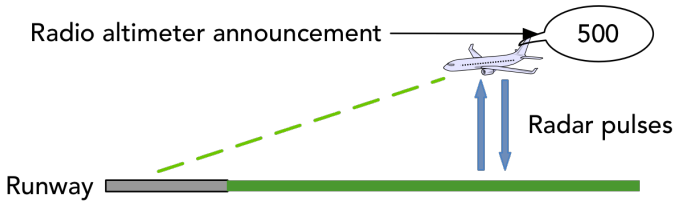

(a) Normal system operation with radio altimeter determining height above ground.

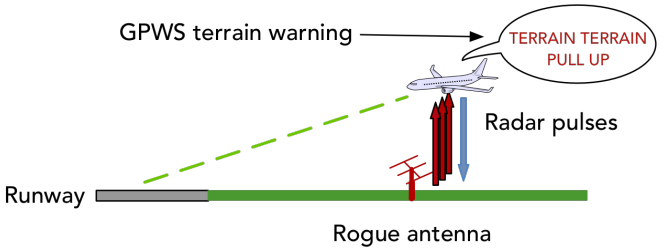

(b) Attacked system with false radar pulses forcing radio altimeter into reporting high terrain closure rate.

Figure 8.6: Representation of normal and under-attack GPWS performance. Speech bubbles represent system annunciations in the cockpit.

closing of ground by transmitting at a ramp of frequencies between 4200 MHz and 4400 MHz. The gradient of this ramp should be created to effectively incrementally reduce the round trip time for the signal. This would create the same effect of the ground approaching rapidly—we illustrate normal operation in Figure 8.7a and the system under attack in Figure 8.7b.

This approach requires some prediction of the phase of the signal from the aircraft, as well as knowledge of the frequency of the sweep itself. Since Mode 2 alerts are based on the rate of change of descent, this could be approximated and based on calculations of what is required to cause an alarm. As an example, a closure rate of 3000 ft/min at 500 ft AGL will trigger an alarm based on Figure A2b in [107]. We first note that calculating the round trip time, t_{rtt} (seconds) of a pulse is simply through

$$t_{rtt} = 2\frac{h}{c} \tag{8.1}$$

with h being the height above ground (metres) and c being the speed of light (metres/second). We can extend this to calculate the change in round trip time as the height of the aircraft changes as follows:

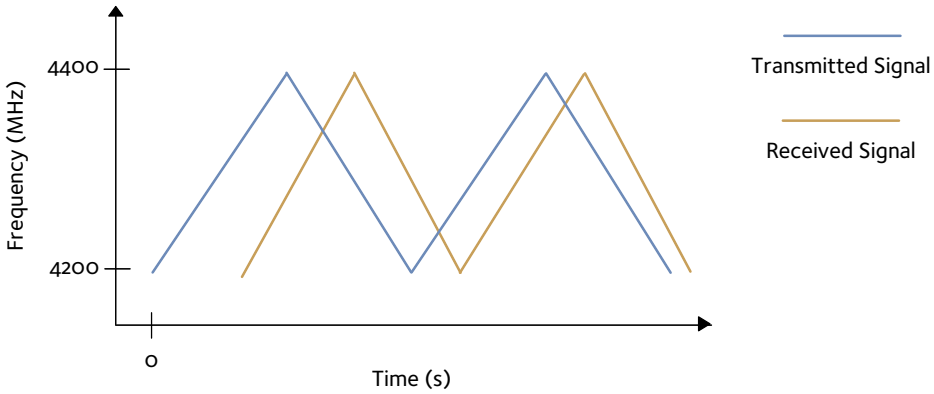

(a) Descending frequency-modulated continuous wave (FMCW) radar operation, for both the transmitted signal and received, reflected signal.

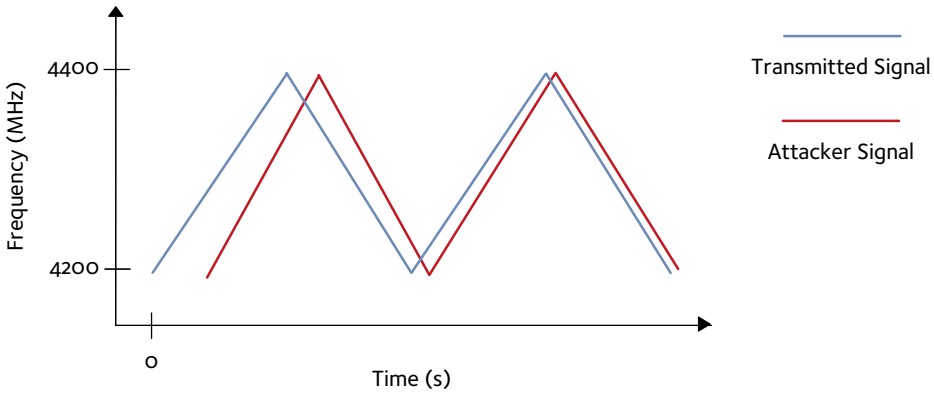

(b) Under-attack frequency-modulated continuous wave (FMCW) radar operation, for both the aircraft transmitted signal and attacker transmitted signal. Note that the attacker transmits signals such the rate of closure is higher than true rate of descent, such that the radar perceives that the aircraft was closing on the ground faster than it is.

Figure 8.7: Representation of FMCW radar operation for radio altimetry when descending and descending under attack, based on diagrams in [230].

$$\Delta t_{rtt} = 2\frac{h_1}{c} - 2\frac{h_2}{c} = \frac{2(h_1 - h_2)}{c} \tag{8.2}$$

We can then calculate the target change in round trip time if the aircraft were descending at a closure rate of 3000 ft/min, calculated over the course of one second. An attacker could aim to emulate such a closure rate by transmitting pulses which mimic the change in closure rate. If we presume that the aircraft travels a negligible horizontal distance during a single pulse, we can calculate the rate of change of round trip time as the aircraft changes height. We first use

conversions in Equations 8.3 and 8.4:

$$\text{AGL } h = 500 \text{ ft} \approx 152.4 \text{ m} \tag{8.3}$$

$$\text{Closure rate } s_{closure} = 3000 \text{ ft/min} \approx 15.4 \text{ m/s} \tag{8.4}$$

Using the the change in round trip time given in Equation 8.2, we can now calculate the unit change in round trip time, and so the value the attacker should aim for to trigger an alarm in this scenario, as follows:

$$\Delta t_{rtt} = \frac{2(h_2 - h_1)}{c} = \frac{2s_{closure}}{c} = \frac{2 \times 15.4}{c} \approx 1.03^{-7} \text{s/m} \tag{8.5}$$

Compared with a 'typical' approach at 3°, thus descending at 700 ft/min (i.e. 2.375×10^{-8}s/m), the attack would see the system calculate a considerable perceived increase in rate of descent. The attacker may only have to invoke this for a short period to trigger an alarm.

Another approach is to flood the aircraft with many signals within the radar frequency range, either at a specific frequency or over a random distribution. The behaviour of the radio altimeter in this scenario is unknown as it would be receiving effectively unpredictable pulses. However, such an approach is easier to undertake than the previous one, as it does not require crafted signals.

Requirements

To achieve this attack the attacker will need a number of directional antennae underneath the approach path or with the ability to transmit signals to the radio altimeter reception antenna on the aircraft. These will be fed by SDRs and the relevant software capable of transmitting the signals for the attack. A more sophisticated set up could listen for radio altimeter signals and only transmit when these are received, or use data from flight trackers to identify incoming aircraft. Although an attacker could operate such a system remotely, the hardware would need to be located near to the runway.

Feasibility

Since an attacker would need an SDR and antennae capable of covering an area underneath final approach directly below the aircraft, the ability to deploy depends on the airfield security and perimeter size. For example, larger airports may close off more of their approach path, but in practice this is still a small area relative to the full approach path.

Even so, this attack can take place a way away from the end of the runway. Using a representative 130 knots landing speed on a 3° glideslope (thus descending at 700 ft/min), an attack at 500 ft AGL would occur around 1.7 miles from the end of the runway.[1]

With regard to system resilience, radio altimeter for GPWS is known to have suffered interference from outside sources in the past and is a topic of concern. In [231], instances of unexplained GPWS alarms on approach to an Israeli airport were later explained by emissions from a nearby military radar. Interference on radio altimeter frequencies has been under discussion for a number of years due to the importance of the instrument, and is still actively being addressed [232]. These examples suggest that the system is vulnerable to external interference. Given our threat model of a determined attacker, we feel that this attack could be implemented under real conditions.

Simulator Implementation

For the purposes of our experiment we emulate the attack by triggering the GPWS '*Terrain, Terrain, Pull Up*' alarm starting 500 ft AGL on approach to Runway 33 at Birmingham, increasing by 250 ft for each subsequent attack. With this we are emulating the ability of an attacker to add some unpredictability to the attack. One of the limitations of this approach is that the point at which the attack actually triggers can vary between 450 ft and 500 ft AGL, and the radio altimeter visual does not show an 'under attack' change for the time under attack.

[1] This is calculated based on the aircraft taking approximately 42 s to reach the touchdown zone based on a vertical speed of -700 ft/min for a 3° glideslope.

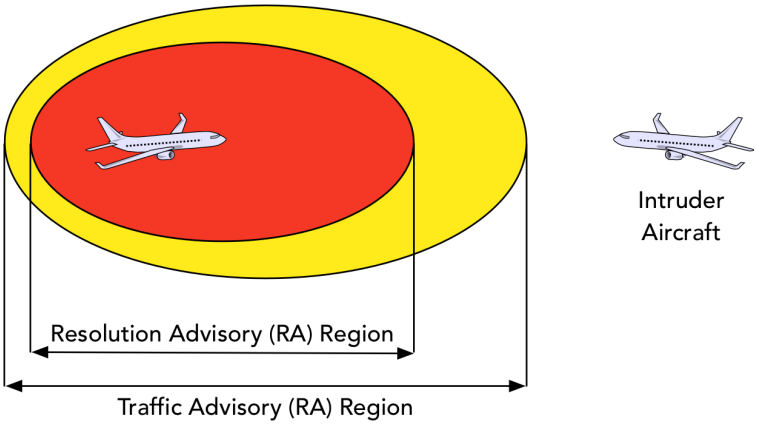

Figure 8.8: Representation of TCAS Traffic (TA) and Resolution Advisory (RA) zones.

Expected Response

Whilst the specifics of the response will depend on the aircraft and airline, there are common principles [233]. In most conditions, the response would be to perform a steep climb to a safe altitude, known as a terrain avoidance manoeuvre. In the simulated scenario, this will lead to a missed approach. However, below 1000 ft above aerodrome level (AAL), with full certainty of position, crew can choose to not follow this. Due to the surprise element, we expect the average response to be a missed approach. On following approaches we expect participants to have identified unexpected behaviour and disregard the warnings.

8.3.2 Traffic Collision Avoidance System

Although ATC manage airspace with high precision, aircraft can still end up closer than is safe. This is called a *loss of separation*, and in the worst case, can result in a mid-air collision. One less extreme example occurred in March 2011, where a Delta aircraft took off with an inactive transponder, becoming too close to three other aircraft before resolving the issue [234]. Traffic Collision Avoidance System (TCAS) provides a technical means by which to avoid collision due to loss of separation, and has been mandated on aircraft with more than 30 seats since 1993 [106, 235].

System Description

TCAS, as covered in Chapter 2, makes use of the Mode C or Mode S transponders fitted to an aircraft (the *object* aircraft in this chapter) to interrogate nearby aircraft [235].

Establishing nearby aircraft with Mode S simply requires the object aircraft to listen for Mode S 'squitters', which are messages transmitted in response to ground-based Mode S interrogations. These contain ICAO transponder IDs, so the object aircraft follows up with Mode S interrogations to establish the position of nearby aircraft, referred to as *targets* in this chapter. The range, heading, altitude data is provided via Mode S by the nearby aircraft and sourced from its instruments. The potential for conflict is then calculated on the object aircraft. Depending on the proximity and closing speed of the target the interrogation rate will vary; at a large distance this will be once per five seconds, increasing to once per second when an aircraft is close [236]. An abstracted protocol diagram for Mode S can be seen in Figure 8.9a.

Mode C operates slightly differently, represented in Figure 8.9b. The object aircraft will issue Mode C-only all-calls, causing nearby aircraft with Mode C transponders to respond at a rate of once per second. If the target has an altimeter then it will respond with its altitude, else TCAS onboard the object aircraft will use response characteristics to estimate altitude as well as range and bearing [237]. TCAS will only provide full alerting if Mode C-equipped aircraft provide altitude. Due to the length of a Mode C-only all-call, TCAS messages from nearby aircraft can cause message garbling. The system handles this with a technique called *whisper-shout*, wherein an interrogation gradually increases in power, preceded by a short, lower power suppression signal. This suppression prevents aircraft who have already responded doing so again.

Through one of these methods, TCAS ascertains how close the target aircraft is both laterally and vertically, before deciding if it is necessary to alert the flight crew. For most systems, especially those on commercial aircraft, alerts are composed of two steps as shown in Figure 8.8. First comes a *traffic advisory* (TA), in which the

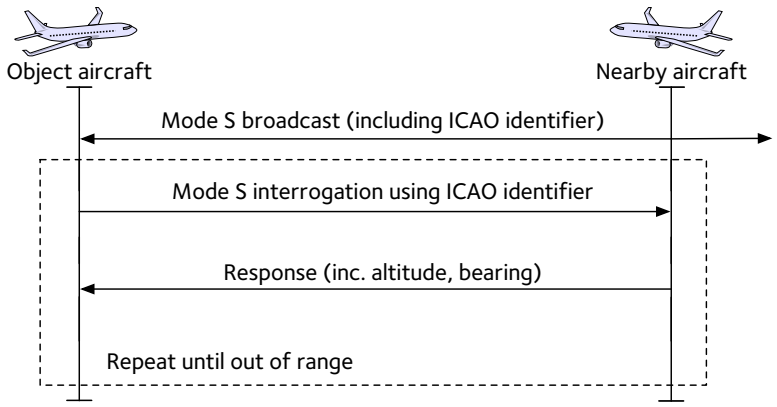

(a) Protocol diagram for TCAS interrogation using the Mode S data link, where nearby aircraft respond with information on their position.

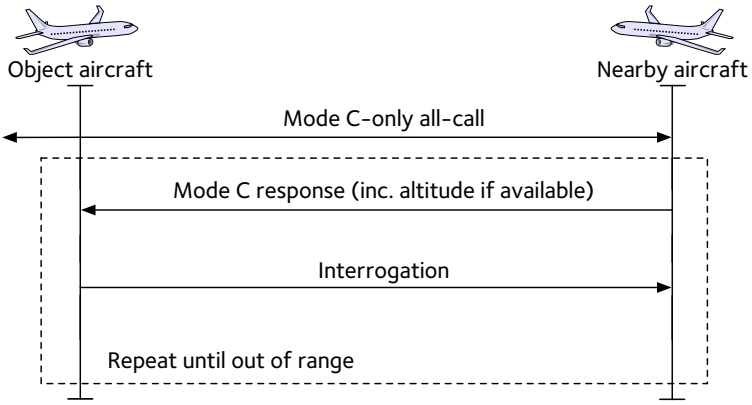

(b) Protocol diagram of TCAS all-call interrogation using Mode C, and response from nearby aircraft with altitude if available. Range and bearing are calculated from response.

Figure 8.9: Representation of TCAS interrogation protocols of nearby aircraft using Mode C and S transponders.

traffic is typically displayed to the pilot as amber and an aural alert of 'traffic' is given. If the intruder becomes closer to the aircraft, a *resolution advisory* (RA) is given. An RA will contain specific instructions for the flight crew, i.e., to climb or descend at a given rate, or hold vertical speed. These instructions are decided between the two aircraft automatically and aim to deconflict the situation. Crew must follow the instructions of an RA within seconds.

In the cockpit, crew have some control over the sensitivity level; they can select *Standby*, *TA-Only*, or *TA/RA*. For most of a flight, TCAS will be set to TA/RA, which automatically calculates sensitivity based on altitude. TA-Only is limited to the lowest sensitivity level and does not issue RAs, whereas Standby performs

no TCAS interrogations and will not resolve conflicts [238].

Whilst in TA/RA, TCAS will calculate the sensitivity based on altitude, with higher altitudes assigned higher sensitivities. This then defines the *tau* value for issuing a TA or RA. Tau is calculated as the time in seconds to the Closest Point of Approach (CPA) between object and target aircraft, either laterally or vertically. When the target aircraft is within tau, the relevant alert is given.[2] For example, between 5000 and 10,000 ft, tau for a TA is 40 s [238].

Attacker Aim

In this scenario, the attacker is aiming to cause arbitrary crew responses to TCAS through triggering false TAs and RAs despite no aircraft being present. This is intended to burn unnecessary fuel, break out of assigned flight levels or clearances and seed doubt in TCAS due to unnecessary RA manoeuvres. This may then result in diversions or switching off the system.

Attack Description

To achieve this attack, an attacker is using the fact that Mode C and S transmissions are sent in the clear with no authentication. This attack is in a similar vein to the work carried out on ADS-B, specifically in [20] and as discussed in Chapter 2. Here, the authors outline theoretical attacks on ADS-B, one of which involves crafting ADS-B packets to create a false aircraft. Since the most widely used ADS-B implementation uses Mode S, we translate these ideas to TCAS.

An attacker will generate a sequence of TCAS responses for a false intruder aircraft which is approaching the object aircraft. These will be generated in such a way to gradually approach from a distance, with the closest point of approach being comfortably within the tau for that altitude. We will refer to the aircraft under attack as *target* and the injected aircraft as *false*.

We firstly presume that we can establish the altitude, heading and speed of the target aircraft from either ADS-B or Mode S messages. Then, the injection

[2]Some adjustments are made to this at lower altitudes, and are covered in detail in [238].

approach will vary slightly between Mode S and Mode C, due to the differing interrogation strategies:

- For Mode S, the attacker would need to transmit a squitter message as if the false aircraft was identifying itself by Mode S. When the target aircraft then responds with subsequent interrogations, the attacker can transmit Mode S responses as if the false aircraft were travelling towards the target, such that the CPA is close enough to trigger a TA then RA.

- For Mode C, the lack of squitter means that the attacker needs to respond to an all-call. The whisper-shout mechanism might cause interrogations to be hard to receive by the attacker, in which case they would need to approximate a response. However, this would be stochastic—the interrogation rates at different ranges are standardised.

A different but unpredictable approach would be to flood the frequencies with Mode C and Mode S responses from a false aircraft. This would be much more unpredictable as it might not align with Mode C/S interrogations, thus not cause the intended effect. Furthermore, the 1090 MHz link is liable to overcrowding thus resulting in message loss [82]. As such, flooding the frequency could simply jam the frequency or eliminate other messages.

Requirements

Transmission by the attacker would require an amplifier and antenna capable of directional transmission, which are readily available; a comprehensive set up might cost £10,000 but could be achieved for less. The attacker would also require two pieces of software; a management system to calculate the required messages to cause TCAS alerts, and a Mode C/S transceiver. A transceiver is needed to both receive interrogations establishing the target aircraft behaviour, and to transmit messages as if from the false aircraft. Management software would then create the attack by establishing the position of the aircraft and required set of messages to cause an alarm.

Physical location is important to maximize range. Indeed, the attackers will also need to choose a place where their exposure to the target aircraft is highest; ideally under an area of dense traffic or on raised ground. Such an attack might be most effective close to airports as the aircraft are at lower altitudes.

Feasibility

By design, Mode C and S messages can be received on the ground—they are intended to be collected by ground-based radars. To achieve this, they use a relatively high transmission power of up to 250 W with directional antennae, meaning that reception from the ground is possible for an attacker [239]. Because of the aircraft being above ground-based obstructions, the potential area in which the attacker can reside is quite large.

The Mode S attack will be easier to carry out as the target aircraft requires the false aircraft to provides range from its own instrumentation, rather than calculating it based on the signal. This would prevent the target calculating the range based on round trip time, which would constrain the attacker.

One of the main feasibility challenges of this attack is that the aircraft may quickly move out of the range of the attacker due to its high speed. A well-resourced attacker might deploy antenna to multiple locations, whereas a more simplistic attacker could instead seek higher ground to maximise range.

For software, some related work in this area exists. Whilst many decoders exist for Mode S signals such as `gr-airmodes` or `dump1090` [240–242], encoding tools are outnumbered by decoders. One example is ADSB-Out, which is intended for use with ADS-B IN systems such as Stratux [243, 244]. A similar tool would be needed for Mode S. Producing this part of the software is arguably the most challenging task, however if the hobbyist community produces an open source tool this will lower the barrier somewhat.

Furthermore, some theoretical analysis of the TCAS II logic suggests that attacks creating situations similar to transponder failure can have a range of effects [245]. Most related to this scenario are intermittent Mode C transmissions or duplicated

Table 8.1: Simulation parameters for TCAS attack. Inter-attack gap is the time between rounds, approach distance is the distance from which a false aircraft approaches at the stated speed, and altitude difference is relative to the target aircraft. Units are in-simulator measures of distance, which do not necessarily correspond to metres or feet.

Attack Round	62Inter-attack Gap (s)	Approach Distance (units)	Altitude Difference to Target (ft)	Speed (Units)
1	10	10000	100	7.0
2	10	20000	-100	10.0
3	10	20000	100	10.0
4	10	15000	-100	7.5
5	10	15000	100	10.0
6	10	10000	-100	7.0
7	10	15000	100	10.0
8	10	20000	-100	7.0
9	10	20000	100	7.5
10	10	20000	-100	10.0

Mode S addresses. For the former, the target aircraft may trigger late alerts on the object aircraft due to the system treating it as not providing altitude. Of the latter, TCAS will ignore the more distant duplicate address, which could be used to deny situational awareness by an attacker.

Simulator Implementation

Within the simulator, we use a strong attacker who covers a large geographic area, attempting to trigger 10 alerts over the course of the flight. To avoid predictability, we varied the angle of approach by the false aircraft, and the speed at which it moved towards the target. For the purposes of measuring response, we configured these to be consistent for each participant, with the parameters provided in Table 8.1.

False aircraft were injected when the target aircraft flew above 2000 ft, after which the first injection began. If the participant chose to turn the TCAS sensitivity to TA-Only, they would still receive TAs but not RAs.

This attack was implemented using an invisible aircraft model, moved by the attack module, which travelled towards the target aircraft. Whilst reliable, this did create some limitations. Since the aircraft was not powered it would change altitude rapidly when outside of TCAS alerting range, sometimes doing this whilst visible to

participants. This could be fixed in future work by moving the object far away from the aircraft once it has passed the CPA. Furthermore, a limit in the A330 model used meant that crew would still see RA visual indicators if the transponder was set to TA-Only. This is much harder to address without modifying the aircraft model; due to rights management software, this was not possible for the model used.

Expected Response

Since following an RA is a compulsory action, we expect that most pilots will comply with at least the first instance [246]. This will result in them following the instructed manoeuvre. From there onwards, we expect some participants to begin to doubt the RAs and eventually turn the alert level of the system down to TA-Only. On average, we expect participants to follow the first 3-4 RAs before reducing the alert level or switching the system off.

8.3.3 Instrument Landing System

The Instrument Landing System (ILS) allows precision landings to take place under all conditions, for example under instrument rules when visibility is very low, or cloud cover is dense. In the most extreme cases, ILS allows aircraft to perform automatic landings at sufficiently equipped airfields.

System Description

As covered in Chapter 2, ILS consists of two components: localizer (LOC) and glideslope (GS) [247]. A localizer provides lateral guidance and alignment, centred on the runway centreline, whereas the GS provides vertical guidance to the touchdown zone on the runway. Typically, the GS will provide a 3° approach path, depending on the specific approach and airport. It is supplemented by Distance Measuring Equipment (DME), which provides the direct distance to a beacon located on the airfield, without directionality.

Transmission power of both the glideslope and localiser is moderate to low, at around 5 W for GS and 100 W for LOC [247]. On the carrier frequencies for the GS and LOC, overlapping lobes modulated at 90 Hz and 150 Hz provide guidance;

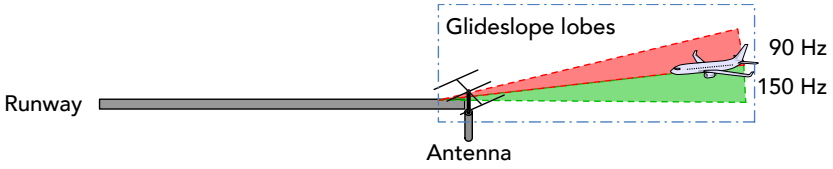

(a) Glideslope under normal operation.

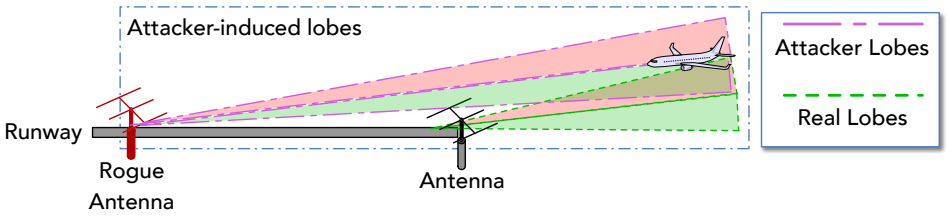

(b) Glideslope under attack with rogue antenna. Note how the aircraft touchdown zone is now at the far end of the runway. This means that if the glideslope is followed to touchdown, there may not be enough runway to slow down.

Figure 8.10: Representation of normal and under-attack glideslope operation, based on diagrams from [247].

the overlapping region represents the correct approach path. The aircraft will use the relative strength of these lobes to identify where it is with respect to the optimal glideslope and centreline of the runway. A diagram of a glideslope can be seen in Figure 8.10a.

Different types of ILS exist, known as CAT I, II and III. These are defined based on weather conditions, with CAT III enabling the landing of an aircraft in very poor visibility, such as when visibility is down to 200 m [248]. With each upgrade of category from I to III, the ILS installation has further constraints on accuracy and monitoring. Generally, the higher the category, the more protected the ILS is and accurate it must be. Glideslopes and localisers are monitored for accuracy to at least 10 nmi beyond the runway threshold, as well as being protected from interference to 25 nmi [66, 70]. It is important to note that here, protection from interference means avoiding other systems using nearby frequencies, rather than interference in a security sense.

Since aircraft must follow specific arrival patterns into an airport, ILS is an important part of managing pilot workload and is the default approach type for most airports. If ILS is not available, other options exist including area navigation (RNAV), which is based on GPS, surveillance radar approach (SRA), which relies

on ATC feedback, or reverting to a visual approach. Furthermore, the LOC and DME components of the ILS can be used without GS if needed.

Although not part of ILS, approach lighting provides an out-of-band check for crew on their approach path. One common form of approach lighting is the Precision Approach Path Indicators (PAPIs), which are configured to match to the angle of the glideslope. When an aircraft is on the correct GS, the PAPIs will show two red and two white lights; if the aircraft is too high it will show three or four white, and too low will be three or four red [249].

Attacker Aim

For this attack, the attacker is aiming to cause unnecessary go-arounds as a result of a tampered glideslope. In turn, this will use additional fuel, introduce delay and potentially force the target aircraft to divert to a different airport. Furthermore, a supplementary aim might be to seed doubt in the ILS to force crew to revert to a different approach method.

Attack Description

The technical component of this attack is arguably the most straightforward of the three, since the attacker aims to replicate the real glideslope but with the touchdown zone in a different location. Since the attacker will not be able to station themselves on the runway, they will have to locate away from the airfield perimeter. As the legitimate GS signal already is transmitted from aside the runway, transmitting a false signal from further aside the runway may be acceptable. As such, an attacker would transmit a false GS with a touchdown zone short or long of the legitimate touchdown zone, using antenna to the side of the runway.

Crucially the signals would be the same as a real GS, so would not be identifiable by a high rate of descent as reflected lobes are. Indeed, a marked difference is induced by the attacker but may not be spotted immediately; if we move a typical 3° GS by 1 km, the height difference between the real and false GS at a given

point will be approximately 52 m, or 172 ft.[3] Along the approach, this could fall within a margin of error.

Requirements

To carry out this attack, an attacker will need an SDR, amplifier and directional antennae to replicate the antennae arrays used for the legitimate GS. Software for this would need to be created but is likely to be achievable by a moderately resourced attacker as it involves implementing a standardised system. Unlike the other attacks, this would not require reception or to be reactive to what the target is doing.

Since the transmission power of a legitimate GS is typically less than 10 W, this is achievable with consumer amplifiers. For reference, even the lowest level of licensed UK amateur radio operators can transmit in frequency bands close to aviation frequencies at powers of up to 10 W [250].

The attacker will have to locate relatively close to the airport perimeter as in the GPWS attack. This may be easier for smaller airports with one runway, where it is easier to position close to the runway but off airport premises.

Feasibility

Considering technical feasibility, no significant barriers exist for this attack. This is possible due to the simplistic nature of the system—whilst the systems are monitored for integrity as defined in ICAO Annex 10, this is for deviations in the legitimate signal rather than malicious interference [251]. An ILS system will normally shut down or notify ATC if excessive deviation is identified.

Although extensively used and relied upon, ILS signals face challenges due to their relatively simplicity. The two particularly relevant to this work, as described in [252], are:

- **False lobes**, which occur due to the GS signals reflecting off the ground. These will appear as legitimate glideslopes according to the instrumentation, but can be identified by their steep angle, typically in the range of $9 - 12°$.

[3] As covered in [66], some glideslopes may be slightly shallower or steeper, but 3° is common.

- **Interference** due to reflections off buildings or vehicles, or aircraft moving through areas critical to ILS [253]. This is often accounted for when flying an ILS, with pilots being taught to expect interference [254].

The main challenges in carrying out the attack instead lie in physical location and limitations introduced by monitoring. As discussed above, being proximate with the runway has a large impact on whether the attack can take place. This is likely to be more difficult at large airports with high levels of security, and the ability to carry out the attack without detection somewhat depends on this.

Whilst potentially more effective, the ability to create a false GS where the touchdown zone would be short of the runway is unclear. This is due to the lack of public information on GS monitoring positions; attacker signals in this instance may interfere with the legitimate glideslope and cause an effective deviation, thus shutting the ILS down. We instead consider the 'long' version of the attack, where an attacker moves the intersection of the GS and the runway further from the threshold. This is done to cause go-arounds by leaving the aircraft too high to correct the approach and land with enough runway to stop. However, this does introduce a limitation in that it relies on aircraft intercepting the GS from above, which whilst not preferred over interception from below, is a valid method of intercepting the glideslope [66].

Simulator Implementation

In the simulator, we produce a simplistic version of this attack. An attacker transmits a false GS at the far end of the runway with an effective shift of 2.05 km (1.27 miles) creating a height difference between the false and true GS of 107 m (352 ft). Due to the way in which ILS is implemented in the simulator software, we could not replicate also having a 'real' glideslope. To account for this we operated on an assumption that the attacker transmits at a higher power than the real GS in an effort to force capture on to the false GS.

The manipulation remains in place regardless of how many approaches are made. We treat the participant aircraft as if it is the first to encounter the attack, with ATC not observing previous aircraft having difficulties.

Expected Response

Since this attack will see the attacker GS track slightly above the real GS, it is unlikely to be immediately obvious that it is incorrect. We expect most participants to follow the GS until they are below cloud at around 1000 ft, at which point they will notice a continued slight discrepancy in the AGL according to approach charts. They may also notice such a discrepancy using the PAPI, as they will show four white lights. At this point, we expect them to be between 500–1000 ft AGL and opt for a missed approach and go around.

8.4 Results

We discuss the data collected from the simulation and through interviewing participants after each scenario. We provide interview response data for all scenarios in Table 8.2 and Figure 8.11. The full data for Figure 8.11 is included in Appendix I.

For reference, participant responses are recorded on the following scales:

- **Confidence** in the response being the correct one, on a scale from 1, *very confident*, to 5, *very unconfident*.
- **Workload** due to the attack, on a scale from 1, *no increase*, to 3, *significant increase*.
- **Trust** in systems affect due to the attack, on a scale from 1, *much more trust*, to 5, *much distrust*.
- **Impact** on the flight due to the attack, on a scale from 1, *significant impact*, to 4, *no impact*.

Full scales and questions asked can be seen in Appendix H, and we allowed responses in between scale points.

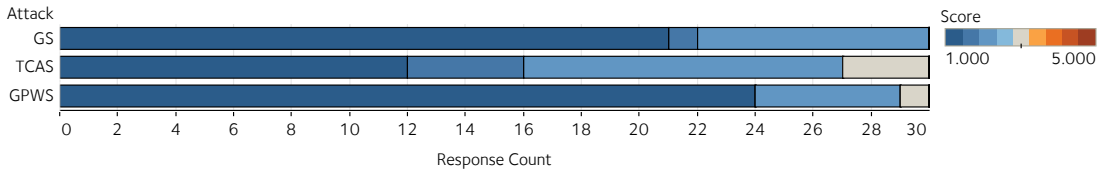

(a) Confidence in response, on a scale from 1, *very confident*, to 5, *very unconfident*.

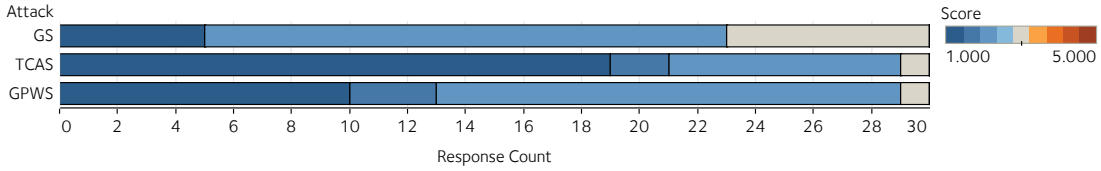

(b) Trust in systems post-attack, on a scale from 1, *much distrust*, to 5, *much more trust*.

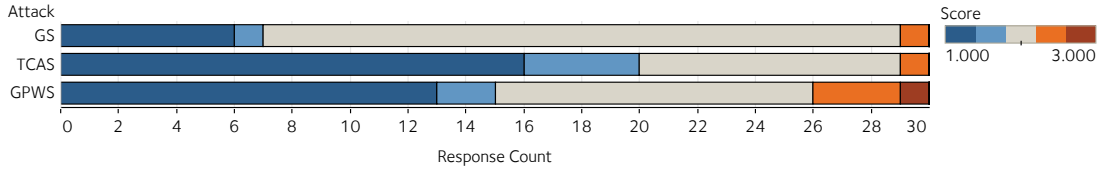

(c) Workload change due to attack, on a scale from 1, *significant increase*, to 3 *no increase*.

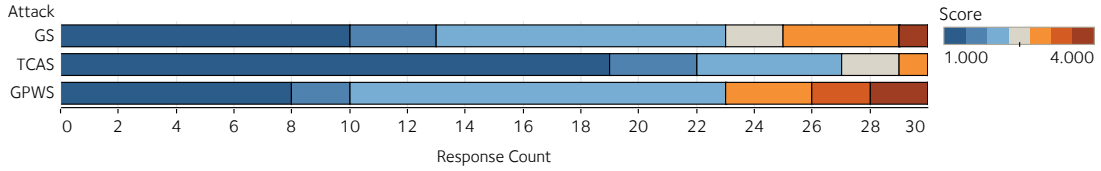

(d) Impact due to attack, on a scale from 1, *significant increase*, to 4, *no increase*.

Figure 8.11: Stacked bar charts for participant scale responses on confidence in response, trust in systems, workload change and impact due to the attack. Orange colours represent the most 'negative' responses, i.e. no effect/change, with blue the most 'positive' responses, i.e. significant effect/change. Data on which this is based is provided in Appendix I.

8.4.1 GPWS Radio Altimeter Spoof

First, we look at the GPWS scenario. We assess attack feasibility based on participant actions, i.e. perform a missed approach and a go-around, land, switch GPWS off, before considering their scale responses.

Control Response

Participants generally responded as expected, with a split between those opting for a terrain avoidance manoeuvre, thus a missed approach, and those disregarding the warning in order to land. The first approach is plotted in Figure 8.12. In Table 8.3, we can see that two-thirds of participants went around on the first approach as a

Table 8.2: Summary of participant actions and responses to yes/no questions as part of the debrief interview. For some participants, the question was not applicable due to previous actions e.g. having ignored an alarm and landed, hence N/A. Percentages are of all participants, for each question.

		Response					
		Yes		No		N/A	
Attack Type	Question	Count	%	Count	%	Count	%
GS	Trust later	1	3.3	25	83.4	4	13.3
	Less safe	19	63.3	11	36.7	-	-
	Same in real AC	28	93.3	2	6.7	-	-
TCAS	Trust later	4	13.3	22	73.4	4	13.3
	Less safe	28	93.3	2	6.7	-	-
	Same in real AC	30	100.0	0	0.0	-	-
GPWS	Trust later	0	0.0	12	40.0	18	60.0
	Less safe	14	46.7	16	53.3	-	-
	Same in real AC	27	90.0	3	10.0	-	-

Table 8.3: Division of action taken by participants in response to the GPWS attack. If a participant lands, they are not included in the numbers of the following approach. Percentages are of the participants involved in that approach.

		Action Count		
Approach	Action	#	%	# Participants
1	Land	10	33.3	30
	Go around	20	66.7	
2	Turn off	11	55.0	20
	Land	8	40.0	
	Go around	1	5.0	
3	Turn off	1	100.0	1

result of the alarm; these participants generally remarked that their choice was an automatic one. This is crucial as it shows that an attacker who can trigger such an attack can cause arbitrary go-arounds with reasonable chance of success. In one instance, the attack triggered late; however, in debrief, the participant noted that they would have had the same course of action regardless.

On the first approach, we found that for those opting to go around, the mean height at which the go around began was 403.9 ft, with a standard deviation of 51.1 ft. We plot this in Figure 8.13; some outliers in the form of later responses do

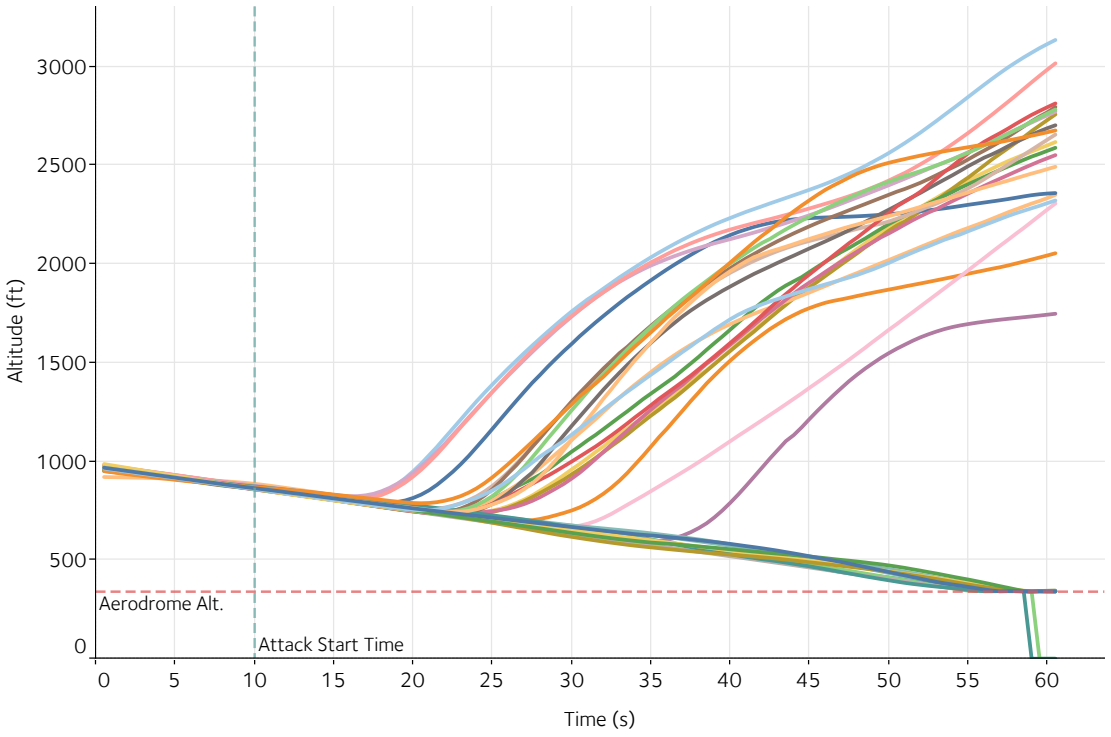

Figure 8.12: Plot of time against altitude for the first approach under GPWS attack. Each line is a participant. Note that 8 participants land and disregard the alarm, on account of being sure of their position.

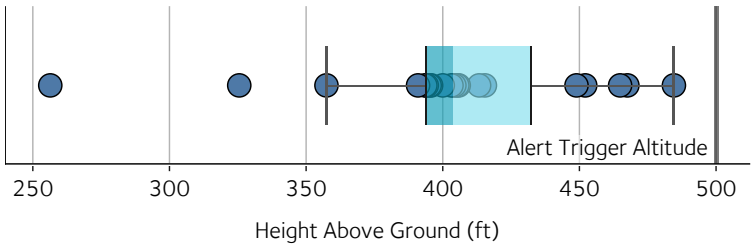

Figure 8.13: Box plot of minimum heights reached by participants opting to go around in first approach of the GPWS attack.

exist, shown in Figure 8.12. Most participants responded within 100 ft of the alarm with an interquartile range of 29.7 ft. This aligns with our expectation that pilots follow the instructed terrain warning and execute a well-drilled manoeuvre, not allowing the aircraft to become unsafe. It also matches attacker aims, exploiting GPWS alarms to cause go-arounds.

In terms of handling the attack, Table 8.3 shows that on a second approach, 11 (55.0%) participants chose to switch the system off due to it causing a distraction. One participant did this on the first approach and went on to land. As GPWS

is a safety system, an attack causing it to be switched off has the potential for erosion of safety on top of disruption—indeed, of the 12 who switched it off, none said they would trust the system later in the flight. Most participants on short final of the second approach sought to silence the alarm, though at this point they were sure of their position.

Perception

As seen in Table 8.2, 14 (46.7%) participants felt that this attack put the aircraft in a less safe situation. The numbers are lower compared to other attacks as the response is a safety manoeuvre, though some pilots felt that due to the extreme nature of the manoeuvre the aircraft is at additional risk.

This scenario has the least impact as assessed by the participant, as shown in Figure 8.11—it was judged to have 'some impact' on average, with 8 (26.7%) saying it was 'significant'. For workload, there was on average 'some increase' with 13 (43.3%) feeling there was a 'significant' increase. On top of this, some remarks were made about startling on what appeared to be a normal approach. Trust in the system was eroded during the scenario, which matches with our assessment in Table 8.2; 29 (96.7%) participants felt at least 'some distrust' of the system in this scenario.

Generally, participant confidence in their response was very high, with an average score of 1.28, or 'very confident'. The majority of participants (27, 90%) felt that they would take the same course of action in a real aircraft. Those who did not feel this way suggested they might have opted for a missed approach rather than landing.

8.4.2 TCAS Injection

Next, we consider the attack on TCAS. Results indicate that this is the most concerning attack to the participants.

Control Response

A summary of the actions can be found in Table 8.4. We provide the 'end-state' for the selected TCAS mode (e.g., if a participant selects TA Only then Standby,

they are in the Standby category) against any actions taken outside of normal flight. Actions are categorised into:

- *Continue on route* i.e. no extra action taken,
- *Avoidance manoeuvre*, in which the participant changes course slightly on top of the RA response, or
- *Divert to origin*, where they request to return to the departure airport.

Some 26 participants (87%) turned TCAS sensitivity down to TA-Only, with 11 (37%) switching to Standby. Participants switched to TA-Only with a mean 4.5 RAs (standard deviation 1.7), then down to Standby after another mean 2.8 TAs (standard deviation 2.1). In the meantime, participants were responding to RAs as normal. Two participants went straight from TA/RA to Standby, one after three RAs, another after six.

Reasons cited for these sensitivity changes were to reduce the additional workload caused by having to fly repeated TAs and RAs, and to remove the distraction factor of repeatedly having to handle to the alerts. Looking at the control response in more detail, three of those eventually turning the transponder to TA-Only and three of those turning it to Standby took avoiding action. The action itself varied per participant but for some, this involved climbing above the planned cruise altitude or making horizontal manoeuvres to try to avoid the attacker's traffic. Two participants diverted back to the origin airport rather than continue with malfunctioning TCAS and three of the remaining participants did feel that TCAS was providing spurious returns but felt the risk of downgrading the system was too high and instead opted to follow the RAs as issued, rather than turn the transponder to TA-Only. The final participant was not aware of the ability to go to TA-Only in the simulator and so remained in TA/RA.

Clearly, this attack not only showed the ability of the attacker to cause multiple arbitrary RAs, but also to push pilots to not use the system due to distraction. The range of responses shows that the attack lies in a procedural grey area, i.e. that pilots could handle it safely but not remove the disruption it caused.

Table 8.4: Participant response to the TCAS attack scenario, mapping the final selected TCAS mode against actions or manoeuvres taken by the pilot. Percentages are of all participants.

	Final Selected TCAS Mode							
	TA/RA		TA-Only		Standby		Total	
Action	Count	%	Count	%	Count	%	Count	%
Continue on route	4	13.3	10	33.3	8	26.7	22	73.3
Avoidance manoeuvre	0	0.0	3	10.0	3	10.0	6	20.0
Divert to origin	0	0.0	2	6.7	0	0.0	2	6.7
Total	4	13.3	15	50.0	11	36.7	30	100.0

Perception

Assessing the impact, Figure 8.11 shows that 27 (90%) pilots felt that the attack had at least 'some impact', with 19 (63%) feeling that it had 'significant impact'. This was coupled with 29 (97%) feeling that there was at least 'some increase' in workload. Typically, this increase in workload was due to having to respond to regular RAs and dealing with periodic distraction. An unduly increased workload creates further problems for the crew managing the situation, and can lead to errors.

Looking to perceived safety, 28 (97%) pilots felt that the attack put the aircraft in an unsafe—or potentially unsafe—situation. Some were concerned for passengers, who might be moving about the cabin, thus injured in an extreme manoeuvre such as an RA. Others noted the possible effects on nearby traffic in the event of following a spurious RA, such as triggering an RA for them too. Similarly, 29 (97%) of participants felt that they had at least 'some distrust' in TCAS during the scenario. Again, this is problematic as it indicates that an attacker can increase workload and seed distrust in critical aircraft safety systems.

8.4.3 Glideslope Spoof

We now look at the glideslope spoof, where an attacker aims to capture a pilot on a false GS.

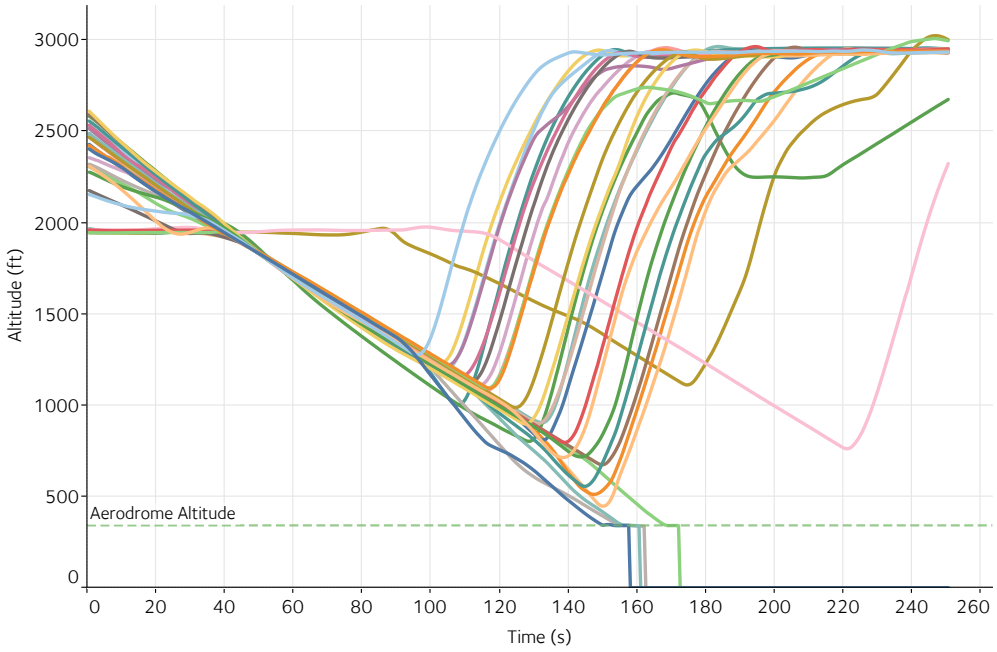

Figure 8.14: Plot of time against altitude for the first approach under glideslope attack. Each line is a participant. Note that four participants opt to correct the approach path and land.

Control Response

For our analysis, we focus on the first approach when participants knew least about the attack. On encountering the attack, 4 (13.3%) participants chose to land anyway on account of having a good visual picture. Of the 26 (86.7%) participants choosing to go around, three went around a further time. The plot in Figure 8.14 displays the path of each participant on the first approach. Eventually, all participants identified some issue with the glideslope—some identifying the issue quite accurately—before choosing a different approach type. The choices after participants identified a problem were as follows:

- 1 (3.3%) used a VHF Omnidirectional Range (VOR) approach,
- 2 (6.7%) used a Surveillance Radar Approach (SRA), which relies on higher involvement with ATC,
- 8 (26.7%) flew a localizer only approach (LOC DME) due to identifying GS problems,

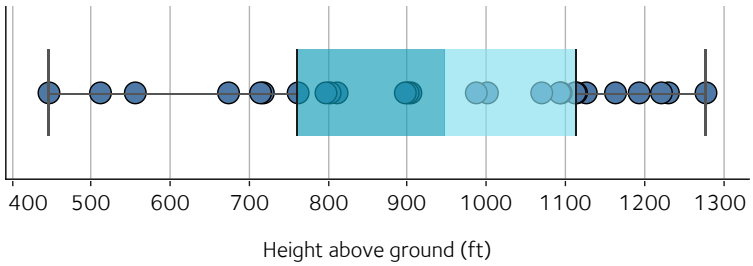

(a) Height above ground level at point of first go-around.

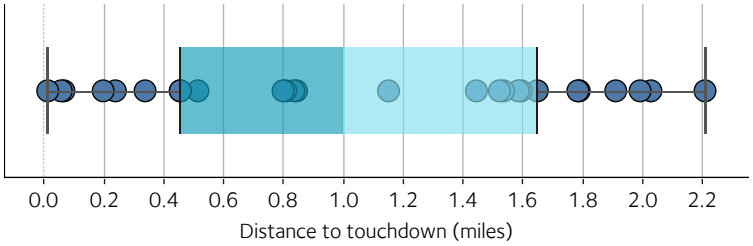

(b) Distance from runway touchdown zone at minimum height, i.e. point of first go-around.

Figure 8.15: Box plots of participants performing a go-around on the first approach under the glideslope attack.

- 9 (30.0%) dropped ILS completely, and used an Area Navigation (RNAV) approach, which is based on GPS,
- 6 (20.0%) flew a visual approach due to good conditions.

The split is interesting as it highlights that whilst participants could identify a problem, there was no consensus on the extent of it, i.e. whether it just affected GS or all ground-based systems. A reasonable portion of participants (11, SRA and RNAV approaches) chose to forgo ILS completely as they could not identify exactly what the issue was. However, eight were happy to use LOC DME since they were confident that it was just the GS affected.

In Figures 8.15a and 8.15b we can see box plots for the mean height above ground level (in feet) and distance from the runway touchdown zone (in miles), respectively. We include each participant opting to go-around in the chart, at the point of aborting the approach, i.e. the lowest height. The mean go-around altitude was 930.0 ft (283.5 m), with a standard deviation of 235.8 ft (71.9 m), and distance of 1.1 mi (1.8 km), with a standard deviation of 0.7 miles (1.1 km). If we account

for the fact that preparation for a go-around takes a few seconds, the mean point is just as the participants descended below 1000 ft (304.8 m). This is often a final check of stability on the approach and so becomes a forced point of decision.

Perception

As shown in Figure 8.11, 13 (43.3%) participants found that the attack had above 'some' impact, with the mean response being 'some impact'. Compared to GPWS, the GS attack had a lower perceived workload increase with 22 (73.3%) participants claiming 'some' increase. For GPWS, 11 (36.7%) have 'some' increase but a further 15 (50%) feel that the increase in workload was even higher. This could be due to the GS attack developing gradually, higher above the ground with PAPIs providing visual checks. Some participants noted that this attack would be harder to deal with in worse weather conditions due to poor visibility. Even so, it did cause 26 (86.7%) participants to perform a go-around through being unsure about the approach.

As with TCAS and GPWS, the attack caused 'some' distrust in aircraft systems, with 23 (76.7%) participants remarking 'some' or 'significant' impact. However, some participants correctly identified that the ground systems were at fault and so did not distrust aircraft systems. Furthermore, Table 8.2 shows that of the 26 (86.7%) participants who did perform a go-around, all but one did not trust the GS on a second approach. This is reinforced by the fact that 19 (63.3%) felt that the attack put the aircraft in a less safe situation.

8.5 Discussion

Clearly, each set of attacks covered in this paper has the potential to cause disruption, but with different levels of severity. If we consider impact as perceived by participants, TCAS was the most serious attack with an average score of 1.38, i.e. at the 'significant' scale point, followed by the GS attack where the average response was 1.85, i.e. 'some' impact. GPWS was considered to have the lowest impact with an average score of 2.03, also on the 'some' impact scale point. This aligns with our findings on control response, and the assessment on safety, workload and trust effects.

8.5.1 Attack Response & Safety Impact

We now consider the response for each system in turn and the impact on safety that the attacks had. Generally, we note that minimal safety loss occurred but impact manifested itself in potential for disruption and added workload.

TCAS. Here, the typical response was to follow RAs then reduce the sensitivity level of the system. As an important safety system, reduced functionality due to an attack is problematic. Having aircraft respond to multiple RAs in busy airspace would cause significant traffic issues to ATC. Even a single RA impacts the wider system, since during the RA procedure, ATC are not allowed to direct the aircraft involved, so effectively lose control of them temporarily [255]. One of the main concerns with this attack is its easy repeatability and scalability, since it does not appear to have extensive physical constraints on the attackers.

Glideslope. Although subtler, this saw over 85% of participants abort their first approach. Although some participants did not identify the exact problem at that stage, the indication was enough for the majority to abandon the GS and choose another approach method. Clearly, we can see that this attack has the potential to cause financial loss to the airlines, as well as disruption and inconvenience since it would now have to land out of sequence, also burning more fuel. Participants noted that attempting to perform this attack on consecutive aircraft may see ATC instructing to avoid using the ILS, neutralising the effect.

GPWS. Whilst clearly of lower impact, it is interesting that the response appears to be determined by training and airline procedure. Some airlines mandate to always go around on this alert, whereas others allow the crew to make an assessment in the situation. Although many participants switched the system off—problematic for a safety system—this was done with certainty of position. Indeed, the go around is a safety manoeuvre. It is noted that a GPWS radio altimeter attack is hardest to completely defeat, since the system relies on fast, non-data carrying pulses. This could be difficult to add meaningful authentication to.

8.5.2 Effect on Safety

On balance, our expected response for each attack was that participants would take the 'safest' option in the circumstances which typically led to disruption. Because of this, we did not necessarily expect much perceived loss of safety. Our results indicated that there was some perceived loss in safety; for TCAS and GS, 93.3% and 63.3% felt the attack made the aircraft less safe. For TCAS, this is likely due to the uncertainty of the situation, with pilots not expecting false alarms. In the case of GS, the safety concern comes from how late the discrepancy is apparent and the situation this leaves the aircraft in. The exception to this was the GPWS attack, in which the terrain avoidance manoeuvre is the de facto safe option so fewer pilots felt safety was affected. Arguably, this highlights the interplay of safety and security the most. Even though most pilots took the safe option, they still felt they were compromised by factors out of their control.

Furthermore, the range of decisions taken indicates that these attacks do indeed create grey areas for responses. In an industry where safety is paramount and relies on well-defined procedure, the ability for an attacker to create situations open to interpretation is concerning. It is especially so when considering that these are systems tied closely to aircraft safety.

Even so, all pilots handled the scenarios in a way which did not severely compromise safety and cause—or create high risk of—collision. This is an important takeaway from the results as it demonstrates that the pilots, as humans in the loop, provide considerable mitigation against attacks compromising safety. However, this is not to say that the attacks compromise safety under no circumstances. As discussed below, some factors may amplify the workload or targeted attacks may create situations in which a reduction is safety is inevitable—identifying and defending against these scenarios is an avenue of future work.

8.5.3 Cost of Disruption

We have demonstrated the ability for these attacks to cause missed approaches and diversions. With this in mind, we can estimate the cost of this attack on an

aircraft. As an example, we use Boeing aircraft due to public fuel usage information to calculate the cost of a missed approach. For their smaller 737-800 aircraft, the missed approach uses 127 kg (41.79 gal) more fuel than a successful one; for the larger 777-200, it is 399 kg (111.55 gal) more [256]. Coupled with a nominal jet fuel cost of 184.58 c/gal, this costs approximately $77 for the 737 or $205 for the 777.[4]. Added to the expense of further time in the air—more difficult to predict as it depends on factors such as the airfield and traffic—plus a second approach, which costs approximately $139 (using 230 kg, or 75.68 gal) or for the 737, or $516 for the 777 (using 850 kg, or 279.69 gal), this becomes expensive for the airline.

Diversions add further expense, with potential effects on scheduling or passenger inconvenience. The UK Civil Aviation Authority estimates that these can cost an airline between £10,000–80,000, depending on the size of the aircraft and location of diversion [258]. For example, passenger disruption causing diversion aboard a Norwegian flight cost €100,000 in 2018 [259, 260]. Closed airports are similarly costly, with drones shutting London Gatwick for two days in December 2018 and costing airline Easyjet £15 million [226]. From these representative figures and our results, it is clear that non-destructive attacks are both possible to enact and can have severe economic consequences, both at the time of attack and in the recovery phase.

8.5.4 Amplifying Factors

In debrief, participants raised a number of other factors which would affect the impact of attacks. Weather conditions were prominent; all scenarios would be more difficult to handle in poor visibility. Particularly for the approach-based GS and GPWS attacks, good visibility allowed participants to arrive at their decisions more quickly. Some participants noted it would be hard to identify the GS attack under automatic landing conditions, leaving much less time for pilots to respond.

Other contributing factors include tiredness and terrain. In response to the GPWS attack, one participant who chose not to go around commented that their action in a real aircraft would depend on tiredness, as well as weather and how

[4]Calculated using IATA Jet Fuel Price Monitor for 18^{th} January 2019 [257].

busy the crew were. Again in the GPWS attack, others identified that terrain surrounding the airport affects their choice—they would be much more likely to abort an approach in challenging terrain, and less if they are familiar with the airport.

8.5.5 Training Benefit

To assess whether responses were realistic, we asked each participant whether their response to each scenario would be the same in a real aircraft. We found that for:

- **GPWS**, 27 (90.0%) would do the same, and the remaining three would go around in the same scenario again,
- **TCAS**, 30 (100.0%) would do the same,
- **Glideslope**, 28 (93.3%) would do the same with the remaining two opting to go around and revert to RNAV.

We asked each participant for their views the value of such experiments or training in preparation for cyber attack. All participants felt the scenarios were useful, and 28 (93.3%) commented that training for cyber attacks using a simulator would be valuable. We also asked if they felt limited by the simulation set up, on a scale of '1–*not*', '2–*somewhat*' and '3–*heavily*' limited, with the average response corresponding to 'somewhat'. The main limits were lack of a second crew member, and the general (rather than Airbus specific) controls. We do note that these figures are subject to some bias due to the experimenter conducting this interview but we feel that the results are sufficiently strong that the effect on our conclusion is minimal.

The results suggest that this method can be valuable both in identifying crew response to attacks and providing cyber attack readiness. Furthermore, the fact that the scenarios in this paper lie in procedural grey areas and do not have a series of steps to resolve them provides an ideal opportunity for training. One point of caution is negative training, with some participants noting that care must be taken to avoid training pilots to ignore or distrust their systems.

8.5.6 Limitations

Since this study is a first-look at using simulation to assess cyber attacks, it has some limitations which should be addressed to improve the generalisability of future work.

Most fundamentally, further research in this area should use a certified flight simulator. This removes unfamiliarity for the pilots, instead placing them in an environment that they are heavily trained to use. In particular, this is important when it comes to physically interacting with controls; many actions in a real cockpit can be completed with dedicated buttons or switches. In a full simulator, more detailed response data could be measured since we would not have to account for unfamiliarity with the controls.

Following on from this, having a single-pilot set up as in this experiment increases the workload of the pilot and so limited the scenarios we could explore such as poor weather or attempts to seriously reduce safety. Extrapolating from a single pilot having to carry out the role of two pilots in an emergency situation would be unrealistic, especially in an unfamiliar environment. This was partly due to our simulator software and hardware only being able to take one set of inputs. Using a certified simulator as above would allow for a full crew.

Another improvement would be to recruit in a less biased way, ideally taking a randomised sample across different companies and interest levels in cyber security. For this study, we advertised to recruit participants and so had to indicate the topic; this naturally biases the sample. Future work could collaborate with an industry body or group of airlines to reduce this bias.

Finally, to examine multiple attacks for each participant we had to keep the flight short, and participants had exposure to multiple attacks; this may increase predictability of the scenarios. This is different to how normal simulator training is done, in which the emergency scenarios are part of a longer flight; future work in this area should look to use this approach instead.

8.6 Mitigations & Recommendations

Many of the security problems enabling the attacks in this chapter are due to a lack of security mechanisms in deployed systems. It is not feasible to redeploy systems with security inbuilt in the near-term due to high cost and long certification timescales. We instead describe potential shorter term mitigations and recommendations to reduce attack impact of these attacks.

8.6.1 Technical Measures

Handling interference from rogue signals is not a new challenge for aviation; the fundamental change is that in the past, this interference would be considered accidental rather than malicious. For example, Eurocontrol refer to potentially (and unlikely) malicious communications on radio frequencies as part of their Air-Ground Communications Safety Plan in 2006 [261]. As discussed, the bar for attackers is now lower and so we must consider how to defend against a modern, capable threat. A number of technical measures could help defend against a threat trying to enact the attacks in this chapter.

Spectrum Monitoring

Both the GPWS and GS attacks require the attacker to transmit signals near to the airfield. For GPWS, this is a localized directional signal typically pointing skywards; for GS, the signal is likely to be higher power than the real GS. Early warning of these attacks could be provided by carrying out spectrum monitoring around the airfield on key aviation frequencies. Looking for unexpected signal sources and high transmission powers could help to spot malicious interference. Some examples of this appear to already exist from established radio companies, however the cost and level of deployment remains unclear [262, 263].

Wide deployment could be achieved with low-cost sensor networks distributed around the airfield, however a reasonable investment into infrastructure to process and store the collected data would be needed. This kind of deployment would be in

a similar vein to the ElectroSense project which is building a crowdsourced wide spectrum monitoring operation using RTL-SDR-style sensors [264].

Crowdsourced Air Traffic Surveillance

A number of crowdsourced networks now exist to collect air traffic surveillance signals, namely Mode S. Key examples include Flightradar24, FlightAware and OpenSky Network [153, 179, 265]. At the hardware level, these consist of geographically distributed low-cost sensors operated by members of the public, who feed data to a central server. A benefit of this kind of network is the ability to cross-check the reception of messages, which in turn can be used to constrain a potential attacker. A theoretical foundation for this can be seen in [100].

An approach like this is very useful in detecting message injection as in the TCAS attack. A picture of the airspace provided by a highly-redundant sensor network could provide a good reference for ATC in trying to identify whether a message injection attack is taking place. For example, if an aircraft is reporting TCAS RAs by voice or data link, but PSR, SSR or crowdsourced networks do not corroborate this, it could be attacker interference.

Security Mechanisms by Design

Longer term, avionics need to be secure by design. Current trends suggest that next generation data links, likely to see early deployment in around 10 years, have some security by default. AeroMACS, the most developed example so far is a derivative of IEEE 802.16e (known as WiMAX) from which it inherits some secure communications [266–268]. Even so, existing technologies such as ACARS will be used for many years.

However, the systems in this chapter will not be replaced soon and security needs to be patched in or developed in some way. Both TCAS and GS/ILS have the opportunity to make such changes. As a minimum, systems such as TCAS need to ensure message integrity beyond basic error correction, but ideally also message authenticity. So far, there appears to be no work in this area. In the case of ILS signals, some research has looked into using Distance Measuring Equipment (DME)

for data carrying [269]. Since DME is used in conjunction with ILS signals, a logical extension might be to carry signal integrity and authenticity data on the ILS signals.

8.6.2 Policy Measures

As well as technical solutions, adjusting or adopting new policies can help with defence against these kinds of attacks.

Security-Focussed System Design

Each attack presented in this chapter exploits system design, specifically taking advantage of a lack of a security mindset during the development of each system. As shown, this can have a range of impacts when attacked, from limited as with GPWS up to severe with TCAS. Conducting vulnerability assessment as part the design process would have helped to identify security issues early on, preventing systems being developed and deployed with fundamental problems. Adopting this approach should see the assessment updated throughout development and deployment in an effort to identify vulnerabilities early.

Developing Attack-Aware Procedures

Responses to the attacks in our study varied, highlighting that the attack effects exist in procedural grey areas. For example in the glideslope scenario, once the glideslope itself was identified as faulty, pilots chose five different approaches for the next landing attempt. Such variability is at odds with other parts of cockpit procedure, which is normally prescribed through checklists or memorised sequences of actions.

Although attacks on avionics might not be predictable, some handling procedures to aid crew under attack could be developed. This could come in a number of forms but should aim to help to diagnose an attack—our results indicate that uncertainty about whether an attack is taking place can itself create variability in response. Ways to achieve this might include a list of systems which could be affected by attack or behaviours which suggest interference rather than normal faults.

Producing these will be slightly different to traditional safety procedure. They will need to be underpinned by a thorough security assessment of systems currently

deployed on an aircraft, so could be expensive and time consuming. Furthermore, they might not always be as absolute or specific as checklists for dealing with faults, for example. However, they could help to standardise the response to an attack to be the safest possible.

Training for Pilot Awareness

As demonstrated in this study, the participants handled the scenarios safely, despite disruption, added workload and some perceived loss in safety. However, the split in responses both in control and interview indicates room for procedure and training to handle attacks. Naturally, processes which help reduce startling and give experience in handling new situations can be beneficial as covered in Section 8.

To develop this into a training exercise, the described limitations would need to be addressed, especially incorporating a full crew and realistic ATC. However, 28 (93.3%) participants felt that this exercise was useful, since it exposed them to unusual situations. As above, it is important to avoid negative training, though this could be done through careful scenario design and context.

8.7 Summary

In this chapter we have explored how flight crew respond to realistic cyber attacks and have shown that they have the potential to cause disruption. In turn, this could lead to financial loss and reputational damage to operators. In some scenarios, we also found that pilots felt the attacks also caused a reduction in safety. Clearly, we can see that effort should be made to explicitly defend wireless systems used in aviation as humans in the loop cannot entirely mitigate the effect of attacks. Our results indicate that the attack on TCAS is the most immediately concerning due to creating inconvenience and the greatest potential safety reduction. Both ILS and GPWS attacks also pose problems, though were easier to identify and mitigate on the flight deck.

More generally, this work suggests that that flight simulation for cyber attack awareness or training has value. Since it is unlikely that preventative or by-default

security will be deployed in the near future, such training could be incredibly beneficial in an environment at increasing risk of attack.

There's never anything you do where you can't think of a way to do it better.

— Chris Boardman

9 Future Work & Summary

Contents

The work in this thesis has looked at some avionic systems which were yet to have their safety and privacy stance assessed. We have found that they mirror the position of systems covered in existing research on ADS-B. Namely, a lack of avionic security mechanisms included at the design stage is now manifesting itself as a problem due to the advent of cheap, easy-to-use radio hardware.

In the first part we presented research into ACARS usage by non-commercial aircraft, focussing on aircraft operating privately. By conducting a measurement study using commodity hardware and software, we demonstrated that many aircraft have grown to use ACARS in such a way that reveals positional information. Given many of these aircraft seek privacy in other ways, the lack of by-default

security—or lack of awareness of the link being unprotected—means that privacy efforts are undermined.

We then moved on to show that in cases where customer demand for securing data link messages existed, a proprietary cipher had emerged. We demonstrated that it was not fit for purpose, with an attacker able to recover message contents with ease. This part of the work is an important lesson in deploying unproven security solutions; relying on them can do more harm than good. Many of the messages carried were again position-related, undoing and contradicting efforts to hide aircraft from public flight trackers.

Finally, we simulated attacks on systems which form critical parts of flying an aircraft to understand their impact. We investigated whether flight crew would be able to neutralise any impact through existing procedure. Using a conservative threat model, we showed that whilst pilots coped well with the attacks, potential for disruption and a reduction in safety exists. In some cases, this would see the participants opt to avoid using a system due to the attack, which is especially problematic for safety systems. We also found that this approach has promise as a method to train crew in dealing with attacks.

9.1 Recommendations for Avionic Security

In Chapters 6, 7 and 8 we provide recommendations specific to the problem at hand, based on our findings. From these, we can derive common high-level recommendations which will be important in making progress towards secure avionics.

9.1.1 Effective and Adaptable Threat Modelling

Arguably the most fundamental component of securing systems is starting with a meaningful threat model. Understanding what to defend against and the likelihood of attack is vital in directing resources effectively. Since aviation has extensive experience of modelling safety risks, adapting experience to include security should be possible. The biggest challenge in this respect is developing an approach to simultaneously model threats from a security and safety perspective.

Threat modelling for security is relevant to all aspects of the industry from design of systems through to usage. In the case of our work on ACARS, the lack of security by default is compounded by those using it not having a good understanding of a realistic threat model. Furthermore, adaptability of the threat model is key. As future research identifies new vulnerabilities and proposes mitigations, the threat model needs to be adjusted according to the latest information. This is an area in which we believe this thesis can directly help through demonstrating the effectiveness of measuring how systems are actually used. Such information can then be fed into threat assessment and inspire mitigations. We discuss this further below.

9.1.2 Adopting a Security Mindset

Closely tied to threat modelling is the adoption of a security mindset. In practice, this means to ensure that security is considered throughout design, implementation and operation of systems used in industry. Whilst aviation's experience in physical security may help in understanding how to defend against attack, cyber security is sufficiently different to require dedicated thinking.

On the other hand, the safety mindset of the industry might prove hard to mesh with. Safety and security are not always complementary; for example, adopting secure communications can lead to communications failure if keys cannot be negotiated. This may lead to an unsafe situation. Addressing this is easier the earlier it is considered in the design process and can be significantly helped by embedding security into industry culture.

9.1.3 Combination of Policy and Technical Measures

It can be tempting to apply single types of measures, i.e. policy or technical, in an effort to secure systems. Instead, a combination of approaches will add defence in depth and is more likely to be successful. This is particularly true for the current situation which the industry finds itself in, having to retroactively secure systems which were not designed with security in mind.

Retroactive technical security measures might not be able to fully secure a system in all situations, meaning that complementary policy measures should be developed. Similarly, policy measures alone are unlikely to secure systems without making them unduly difficult to use or having to remove them from service. Our work in Chapter 8 touches on this by considering the benefit of exposing pilots to cyber attacks through simulated scenarios. Whilst technical measures might help to identify attacks, adapting procedure and providing training to help minimise cockpit impact can reduce the disruption caused whilst the attack is underway.

9.2 A Framework for Assessing Security, Privacy and Safety

Through analysing existing systems, this thesis has identified how real-world avionic usage can be different from intention, which in turn affects security, privacy and sometimes safety. Whilst the studies comprising this work are focussed on specific systems, we believe that the techniques can be used form a more general framework approach. This can help to assess the security and privacy implications of how systems are used in reality.

Measurement of how systems are actually used is important in domains such as aviation where, once a system is deployed, it can remain in place for decades. For example, Chapters 6 and 7, showed ACARS as a system not intended for sensitive data but that is now being used to carry it.

Furthermore, this approach is also useful when human operators play a large role in the way in which systems are used. For example, in Chapter 8 we showed that under certain circumstances pilots will deviate from accepted procedure if they feel that another option offers sound operational benefit such as easier workload management or improvement of safety.

In both cases, a multitude of factors contribute to the eventual impact on security, privacy and safety. In the case of non-commercial usage of ACARS, effort to protect privacy outside the link is a major factor in how the data sent over the

link is interpreted. Namely, whether an aircraft using ACARS is blocked or not changes how we judge whether in-the-clear data transmission from it is acceptable.

Contributing factors can and will evolve over time—some new mechanism to protect data link privacy might emerge, or procedure to minimise the effect of attacks on ILS, TCAS or GPWS could be adopted. At the same time, potential threats will adapt too. This modifies the threat model which triggers the need for reassessment of the security and privacy stance of avionics.

We propose that a measurement approach applied to *how* a system is used in reality complements the existing analyses of how a system *should* be used. As we have shown, the two do not always align, and one way to reveal this is to repeatedly measure operational realities. Various methods can help to achieve this—we used threat model driven data collection and analysis, cryptanalysis and simulation-based user studies. However, other research methods which capture how a system is used are equally valid.

Taking this approach is complemented by deep technical analyses of specific systems. In particular, the measurements can be used to guide further research into these systems. For example, in Chapter 8, our work identifies the impact level of different attacks on avionics whilst accounting for humans in the loop. These findings help to focus future work, such as by researching mitigations for the most vulnerable systems. In the case of our simulator research, this identified TCAS to be of highest concern even though vulnerability of other systems such as GPWS might have appeared more problematic on paper. Some theses already adopt such an approach, such as Strohmeier and McCallie's research into ADS-B security and privacy [270, 271].

9.3 Next Steps

Multiple topics could build on the work in this thesis. As the investigation into avionic data link security is more mature, the next steps can more focussed. In contrast, the work on simulating cyber attacks is more exploratory and so this has many potential directions.

9.3.1 Communications System Security

Although our research focussed on ACARS, the fundamental problem of unsecured data links exists for other current aviation systems, including the non-ACARS uses of SATCOM and VDL2. Our findings suggest that investigating this further could help protect privacy and security for many aircraft operators, leading a number of next steps:

- **Non-standardised security** could provide a fast solution for data link users which need to protect their message content. This should be approached with caution and should not be an excuse to produce weak cryptography, but assessing the feasibility of using other established solutions would be valuable. It is likely that suitability of this approach would depend on the intended usage and available hardware onboard. It needs to allow an 'application layer' security solution on top of the link without modifying the avionics.

- **Measuring commercial data link usage** is a natural complement to this work. We did not consider it in our research as the privacy and security requirements are markedly different to non-commercial aircraft; one of the major differences is that passenger privacy becomes a primary concern. Performing an analysis of this stakeholder group would complement this work and help identify if structured leakage of passenger data was occurring.

- **Message injection and modification**, which so far has only been discussed theoretically. It is unclear what impact an attacker would have by carrying out injection or modification on a link, but since message types such as flight plans or load sheets are transferred, potential for disruption may exist. This work could also generalise to other data links.

9.3.2 Extending Simulated Attacks

A number of opportunities arise from the work in Chapter 8, both in extending the research on the flight deck and looking at other roles.

Attack Proof of Concept and Defence

Since Chapter 8 outlines the attacks, a natural step for this work is to develop proof-of-concept in order to help produce defences. TCAS is one system for which this could be beneficial due to the high impact of the attack. Coupled with the fact that an attack can be carried out at long distances, some Airbus aircraft can now fly TCAS actions automatically, thus limiting the ability of the human-in-the-loop to mitigate said attacks [272].

Improved Cockpit Scenario Simulation

Addressing the limits of the study would provide opportunities to take this work forward. Specifically:

- **Improving simulator fidelity** by matching the simulator to a more realistic and familiar cockpit, with a full crew compliment. This is a significant resource investment as cockpit configurations vary depending on aircraft manufacturer and model. However, running the experiment with full controls and crew would not only make it easier to use, but would allow the participants to act more like they would in a real aircraft. Having two crew members is particularly important in assessing responses to unexpected and dangerous events. This is the main barrier to exploring stronger threat models who are intent on destructive actions, or to taking more precise control response measurements.

- **Investigating attack variants**, particularly with a focus on different threat models. This would include investigating whether attackers could cause the loss of an aircraft, loss of separation or more generally, a reduction in safety. A starting point for this would be to adjust the attacks already investigated; for example, using the GS attack to move the touchdown zone short of threshold.

- **Realistic scenario creation** by including the attacks in longer, more realistic flights, as well as including other traffic and having realistic air traffic control. This would help to remove bias introduced by having participants expecting

and predicting attacks due to short experimental runs, which in turn may affect their response. Having attacks occur as part of a long, otherwise normal flight would also be more closely matched to a realistic attack scenario. These changes will require a significant resource investment as it should be coupled with improvements in fidelity, on top of needing more simulator hours.

- **Development as a training method**, which would adopt a different angle on the work in Chapter 8. Whilst we do provide some analysis on the potential of this approach for training, a next step would be to expand this work extensively and add more rigorous assessment using relevant work in psychology studies on simulation. This could be coupled with trying new procedures. To enable this, the scenarios should be realistic, and the simulator fidelity would need improvement in order to remove barriers to interaction.

Air Traffic Control

Another avenue would be to explore how these attacks might look to ATC operators. This has potential for significant benefit, as it would allow investigations into how attacks affect the wider air traffic system rather than a single aircraft. For example, in the case of the TCAS injection attack, investigating ATC response would help to assess impact of induced RAs on other nearby aircraft. On top of this, simulating the wider system would help to more accurately quantify inconvenience factors such as delay or excess fuel usage, which partially depend on other aircraft.

9.4 Conclusion

Aviation is at a security and privacy turning point. During the time taken to work on this thesis, the industry has begun to accept the cyber security challenges it faces and is now actively engaging with researchers. Airlines, regulators, aircraft, engine and avionics manufacturers are all attempting to better understand the threat of attack and how to handle it.

For privacy, a discussion remains to be had about whether private aircraft movements can be upheld as a concept in aviation. This is becoming pressing,

especially with the mandate for ADS-B meaning that most aircraft will soon be required to broadcast their position and identity. The primary privacy method, blocking, is consistently undermined by both a lack of security in other systems and the new-found ease of collecting messages from aircraft. Furthermore, it relies too heavily on flight tracking company cooperation to be reliable. As sensor networks and flight trackers increase in size and begin to collect different types of messages, blocking as a privacy approach will be effectively rendered void.

Security of current avionics is a larger and arguably more important issue. It is not enough to rely on these systems reaching end-of-life and being decommissioned, rather than addressing their security now. A key factor in managing this risk will be the continuing engagement by industry and regulators with security researchers. This can help to identify where system vulnerabilities lie and how to address them in a timely fashion. Progress is being made in this area but it remains to be seen how advanced it is, with many of the findings being kept private.

One of the biggest ongoing challenges for the aviation industry is factoring security into its safety-focussed mindset. As aircraft and the supporting systems become increasingly connected, building avionics which are both safe and secure will be of paramount importance. Given the crucial role of connectivity in enabling modernisation, it will become difficult to construct safe systems without thinking about security. As the industry is very effective at producing safe systems, it should aim to be of the same 'gold standard' for secure systems too.

A common question arising from our work is whether any of these attacks will actually happen. The nature of the industry means that any security and safety issues become public, from inconveniences such as drunken passengers to catastrophes like the loss of an aircraft. Based on the lack of reports to date, we can infer that attacks similar to those described in this work are still theoretical. If nothing else, this should be a significant motivator to address security problems before an attacker gets there first. After all, yesterday's laser shone into the cockpit is today's drone flying across final approach; tomorrow's new threat may well be an SDR in the wrong hands.

Appendices

A

ACARS Data Collection Frequencies

Table A.1: Collection frequencies for POA, VDL2 and SATCOM ACARS.

Subnetwork	Frequency (MHz)
POA	131.525
	131.725
	131.850
VDL2	136.725
	136.775
	136.875
	136.975
SATCOM	1545.015
	1545.030
	1545.035
	1545.040
	1545.045
	1545.095
	1545.180
	1545.195
	1546.055
	1546.070
	1546.850

B

Opinions on Data Link Security and Privacy Survey

1. What is your role in aviation?

   (a) Captain

   (b) First Officer

   (c) Flight Engineer

   (d) Air Traffic Controller

   (e) Avionics Development/Engineering

   (f) Researcher

   (g) Maintenance

   (h) Prefer not to say

   (i) Other (please specify)

2. What type of organisation do you work for?

   (a) Regional/continental airline

   (b) International airline

   (c) Air Navigation Service Provider (ANSP) (non-governmental)

(d) Governmental/Regulatory body

(e) Avionic Systems Producer

(f) Prefer not to say

(g) Other (please specify)

3. Describe your experience with/knowledge of ACARS briefly (free text)

4. Do you have any experience of ACARS being used to share sensitive or private information? This might include personal data (e.g. names, addresses) or commercially sensitive data.

   (a) Please describe your experience of ACARS being used to share sensitive information, without providing specifics (e.g. company names, data).

5. How suitable would you consider standard ACARS (unencrypted) to be from a safety point-of-view? In other words, to what extent do you think that ACARS is secure enough for safety-related data?

   (a) Very unsuitable

   (b) Somewhat unsuitable

   (c) Neither suitable nor unsuitable

   (d) Somewhat suitable

   (e) Very suitable

   (f) Unsure

   (g) Prefer not to say

6. How suitable would you consider standard ACARS (unencrypted) to be from a privacy point-of-view? For example, for transmitting sensitive data such as names or addresses.

   (a) Very unsuitable

   (b) Somewhat unsuitable

(c) Neither suitable nor unsuitable

(d) Somewhat suitable

(e) Very suitable

(f) Unsure

(g) Prefer not to say

7. ACARS Message Security (AMS, also known as ARINC 823P1) is a standard for providing secure messaging over ACARS. Do you have any experiences with, or knowledge of, it being used in practice?

   (a) Yes (If yes, please comment on the extent of it's use.)

   (b) No

   (c) Prefer not to say

8. Some non-AMS security solutions for ACARS are currently in use, offering varying levels of security. For example, some systems use mono-alphabetic substitution to obfuscate data sent over ACARS. Do you have any experience, or knowledge, of these being used in practice?

   (a) Yes (If yes, please provide further details of what is used, and the extent to which it is used.)

   (b) No

   (c) Prefer not to say

9. Are you aware of the Aircraft Situation Display to Industry (ASDI) programme, run by the US Federal Aviation Administration (FAA)?

   (a) Yes

   (b) No

   (c) Prefer not to say

10. To what extent would you agree with the following statements Statement A "Business aircraft using ASDI blocks regularly undermine their privacy by using ACARS"? We define business aircraft as small/medium sized aircraft (e.g. Gulfstream G650, Embraer Legacy 650) used by individuals or businesses, either owned or chartered. Statement B "State/military aircraft using ASDI blocks/other obscuring mechanisms regularly undermine their privacy by using ACARS"? We define state and military aircraft as any aircraft which operate on an official state basis (and indeed this could be a civilian or military aircraft).

    (a) Statement A

        i. Strongly agree

        ii. Agree

        iii. No opinion

        iv. Disagree

        v. Strongly disagree

    (b) Statement B

        i. Strongly agree

        ii. Agree

        iii. No opinion

        iv. Disagree

        v. Strongly disagree

11. To what extent would you agree with the following statement: "Commercial aircraft use ACARS in a way which could leak sensitive passenger information"? We define commercial aircraft as airliners used to carry passengers, operating in a non-general aviation fashion.

    (a) Strongly agree

    (b) Agree

(c) No opinion

(d) Disagree

(e) Strongly disagree

12. Have you ever had experience with reception of anomalous ACARS messages?

    (a) Yes (If yes, please provide us with some details on this.)

    (b) No

    (c) Prefer not to say

13. If you have any further comments on the use of ACARS for sensitive data or its impact on privacy, enter them below. Please avoid specifics such as company names or sensitive data.

C

ACARS Encryption Frequency Analysis

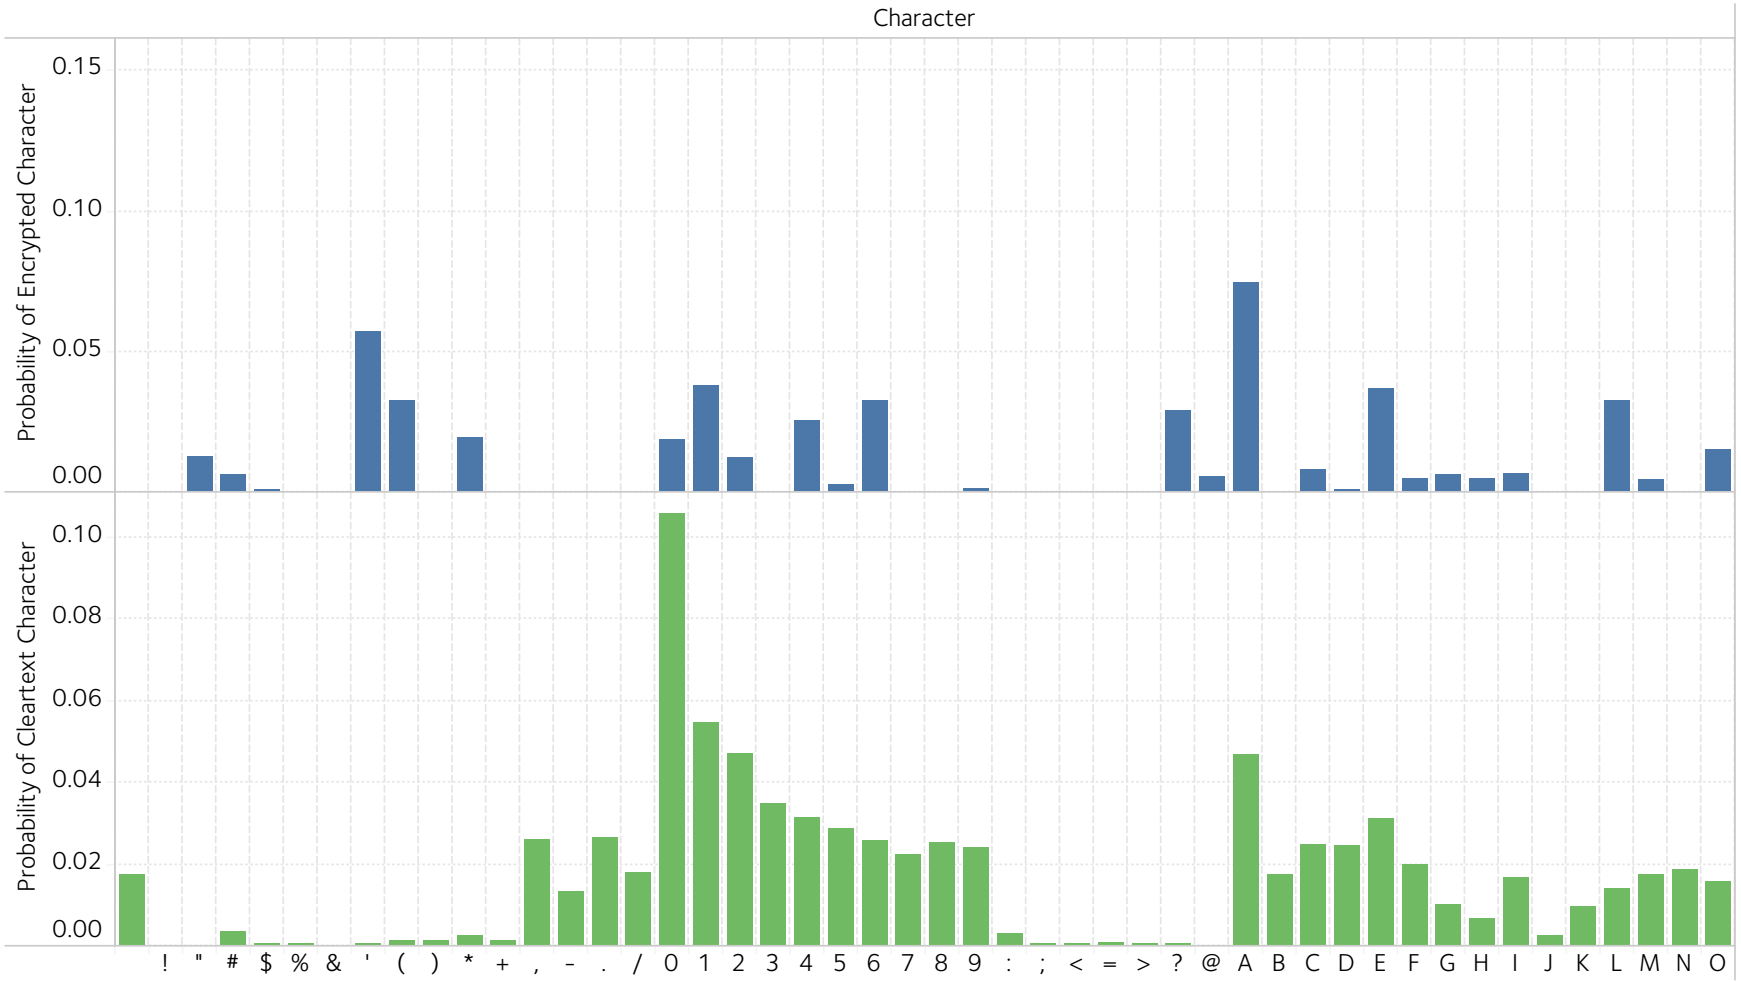

Figure C.1: Frequency analysis for key `01` over first half of the character set. Probability of each character in encrypted messages (top, blue) is compared against probability of the same character appearing in a cleartext message, calculated across all messages (bottom, green).

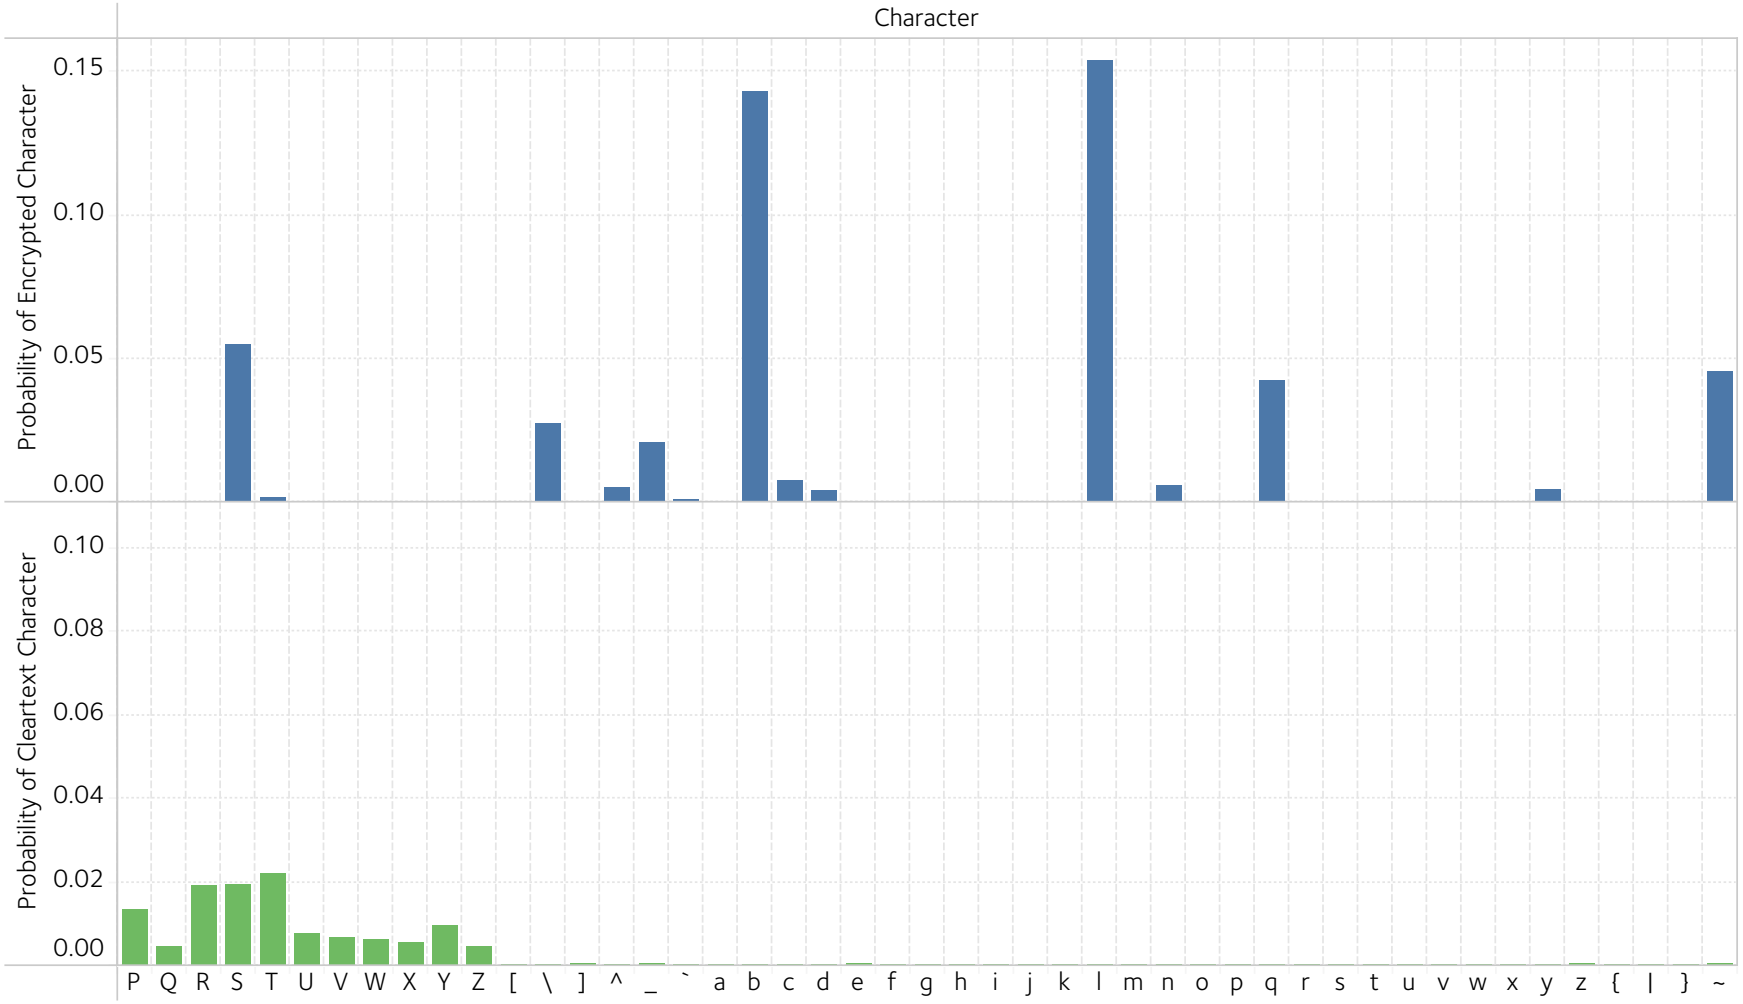

Figure C.2: Frequency analysis for key `01` over second half of the character set. Probability of each character in encrypted messages (top, blue) is compared against probability of the same character appearing in a cleartext message, calculated across all messages (bottom, green).

D

Manufacturer and Model Data for Encrypted Message Usage

Table D.1: Full list of aircraft manufacturers and models using ACARS encryption, with names omitted. Percentages are of all aircraft using this cipher.

Manufacturer			Model		
Name	# Aircraft	%	Name	# Aircraft	%
A	1	0.30	A-1	1	0.30
B	1	0.30	B-1	1	0.30
C	1	0.30	C-1	1	0.30
D	11	3.26%	D-1	2	0.59
			D-2	1	0.30
			D-3	2	0.59
			D-4	1	0.30
			D-5	3	0.89
			D-6	2	0.59
E	20	5.93	E-1	2	0.59
			E-2	3	0.89
			E-3	1	0.30
			E-4	14	4.15
F	296	87.83	F-1	87	25.82
			F-2	154	45.70
			F-3	55	16.32
G	3	0.89	G-1	3	0.89
H	1	0.30	H-1	1	0.30
I	1	0.30	I-1	1	0.30
J	2	0.59	J-1	2	0.59
Total	337	100	—	337	100

E

ACARS Encrypted Messages by Blocked Aircraft

Table E.1: Summary of number of encrypted messages (# in table) by subnetwork and stakeholder, further split into blocked and public, i.e. not blocked. We include commercial as a single seemingly commercial aircraft used this encryption on a handful of messages. Percentages are of all encrypted messages.

	Stakeholder																	
	Business				Military				State				Commercial					
	Blocked		Public		Blocked		Public		Blocked		Public		Blocked		Public		Total	
Subnetwork	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%
POA	1094	29.21	4	0.11	1	0.03	0	0.00	2	0.05	0	0.00	0	0.00	6	0.16	1107	29.56
VDL2	609	16.26	220	5.87	14	0.37	0	0.00	10	0.27	0	0.00	1	0.03	0	0.00	854	22.80
SATCOM	1585	42.32	104	2.78	60	1.60	0	0.00	32	0.85	3	0.08	0	0.00	0	0.00	1784	47.64
Total	3288	87.80	328	8.76	75	2.00	0	0.00	44	1.17	3	0.08	1	0.03	6	0.16	3745	100

F

Flight Simulator Experiment: Procedure

1. Introduction, participant brief on the experiment and signing of consent forms. [10 minutes]
2. Familiarisation flight - 'control' conditions, no attacks, time for pilot to acquaint with controls. [30 minutes]
3. Simulator configuration - set up simulator for ILS attack. [5 minutes]
4. Fly Attack 1 in the simulator - ILS. [30 minutes]
5. Attack 1 debrief. [10 minutes]
6. Simulator configuration - set up simulator for TCAS attack. [5 minutes]
7. Fly Attack 2 in the simulator - TCAS. [30 minutes]
8. Attack 2 debrief. [10 minutes]
9. Simulator configuration - set up simulator for GPWS attack. [5 minutes]
10. Fly Attack 3 in the simulator - GPWS. [30 minutes]
11. Attack 3 debrief. [10 minutes]
12. Final debrief. [10 minutes]

G

Flight Simulator Experiment: Route

Table G.1: Summary of London Heathrow (EGLL) 09R to Birmingham International (EGBB) 33 route used in simulator experiment.

Waypoint/Fix	Maximum Altitude (ft)	Latitude	Longitude
EGLL	83	51.4775	-0.461389
D132B	5100	51.465417	-0.426278
D072J	7000	51.541169	-0.214133
D330T	9000	51.617036	-0.237647
D330W	10000	51.659903	-0.278811
N/A	12000	51.774291	-0.389104
BUZAD	12000	51.942222	-0.552222
DTY	8000	52.180142	-1.113789
D119L	6000	52.2644	-1.374728
N/A	5000	52.303038	-1.495318
HON	3000	52.356678	-1.663725
EGBB	339	52.453889	-1.748056

H

Flight Simulator Experiment: Participant Debrief

Glideslope

1. During this scenario, did the aircraft perform as expected? [Yes/No] In particular, did the ILS approach happen as you would normally expect it to? [Yes/No] Briefly describe the impact of the ILS procedure not occurring as expected, particularly with respect to how this impacted flight and the steps you had to take to account for this.
   (a) Significant impact
   (b) Some impact
   (c) Little impact
   (d) No impact
2. Did you opt to not use it? [Yes/No]
   (a) If so would you use it later? [Yes/No]
3. How confident are you that this was the best decision in the circumstances?
   (a) Very confident

(b) Somewhat confident

(c) Not sure

(d) Unconfident

(e) Very unconfident

4. Do you feel that it put the aircraft in an less safe situation? [Yes/No]

   (a) If so, how?

5. To what extent did this increase your workload?

   (a) Significant increase

   (b) Some increase

   (c) No increase

6. Did this affect your trust in your systems?

   (a) Much distrust

   (b) Some distrust

   (c) No effect

   (d) Some more trust

   (e) Much more trust

7. If this happened in a real aircraft, do you feel you would act in the same way? [Yes/No]

   (a) If not, which different steps would you take? [Yes/No]

TCAS

8. During this scenario, did the aircraft perform as expected? [Yes/No] In particular, did the TCAS system behave as you would normally expect it to? [Yes/No] Briefly describe the impact of TCAS not behaving as expected, particularly with respect to how this impacted flight and the steps you had to take to account for this.

    (a) Significant impact

    (b) Some impact

    (c) Little impact

    (d) No impact

9. Did you turn it off?

    (a) If so would you turn it back on later? [Yes/No]

10. How confident are you that this was the best decision in the circumstances?

    (a) Very confident

    (b) Somewhat confident

    (c) Not sure

    (d) Unconfident

    (e) Very unconfident

11. Do you feel that it put the aircraft in an less safe situation? [Yes/No]

    (a) If so, how?

12. To what extent did this increase your workload?

    (a) Significant increase

    (b) Some increase

(c) No increase

13. Did this affect your trust in your systems?

    (a) Much distrust

    (b) Some distrust

    (c) No effect

    (d) Some more trust

    (e) Much more trust

14. If this happened in a real aircraft, do you feel you would act in the same way? [Yes/No]

    (a) If not, which different steps would you take? [Yes/No]

GPWS

15. During this scenario, did the aircraft perform as expected? [Yes/No]In particular, did the GPWS system behave as you would normally expect it to? [Yes/No] Briefly describe the impact of the GPWS not behaving as expected, particularly with respect to how this impacted flight and the steps you had to take to account for this.

    (a) Significant impact

    (b) Some impact

    (c) Little impact

    (d) No impact

16. Did you turn it off? [Yes/No]

    (a) If so would you turn it back on later? [Yes/No]

17. How confident are you that this was the best decision in the circumstances?

(a) Very confident

(b) Somewhat confident

(c) Not sure

(d) Not confident

(e) Significantly not confident

18. Do you feel that it put the aircraft in an less safe situation? [Yes/No]

    (a) If so, how?

19. To what extent did this increase your workload?

    (a) Significant increase

    (b) Some increase

    (c) No increase

20. Did this affect your trust in your systems?

    (a) Much distrust

    (b) Some distrust

    (c) No effect

    (d) Some more trust

    (e) Much more trust

21. If this happened in a real aircraft, do you feel you would act in the same way? [Yes/No]

    (a) If not, which different steps would you take? [Yes/No]

Final debrief

22. Did you find the scenarios to be a useful exercise? [Yes/No]

23. To what extent did you feel limited by the simulator?

    (a) Heavily limited

    (b) Somewhat limited

    (c) Not limited

24. Have you encountered any of the scenarios in the wild? [Yes/No] If so, provide detail.

25. Do you feel that this could be a useful training tool for pilots? [Yes/No]

I

Flight Simulator Experiment: Interview Debrief Responses

Table I.1: Summary of participant interview responses for attack scenarios. Scale points are normalized so that 1 represents the most 'positive' point, and the lowest value represents the most 'negative', e.g. for impact, 1 is the 'significant impact' response. Dash indicates where no scale value existed, and representative scale point is taken as scale response at the rounded mean, e.g. for impact, 1.4 will be 'significant impact'.

		Scale Points											
Attack	Question	1	1.5	2	2.5	3	3.5	4	4.5	5	Mean	Representative Scale Point	Std. Dev
GS	Impact	10	3	10	2	4	0	1	-	-	1.85	Some impact	0.787
	Confidence	21	1	8	0	0	0	0	0	0	1.28	Very confident	0.441
	Workload	6	1	22	1	0	-	-	-	-	1.80	Some increase	0.420
	Trust	5	0	18	0	7	0	0	0	0	2.07	Some distrust	0.629
TCAS	Impact	19	3	5	2	1	0	0	-	-	1.38	Significant impact	0.573
	Confidence	12	4	11	0	3	0	0	0	0	1.63	Somewhat confident	0.639
	Workload	16	4	9	1	0	-	-	-	-	1.42	Significant increase	0.484
	Trust	19	2	8	0	1	0	0	0	0	1.36	Much distrust	0.531
GPWS	Impact	8	2	13	0	3	2	2	-	-	2.03	Some impact	0.894
	Confidence	24	0	5	0	1	0	0	0	0	1.23	Very confident	0.496
	Workload	13	2	11	3	1	-	-	-	-	1.62	Some increase	0.601
	Trust	10	3	16	0	1	0	0	0	0	1.65	Some distrust	0.519

References

[1] Harro Ranter. *ASN data shows 2017 was safest year in aviation history*. Online. Accessed on 2018-11-26. Dec. 2017. URL: https://news.aviation-safety.net/2017/12/30/preliminary-asn-data-show-2017-safest-year-aviation-history/.

[2] Ellis F. Hitt. "Digital Avionics Handbook". In: ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 5, 5.1–5.3.

[3] European Commission and Eurocontrol. *Discover SESAR*. Online. Accessed on 2018-11-26. 2018. URL: https://www.sesarju.eu/discover-sesar.

[4] Federal Aviation Administration. *Modernization of U.S. Airspace*. Online. Accessed on 2018-11-26. Nov. 2018. URL: https://www.faa.gov/nextgen/.

[5] Internation Civil Aviation Organisation. *Global Air Navigation Plan 2016-2030*. Tech. rep. Doc 9750-AN/963. Internation Civil Aviation Organisation, 2016. URL: https://www.icao.int/airnavigation/Documents/GANP-2016-interactive.pdf.

[6] SESAR JU. "European ATM Master Plan". In: 2015. Chap. 3, p. 22. URL: https://www.sesarju.eu/masterplan.

[7] BBC. *Pilot laser attacks 'dangerously high'*. Online. Accessed on 2018-12-02. Feb. 2017. URL: https://www.bbc.co.uk/news/uk-scotland-glasgow-west-39107853.

[8] Kristine Phillips. Online. Accessed on 2018-12-02. Jan. 2018. URL: https://www.washingtonpost.com/news/dr-gridlock/wp/2018/01/20/a-woman-called-the-serial-stowaway-sneaked-past-airport-security-again-and-flew-to-london/?utm_term=.d32e7526af86.

[9] BBC. *Heathrow plane stowaway makes 'significant recovery'*. Online. Accessed on 2018-12-02. Sept. 2015. URL: https://www.bbc.co.uk/news/uk-england-london-34127735.

[10] CNN Library. *Terrorism and War-Related Airplane Crashes Fast Facts*. Online. Accessed 2018-12-02. May 2018. URL: https://edition.cnn.com/2016/03/24/world/terrorism-and-war-related-airplane-crashes-fast-facts/index.html.

[11] BBC. *Cathay Pacific data hack hits 9.4 million passengers*. Online. Accessed on 2018-12-02. Oct. 2018. URL: https://www.bbc.co.uk/news/business-45974020.

[12] James Sillars. *BA hacked: 380,000 card payments 'compromised' in breach*. Online. Accessed on 2018-12-02. Sept. 2018. URL: https://news.sky.com/story/ba-calls-in-police-over-customer-data-theft-from-website-11491980.

[13] Matt Thurber. *The aircraft certification process*. Online. Accessed on 2016-12-16. Dec. 2006. URL: https://www.ainonline.com/aviation-news/aviation-international-news/2006-12-18/aircraft-certification-process.

[14] Steve Markgraf, Dimitri Stolnikov, and Hoernchen. *RTL-SDR*. Online. Accessed on 2018-10-01. Oct. 2018. URL: https://www.rtl-sdr.com/about-rtl-sdr/.

[15] Andrei Costin and Aurélien Francillon. "Ghost in the Air(Traffic): On insecurity of ADS-B protocol and practical attacks on ADS-B devices". In: *Black Hat USA*. July 2012, pp. 1–10.

[16] Brad 'RenderMan' Haines. *Hacker + Airplanes = No Good Can Come Of This*. Presentation. Las Vegas, 2013. URL: https://media.defcon.org/DEFCON20/DEFCON20slides/DEFCON20HackingConferencePresentationByRenderMan-HackerandAirplanesNoGoodCanComeOfThis-Slides.m4v.

[17] Hugo Teso. "Aicraft hacking: Practical Aero Series". In: *4th Hack in the Box Security Conference in Europe*. Amsterdam, Apr. 2013.

[18] Andy Greenberg. *Researcher Says He's Found Hackable Flaws In Airplanes' Navigation Systems (Update: The FAA Disagrees)*. Online. Accessed on 2018-12-16. Apr. 2013. URL: http://www.forbes.com/sites/andygreenberg/2013/04/10/researcher-says-hes-found-hackable-flaws-in-airplanes-navigation-systems/.

[19] Kellyn Wagner Ramsdell. *Boeing 757 Hacked in DHS Test*. Online. Accessed on 2018-11-26. Feb. 2018. URL: https://www.ainonline.com/aviation-news/air-transport/2018-02-01/boeing-757-hacked-dhs-test.

[20] Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. "Experimental Analysis of Attacks on Next Generation Air Traffic Communication". In: *Lecture Notes in Computer Science* 7954 LNCS (2013), pp. 253–271. URL: http://link.springer.com/chapter/10.1007/978-3-642-38980-1_16.

[21] Martin Strohmeier, Vincent Lenders, and Ivan Martinovic. "On the Security of the Automatic Dependent Surveillance-Broadcast Protocol". In: *IEEE Communications Surveys and Tutorials* 17.2 (2014), pp. 1066–1087. arXiv: 1307.3664. URL: http://arxiv.org/abs/1307.3664.

[22] 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 (2011), pp. 78–87. URL: http://linkinghub.elsevier.com/retrieve/pii/S1874548211000229.

[23] Flightradar24 AB. *Flightradar24*. Online. Accessed on 2018-12-16. 2018. URL: https://www.flightradar24.com.

[24] FlightAware. *FlightAware*. Online. Accessed on 2018-12-16. 2018. URL: https://www.flightaware.com/.

[25] The OpenSky Network. *OpenSky Network*. Online. Accessed on 2018-11-26. 2018. URL: https://opensky-network.org/.

[26] Woodrow Bellamy III. *What is the Answer to Business Aviation's ADS-B Privacy Concern?* Online. Accessed 2018-10-29. Oct. 2018. URL: http://interactive.aviationtoday.com/avionicsmagazine/october-november-2018/what-is-the-answer-to-business-aviations-ads-b-privacy-concern/.

[27] Martin Strohmeier et al. "The Real First Class? Inferring Confidential Corporate Mergers and Government Relations from Air Traffic Communication". In: *IEEE European Symposium on Security and Privacy (EuroS&P) 2018*. IEEE. Apr. 2018.

[28] Peter Aldhous. *We Trained A Computer To Search For Hidden Spy Planes. This Is What It Found.* Online. Accessed on 2018-11-27. Aug. 2017. URL: https://www.buzzfeednews.com/article/peteraldhous/hidden-spy-planes.

[29] Julien Pilet and François Pilet. *GVA Dictator Alert*. Online. Accessed on 2018-12-16. 2018. URL: https://twitter.com/GVA_Watcher.

[30] Matthew Smith et al. "Undermining Privacy in the Aircraft Communications Addressing and Reporting System (ACARS)". In: *18th Privacy Enhancing Technologies Symposium (PETS 2018)*. July 2018.

[31] Matthew Smith et al. "Economy Class Crypto: Exploring Weak Cipher Usage in Avionic Communications via ACARS". In: *21st International Conference on Financial Cryptography and Data Security*. Malta, 2017.

[32] Matthew Smith et al. "Safety vs. Security: Attacking Avionic Systems with Humans in the Loop". In: *CoRR* abs/1905.08039 (2019). arXiv: 1905.08039. URL: http://arxiv.org/abs/1905.08039.

[33] Commercial Aviation Safety Team. *Phase of Flight*. Tech. rep. 1.3. International Civil Aviation Organisaton, Apr. 2013. URL: http://www.intlaviationstandards.org/Documents/PhaseofFlightDefinitions.pdf.

[34] Dorothy Saul-Pooley. "Air Law & Meteorology". In: ed. by Dorothy Saul-Pooley et al. Controlled airspace def. Pooley's, 2017. Chap. 5, pp. 76–81.

[35] Dorothy Saul-Pooley. "Air Law & Meteorology". In: ed. by Dorothy Saul-Pooley et al. FIR definition. Pooley's, 2017. Chap. 5, pp. 75–76.

[36] Dorothy Saul-Pooley. "Air Law & Meteorology". In: ed. by Dorothy Saul-Pooley et al. uncontrolled airspace defintion. Pooley's, 2017. Chap. 5, pp. 81–86.

[37] Dorothy Saul-Pooley. "Air Law & Meteorology". In: ed. by Dorothy Saul-Pooley et al. VFR defintion. Pooley's, 2017. Chap. 7, pp. 117–119.

[38] Dorothy Saul-Pooley. "Air Law & Meteorology". In: ed. by Dorothy Saul-Pooley et al. IFR defintion. Pooley's, 2017. Chap. 8, p. 123.

[39] LHR Airports Limited. *Monthly traffic statistics up to October 2018, excluding Gatwick, Stanstead, Edinburgh, Naples, Aberdeen, Glasgow and Southampton.* Online. Accessed on 2018-11-26. Oct. 2018. URL: https://www.heathrow.com/company/investor-centre/reporting/traffic-statistics.

[40] International Civil Aviation Organization. *Guidance Material on Advice to Military Authorities Regarding ADS-B Data Sharing*. Online. Accessed on 2018-12-16. Sept. 2012. URL: https://www.icao.int/APAC/Documents/edocs/advice%20to%20%20military%20%20authorities.pdf.

[41] John Bourn. *Royal Travel by Air and Rail*. Tech. rep. Accessed on 2017-01-18. London: United Kingdom National Audit Office, 2001. URL: https://www.nao.org.uk/wp-content/uploads/2001/06/010225.pdf.

[42] Roy T Oishi and Ann Heinke. "Digital Avionics Handbook". In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, 2.4–2.7.

[43] European Commission. *Commission Implementing Regulation (EU) No 1079/2012 of 16 November 2012 laying down requirements for voice channels spacing for the single European sky Text with EEA relevance*. Nov. 2012. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32012R1079.

[44] Rockwell Collins. *VHF & HF Air/Ground Voice Services*. Online. Accessed on 2018-11-16. 2018. URL: https://www.arincdirect.com/what-we-do/flight-deck-communications/voice-services/.

[45] International Civil Aviation Organization. "Annex 10 to the Convention on International Civil Aviation—Aeronuatical Telecommunications". In: vol. 1. 2006. Chap. 5, p. 5.19.

[46] H Hering, M Hagmüller, and G Kubin. "Safety and Security increase for Air Traffic Management Through Unnoticeable Watermark Aircraft Identification Tag Transmitted with the VHF Coice Communication". In: *Digital Avionics Systems Conference, 2003. DASC '03. The 22nd* 1 (2003), 4.E.2–41–10 vol.1.

[47] M. Hagmüller et al. "Speech watermarking for air traffic control". In: *2004 12th European Signal Processing Conference*. Sept. 2004, pp. 1653–1656.

[48] R. Fantacci et al. "A secure radio communication system based on an efficient speech watermarking approach". In: *Security and Communication Networks* 2.4 (2008), pp. 305–314. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/sec.70.

[49] J. Prinz, M. Sajatovic, and B. Haindl. "S/sup 2/EV- safety and security enhanced ATC voice system". In: *2005 IEEE Aerospace Conference*. Mar. 2005, pp. 1924–1930.

[50] T. H. Stelkens-Kobsch et al. "Towards a more secure ATC voice communications system". In: *2015 IEEE/AIAA 34th Digital Avionics Systems Conference (DASC)*. Sept. 2015, pp. 4C1–1–4C1–9.

[51] Michael Finke and Tim H. Stelkens-Kobsch. "A practical example for validation of ATM security prototypes". In: *CEAS Aeronautical Journal* 9.1 (Mar. 2018), pp. 157–170. URL: https://doi.org/10.1007/s13272-017-0275-y.

[52] Michael Neffe et al. "Speaker Segmentation for Air Traffic Control". In: *Speaker Classification II: Selected Projects*. Ed. by Christian Müller. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 177–191. URL: https://doi.org/10.1007/978-3-540-74122-0_15.

[53] Dustin Hoffman and Semon Rezchikov. "Busting the BARR: Tracking "Untrackable" Private Aircraft for Fun & Profit". In: *DEF CON 20*. Las Vegas, 2012.

[54] Michael Zenkovich. *Data Link Communications*. Advisory Circular 90-117. Federal Aviation Administration, Mar. 2017.

[55] Roy T Oishi and Ann Heinke. "Digital Avionics Handbook". In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, 2.7–2.13.

[56] SKYbrary. *North Atlantic Operations - ATC Clearance*. Online. Accessed on 2018-11-16. Dec. 2017. URL: https://www.skybrary.aero/index.php/North_Atlantic_Operations_-_ATC_Clearance.

[57] EUROCONTROL. *Don't text while driving. Text while Flying! Why use Controller-Pilot Data Link Communications?* June 2018. URL: https://www.eurocontrol.int/sites/default/files/publication/files/factsheet-cpdlc.pdf.

[58] EUROCONTROL. *CPDLC General Safety Considerations*. Online. Accessed on 2018-10-05. Dec. 2017. URL: https://www.skybrary.aero/index.php/CPDLC_General_Safety_Considerations.

[59] Airways New Zealand. *FANS1/A Problem Reporting*. Online. Accessed on 2018-11-30. 2018. URL: http://www.fans-cra.com/.

[60] Andrei Gurtov, Tatiana Polishchuk, and Max Wernberg. "Controller–Pilot Data Link Communication Security". In: *Sensors* 18.5 (2018). URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5982923/.

[61] O. Osechas et al. "Addressing vulnerabilities of the CNS infrastructure to targeted radio interference". In: *IEEE Aerospace and Electronic Systems Magazine* 32.11 (Nov. 2017), pp. 34–42.

[62] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. NDB ref. Pooley's, 2017. Chap. 11, pp. 205–210.

[63] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. Pooley's, 2017. Chap. 11, pp. 210–217.

[64] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. VOR. Pooley's, 2017. Chap. 14, pp. 277–285.

[65] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. DME. Pooley's, 2017. Chap. 10, pp. 187–197.

[66] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. Glideslope. Pooley's, 2017. Chap. 15, pp. 325–329.

[67] NATS Aeronautical Information Service. *London City Aerodrome – Textual Data*. Online. Accessed on 2018-10-02. Dec. 2017. URL: http://www.ead.eurocontrol.int/eadbasic/pamslight-188C68846080EFD633326C47E18B6948/7FE5QZZF3FXUS/EN/AIP/AD/EG_AD_2_EGLC_en_2018-09-13.pdf.

[68] British Airways. *Capital views: approach into London City Airport*. Online. Accessed on 2018-12-14. Aug. 2016. URL: https://youtu.be/1S8bjpHLzog.

[69] Ofcom. *UK Interface Requirement 2055—Ground based instrument landing system (ILS) glide path transmitter radio of the aeronautical radionavigation service.* Tech. rep. 2018.

[70] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. Localiser. Pooley's, 2017. Chap. 15, pp. 325–321.

[71] Ofcom. *UK Interface Requirement 2056—Ground based instrument landing system (ILS) localiser radio equipment at aeronautical stations of the aeronautical radionavigation service.* Tech. rep. 2018.

[72] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. Pooley's, 2017. Chap. 17, pp. 359–363.

[73] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. Pooley's, 2017. Chap. 17, pp. 374–386.

[74] Alan Grant et al. "GPS Jamming and the Impact on Maritime Navigation". In: *Journal of Navigation* 62.2 (2009), 173–187.

[75] Hui Hu and Na Wei. "A study of GPS jamming and anti-jamming". In: *2009 2nd International Conference on Power Electronics and Intelligent Transportation System (PEITS).* Vol. 1. Dec. 2009, pp. 388–391.

[76] Nils Ole Tippenhauer et al. "On the Requirements for Successful GPS Spoofing Attacks". In: *Proceedings of the 18th ACM Conference on Computer and Communications Security.* CCS '11. Chicago, Illinois, USA: ACM, Oct. 2011, pp. 75–86.

[77] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. PSR. Pooley's, 2017. Chap. 9, pp. 169–175.

[78] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. SSR. Pooley's, 2017. Chap. 9, pp. 175–177.

[79] International Civil Aviation Organization. "Manual on the Secondary Surveillance Radar (SSR) Systems". In: Third. Doc 9684 AN/951. 2004. Chap. 1, 1.1–1.2.

[80] EUROCONTROL. "Transponders in aviation". In: *NetAlert* 19 (May 2014), pp. 1–3. URL: `https://www.eurocontrol.int/sites/default/files/publication/files/NetAlert-19.pdf`.

[81] Martin Strohmeier et al. "Realities and Challenges of NextGen Air Traffic Management: The Case of ADS-B". In: *IEEE Communications Magazine* 52.5 (2014), pp. 111–118.

[82] Matthias Schäfer et al. "Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research". In: *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks.* IPSN '14. Berlin, Germany: IEEE Press, Apr. 2014, pp. 83–94.

[83] International Civil Aviation Organization. "Manual on the Secondary Surveillance Radar (SSR) Systems". In: Third. Doc 9684 AN/951. 2004. Chap. 10, 10.1–10.3.

[84] Matthias Schäfer et al. “OpenSky’s Report 2016: Facts, Figures and Trends in Wireless ATC Communication Systems”. In: *35th Digital Avionics Systems Conference - Proceedings*. IEEE/AIAA, 2016.

[85] M. Schäfer et al. “OpenSky Report 2017: Mode S and ADS-B Usage of Military and Other State Aircraft”. In: *2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*. Sept. 2017, pp. 1–10.

[86] M. Schäfer et al. “OpenSky Report 2018: Assessing the Integrity of Crowdsourced Mode S and ADS-B Data”. In: *2018 IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*. 2018.

[87] Joel M. Wichgers. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 23, 23.1–23.2.

[88] Federal Aviation Administraton. *Automatic Dependent Surveillance-Broadcast (ADS-B) Out equipment performance requirements.* Federal Regulation 14 CFR 91.225. US Department for Transport, May 2010.

[89] European Commission. *Commission Implementing Regulation (EU) 2017/386 of 6 March 2017 amending Implementing Regulation (EU) No 1207/2011 laying down requirements for the performance and the interoperability of surveillance for the single European sky*. Regulation 2017/386. European Parliament, Mar. 2017.

[90] Joel M. Wichgers. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 23, 23.4–23.6.

[91] Joel M. Wichgers. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 23, 23.7–23.9.

[92] Joel M. Wichgers. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 23, 23.9–23.11.

[93] Emily Cook. “ADS-B, friend or foe: ADS-B message authentication for NextGen aircraft”. In: *Proceedings - 2015 IEEE 17th International Conference on High Performance Computing and Communications, 2015 IEEE 7th International Symposium on Cyberspace Safety and Security and 2015 IEEE 12th International Conference on Embedded Software and Systems* (2015), pp. 1256–1261.

[94] Omar Yeste-Ojeda and Rene Landry. “ADS-B Authentication Compliant with Mode-S Extended Squitter Using PSK Modulation”. In: *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*. 2015, pp. 1773–1778.

[95] Richard C Agbeyibor. “Secure ADS-B: Towards Airborne Communications Security in the Federal Aviation Administration’s Next Generation Air Transportation System”. PhD thesis. US Air Force Institute of Technology, 2014. URL: `http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA600893`.

[96] Ning Xu, Rick Cassell, and Carl Evers. “Performance assessment of multilateration systems - a solution to NextGen surveillance”. In: *Integrated Communications Navigation and Surveillance Conference (ICNS)*. IEEE/AIAA, 2010, pp. 2–9.

[97] Federal Aviation Administration. *Wide Area Multilateration (WAM) Project*. Accessed on 2018-12-12. 2018. URL: https://www.faa.gov/nextgen/programs/adsb/atc/wam/.

[98] 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 '15)*. ACM, Apr. 2015, pp. 49–60.

[99] Alexandre Barreto and Paulo Costa. "Detecting Malicious ADS-B Transmitters Using a Low-Bandwidth Sensor Network". In: *18th International Conference on Information Fusion*. 2015, pp. 1696–1701.

[100] Martin Strohmeier et al. "Crowdsourcing Security for Wireless Air Traffic Communications". In: *Cyber Conflict (CYCON), 2017 9th International Conference on*. IEEE. May 2017.

[101] Martin Strohmeier and Ivan Martinovic. "On Passive Data Link Layer Fingerprinting of Aircraft Transponders". In: *Cps-Spc* (2015).

[102] International Civil Aviation Organization. "Global Operational Data Link Document (GOLD)". In: 2nd ed. Apr. 2013. Chap. 2, 2–46–65.

[103] Anita Pratap and Jamie McIntyre. *Officials say hundreds feared killed in airline collision over India*. Online. Accessed 2018-10-08. Nov. 1996. URL: https://web.archive.org/web/20020911221032/http://www.cnn.com/WORLD/9611/12/india.air.crash/index.html.

[104] International Civil Aviation Organization. "Annex 10 to the Convention on International Civil Aviation—Surveillance and Collision Avoidance System". In: 4th ed. Vol. 4. July 2007. Chap. 4, 4.1–4.47.

[105] Federal Aviation Administration. "Introduction to TCAS II Version 7.1". In: U.S. Department of Transport, 2011. Chap. 1, 5–10.

[106] Steve Henely. "Digital Avionics Handbook". In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 22, 22.1.

[107] UK Civil Aviation Authority. *Ground Proximity Warning Systems*. Accessed on 2018-12-16. 1976. URL: https://publicapps.caa.co.uk/docs/33/CASPEC14.PDF.

[108] David Jensen. *EGPWS: Look What It Can Do Now*. Online. Accessed on 2018-12-01. Nov. 2000. URL: https://www.aviationtoday.com/2000/11/01/egpws-look-what-it-can-do-now/.

[109] Department for Transport. *Aviation Cyber Security Strategy*. Tech. rep. Her Majesty's Government, 2018. URL: https://www.gov.uk/government/publications/aviation-cyber-security-strategy.

[110] Ettus Research. *About Us*. Online. Accessed on 2018-09-26. 2018. URL: https://www.ettus.com/about.

[111] The GNU Radio Foundation, Inc. *GNU Radio*. Online. Accessed on 2018-10-01. Oct. 2018. URL: https://www.gnuradio.org/.

[112] FlightAware. *FlightAware Pro Stick and Pro Stick Plus - High Performance USB SDR ADS-B Receivers*. Online. Accessed on 2018-10-01. Oct. 2018. URL: https://uk.flightaware.com/adsb/prostick/.

[113] Great Scott Gadgets. *HackRF One*. Online. Accessed on 2018-10-01. Oct. 2018. URL: https://greatscottgadgets.com/hackrf/.

[114] KjM and Rene Hesse. *acarsd - ACARS Decoder for Linux & Windows*. Online. Accessed on 2018-12-07. 2007. URL: https://www.acarsd.org/.

[115] Anthony Tyler. *The Impact of September 11 2001 on Aviation*. Tech. rep. International Air Transport Association, 2011. URL: https://www.iata.org/pressroom/documents/impact-9-11-aviation.pdf.

[116] Harumi Ito and Darin Lee. "Assessing the impact of the September 11 terrorist attacks on U.S. airline demand". In: *Journal of Economics and Business* 57.1 (2005), pp. 75 –95.

[117] Bruce Schneier. *More on Chris Roberts and Avionics Security*. 2015. URL: https://www.schneier.com/blog/archives/2015/05/more_on_chris_r.html (visited on 07/31/2015).

[118] Kim Zetter. *Is it possible for passengers to hack commercial aircraft?* Online. Accessed on 2018-10-01. May 2015. URL: https://www.wired.com/2015/05/possible-passengers-hack-commercial-aircraft/.

[119] David Gloven and David Voreacos. *Dream Insider Informant Led FBI From Galleon to SAC*. Online. Accessed on 2018-12-05. Dec. 2012. URL: http://www.bloomberg.com/news/articles/2012-12-03/dream-insider-informant-led-fbi-from-galleon-to-sac.

[120] Damien Gayle and Amy Gibbons. *Activists accused of blocking Stansted flight go on trial over terror charge*. Online. Accessed on 2018-11-27. Oct. 2018. URL: https://www.theguardian.com/uk-news/2018/oct/02/trial-begins-for-anti-deportation-activists-accused-of-blocking-stansted-flight.

[121] Jeremy C. Davis. *Sweden's airspace shut down by Russian APT, not a solar storm*. Accessed on 2018-08-10. 2016. URL: https://www.scmagazine.com/home/security-news/swedens-airspace-shut-down-by-russian-apt-not-a-solar-storm/.

[122] Roy T Oishi and Ann Heinke. "Digital Avionics Handbook". In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, 2.1–2.3.

[123] ARINC, Inc. *ICAO South American Region Data Link Applications Workshop*. Online. Accessed 2018-10-11. Sept. 2012. URL: https://www.icao.int/SAM/Documents/DATALINK11/Sesion02%2008%20ARINC%20AirGroundDataLinkATNAppsServices.pdf.

[124] ARINC. *Air/Ground Character-Oriented Protocol Specification*. Tech. rep. 618-7. 2013.

[125] Roy T Oishi and Ann Heinke. "Digital Avionics Handbook". In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, p. 2.9.

[126] Roy T Oishi and Ann Heinke. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, 2.10–2.11.

[127] Roy T Oishi and Ann Heinke. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, 2.9–2.10.

[128] Inmarsat plc. *SwiftBroadband-Safety*. Online. Accessed on 2018-10-11. Oct. 2018. URL: https://www.inmarsat.com/aviation/complete-aviation-connectivity/classic-aero/.

[129] Inmarsat plc. *Classic Aero*. Online. Accessed on 2018-10-11. Oct. 2018. URL: https://www.inmarsat.com/aviation/complete-aviation-connectivity/classic-aero/.

[130] Inmarsat plc. *Our coverage*. Online. Accessed on 2018-10-11. Oct. 2018. URL: https://www.inmarsat.com/about-us/our-satellites/our-coverage/.

[131] Inmarsat plc. *Alphasat and the I-4s*. Online. Accessed on 2018-10-11. Oct. 2018. URL: https://www.inmarsat.com/about-us/about-usour-satellites/inmarsat-4/.

[132] Roy T Oishi and Ann Heinke. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, p. 2.11.

[133] Carolyn Bray. *HFDL and Polar Flights – CPWG Meeting*. Online. Accessed 2018-10-11. May 2017. URL: https://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/systemops/ato_intl/documents/cross_polar/CPWG23/CPWG23_Brf_HFDL_and_Polar_Flights.pdf.

[134] ARINC. *Air/Ground Character-Oriented Protocol Specification*. Standard 753. Feb. 2001.

[135] Roy T Oishi and Ann Heinke. “Digital Avionics Handbook”. In: *Digital Avionics Handbook*. Ed. by Cary R. Spitzer, Uma Ferrell, and Thomas Ferrell. Third. CRC Press, 2015. Chap. 2, p. 2.10.

[136] ARINC. *Datalink Ground System Standard and Interface Specification*. Technical Standard 620-8. 2014.

[137] Woodrow Bellamy III. *TCPDLC: Transitioning to Controller Pilot Data Link Communications*. Online. Accessed on 2018-12-08. Dec. 2014. URL: https://www.aviationtoday.com/2014/12/01/tcpdlc-transitioning-to-controller-pilot-data-link-communications/.

[138] Aeronautical Radio Inc. (ARINC). *DataLink Security, Part 1 - ACARS Message Security*. Technical Standard 823P1. 2007.

[139] Aloke Roy. “Secure aircraft communications addressing and reporting system (ACARS)”. In: *20th Digital Avionics Systems Conference* 2 (2001), 7A2/1–7A2/11.

[140] Michael Olive. *ACARS Message Security (AMS) as a Vehicle for Validation of ICAO Doc. 9880 Part IV-B Security Requirements*. Online. Accessed 2018-10-16. June 2009. URL: https://www.icao.int/safety/acp/ACPWGF/ACP-WG-M-14/ACP-WGM-IP07%20-%20AMS%20for%209880%20Security%20Validation_20090604%20(HON-Olive).pdf.

[141] R Housley and L Zieglar. *Reclassification of Suite B Documents to Historic Status*. Online. Accessed on 2018-12-02. Feb. 2018. URL: https://tools.ietf.org/html/draft-housley-suite-b-to-historic-04.

[142] B. Blanchet. "Symbolic and Computational Mechanized Verification of the ARINC823 Avionic Protocols". In: *IEEE 30th Computer Security Foundations Symposium (CSF)*. Aug. 2017, pp. 68–82.

[143] Paul E. Storck. "Benefits of Commercial Data Link Security". In: *Integrated Communications, Navigation and Surveillance Conference, ICNS*. Herndon: IEEE, 2013.

[144] Aloke Roy. "Security Strategy for US Air Force to Use Commercial Data Link". In: *19th Digital Avionics Systems Converence*. IEEE, 2000, pp. 1–8.

[145] Curtis Risley, James McMath, and Brian Payne. "Experimental encryption of Aircraft Communications Addressing and Reporting System (ACARS) Aeronautical Operational Control (AOC) messages". In: *20th Digital Avionic Systems Conference*. Daytona Beach, 2001, pp. 1–8.

[146] Charlotte Adams. "Securing ACARS: Data Link in the Post 9/11 Environment". In: *Avionics Magazine* (June 2006), pp. 24–26.

[147] Meng Yue and Xiaofeng Wu. "The approach of ACARS data encryption and authentication". In: *Proceedings - 2010 International Conference on Computational Intelligence and Security, CIS 2010* (2010), pp. 556–560.

[148] Martin Strohmeier et al. "On Perception and Reality in Wireless Air Traffic Communication Security". In: *IEEE Transactions on Intelligent Transportation Systems* 18.6 (June 2017), pp. 1338–1357. URL: http://ieeexplore.ieee.org/document/7589091/.

[149] Sirio Antenne. *GP- LB Series*. Online. Accessed on 2018-12-26. 2018. URL: http://www.sirioantenne.it/en/products/vhf/gp-lb-series.

[150] Thierry Leconte. *ACARSDec ACARS Decoder*. Online. Accessed on 2018-12-18. 2015. URL: https://github.com/TLeconte/acarsdec.

[151] Tomasz Lemiech. *dumpvdl2*. Accessed on 2016-12-18. Aug. 2018. URL: https://github.com/szpajder/dumpvdl2.

[152] Jonathan Olds. *JAERO*. Accessed on 2018-12-16. 2017. URL: https://github.com/jontio/JAERO.

[153] Matthias Schäfer et al. "Bringing Up OpenSky: A Large-scale ADS-B Sensor Network for Research". In: *IPSN 2014 - Proceedings of the 13th International Symposium on Information Processing in Sensor Networks* (2014), pp. 83–94.

[154] FlightAware. *Global Flight Tracking*. Online. Accessed on 2018-12-04. 2018. URL: http://uk.flightaware.com/commercial/global.

[155] Junzi Sun. *World Aircraft Database 2018*. Online. Accessed on 2018-10-29. 2018. URL: https://junzisun.com/adb/data.

[156] Ralph D. Kloth. *Airframes.org*. Accessed on 2018-12-14. 2018. URL: http://www.airframes.org/.

[157] Federal Aviation Administration. *Aircraft Registry - Releasable Aircraft Database Download*. Accessed on 2017-02-19. 2018. URL: https://www.faa.gov/licenses_certificates/aircraft_certification/aircraft_registry/releasable_aircraft_download/.

[158] Darren R. Crocker. *AvDelphi*. Online. Accessed on 2018-12-01. 2018. URL: https://www.avdelphi.com/.

[159] John Croft. *NBAA: The old BARR is back*. Online. Accessed on 2018-10-29. Nov. 2011. URL: https://www.flightglobal.com/news/articles/nbaa-the-old-barr-is-back-365331/.

[160] Flightradar24 AB. *Flightradar24*. Accessed on 2018-12-04. 2018. URL: https://www.flightradar24.com.

[161] Her Majesty's Government. *Wireless Telegraphy Act 2006*. Statute. Nov. 2006.

[162] Amar Toor. *This Twitter Bot is Tracking Dictators' Flights In and Out of Geneva*. Retrieved on 2017-01-27. 2016. URL: http://www.theverge.com/2016/10/13/13243072/twitter-bot-tracks-dictator-planes-geneva-gva-tracker.

[163] Julien Pilet and François Pilet. *GVA Dictator Alert*. Online. Accessed on 2018-12-16. 2018. URL: https://twitter.com/GVA_Watcher.

[164] David Cenciotti. *Online flight tracking provides interesting details about Russian air bridge to Syria*. Online. Accessed on 2018-12-08. Sept. 2015. URL: https://theaviationist.com/2015/09/11/ads-b-exposes-russian-air-bridge-to-syria/.

[165] Mary Kirby. *How Ryanair monitors health of Boeing 737s without ACARS*. Online. Accessed on 2018-12-16. Sept. 2014. URL: https://runwaygirlnetwork.com/2014/09/06/how-ryanair-monitors-health-of-boeing-737s-without-acars/.

[166] Federal Office of Civil Aviation FOCA. *Swiss Aircraft Registry*. Online. Accessed on 2017-02-28. 2018. URL: https://www.bazl.admin.ch/bazl/en/home/specialists/aircraft/swiss-aircraft-registry.html.

[167] Mark Maremont and Tom McGinty. *FAA Is Set to Give Investors a Peek at M&A Air*. Online. Accessed on 2018-12-16. June 2011. URL: https://www.wsj.com/articles/SB10001424052702303499204576389923856575528.

[168] Woodrow Bellamy III. *ADS-B Security Risk Remains Unresolved for US Military*. Online. Accessed on 2018-10-28. Oct. 2018. URL: https://www.aviationtoday.com/2018/10/04/ads-b-security-risk-remains-unresolved-u-s-military/.

[169] "PM charters flight as royals use official RAF Airbus". In: *BBC* (Apr. 2017). Accessed 2018-10-28. URL: https://www.bbc.co.uk/news/uk-39510024.

[170] Bilderberg Meetings. *Participants*. Online. Accessed 2018-10-29. June 2018. URL: https://www.bilderbergmeetings.org/participants.html.

[171] Fon Mathuros Chantanayingyong. *Who's coming to Davos 2018?* Online. Accessed 2018-10-29. Jan. 2018.

[172] David Cenciotti. *Exclusive: all the details about the air ops and aerial battle over Turkey during the military coup to depose Erdogan*. Online. Accessed on 2019-03-21. July 2016. URL: https://theaviationist.com/2016/07/18/exclusive-all-the-details-about-the-aerial-battle-over-turkey-during-the-military-coup/.

[173] Olga Gurtovaya. "Maintaining Privacy in a World of Technological Transparency: The BARR Program's ups and downs in Changing Times". In: *Journal of Air Law and Commerce* 77 (2012), pp. 569–603.

[174] Federal Aviation Administration. *Traffic Flow Management System*. Online. Accessed on 2018-12-08. July 2018. URL: https://aspmhelp.faa.gov/index.php/Traffic_Flow_Management_System_(TFMS).

[175] Federal Aviation Administration. *Traffic Flow Management System*. Online. Accessed on 2018-12-08. 2018. URL: https://www.faa.gov/about/office_org/headquarters_offices/ang/offices/tc/library/storyboard/detailedwebpages/tfms.html.

[176] Tony Webster. *FAA: List of Blocked Aircraft (BARR List/ASDI Block List)*. Online. Accessed on 2018-10-10. Oct. 2017. URL: https://www.muckrock.com/foi/united-states-of-america-10/faa-list-of-blocked-aircraft-barr-listasdi-block-list-34713/.

[177] FlightAware. *FlightAware Global – Frequently Asked Questions*. Online. Accessed on 2018-10-29. Oct. 2018. URL: https://uk.flightaware.com/commercial/global.

[178] Flightradar24. *Aircraft Unblocking*. Online. Accessed on 2018-10-29. Oct. 2018. URL: https://www.flightradar24.com/business/aircraft-unblocking.

[179] Flightradar24 AB. *Flightradar24 FAQs*. 2018. URL: https://www.flightradar24.com/faq (visited on 03/15/2018).

[180] *ADS-B Exchange*. Online. Accessed 2018-10-29. Oct. 2018. URL: https://www.adsbexchange.com.

[181] NATS. *Airspace Explorer FAQs*. Accessed on 2018-11-24. 2018. URL: https://www.nats.aero/ae-home/faqs/.

[182] Federal Aviation Administration. *Access to Aircraft Situation Display to Industry (ASDI) and National Airspace System Status Information (NASSI) Data.* Accessed on 2018-12-16. May 2012. URL: https://www.govinfo.gov/content/pkg/FR-2012-05-09/html/2012-11251.htm.

[183] Chad Trautvetter. *FltPlan Flight Privacy Program Exposes Tangled FAA Policy.* Accessed on 2018-10-24. 2011. URL: https://www.ainonline.com/aviation-news/aviation-international-news/2011-08-31/fltplan-flight-privacy-program-exposes-tangled-faa-policy.

[184] FltPlan.com. *Flying in Private*. Online. Accessed on 2018-10-29. Oct. 2018. URL: https://flttrack.fltplan.com/FltPlanInfo/DCMCallSigns.htm.

[185] The Economist. *Latin American leaders are embarrassed by their aeroplanes*. Online. Accessed on 2018-12-08. Sept. 2018. URL: https://www.economist.com/the-americas/2018/09/27/latin-american-leaders-are-embarrassed-by-their-aeroplanes.

[186] The Economist. *All-you-can-fly membership models are slowly catching on*. Online. Accessed on 2018-12-08. Oct. 2018. URL: https://www.economist.com/gulliver/2017/10/10/all-you-can-fly-membership-models-are-slowly-catching-on.

[187] Directorate of Air Traffic Management. *Air Traffic Management Circular – Automatic Dependent Surveillance – Broadcast (ADS-B)*. Tech. rep. 15. Accessed on 2018-10-28. New Delhi, India: Airports Authority of India, Dec. 2014. URL: https://www.aai.aero/en/system/files/resources/ATMC15of2014.pdf.

[188] K. Sampigethaya and R. Poovendran. "Flight Privacy in the NextGen: Challenges and Opportunities". In: *Integrated Communications, Navigation and Surveillance Conference (ICNS), 2013*. Apr. 2013, pp. 1–15.

[189] FreeFlight Systems. *ADS-B and Privacy – Not all Datalinks are Created Equal*. Online. Accessed 2018-10-29. July 2018. URL: https://www.freeflightsystems.com/blog/2018/07/10/ads-b-and-privacy-not-all-datalinks-are-created-equal/.

[190] Mike Collins. *ADS-B: Broadcast Clarity – Nonelectrical Aircraft, NPE, Rebate Updates*. Online. Accessed on 2017-10-29. Apr. 2017. URL: https://www.aopa.org/news-and-media/all-news/2017/april/pilot/adsb-broadcast-clarity.

[191] Jim Wolper. *Security Risks of Laptops in Airline Cockpits*. Online. Accessed on 2018-10-12. Dec. 1998. URL: http://catless.ncl.ac.uk/Risks/20/12#subj4.

[192] Ken Pascoe. *ACARS and Error Checking*. Online. Accessed on 2018-12-04. Dec. 2015. URL: http://www.flight.org/acars-and-error-checking.

[193] Airline Pilots Association. *Aviation Cyber Security: The Pilot's Perspective*. Tech. rep. Washington: Air Line Pilots Association Int'l, 2017. URL: https://www.rtca.org/sites/default/files/symposium_2017_cybersecurity_white_paper_digital.pdf.

[194] Raju Patel. *Managing Cybersecurity Risk in Weapons Systems*. Presentation. Accessed on 2018-12-16. Mar. 2017. URL: https://www.omg.org/news/meetings/tc/va-17/special-events/cybersecurity-pdf/Dr-Raju-Patel-Managing-Cybersecurity-Risk-in-Weapons-Systems-3-21-17.pdf.

[195] Dorothy Saul-Pooley. "Radio Navigation & Instrument Flying". In: ed. by Dorothy Saul-Pooley and Jonathan Shooter. Waypoint. Pooley's, 2017. Chap. 17, p. 359.

[196] International Civil Aviation Organization. "Annex 11 to the Convention on International Civil Aviation—Air Traffic Services". In: 13th. 2001. Chap. 4, 4.4.

[197] Federal Aviation Administration. *Equip ADS-B*. Online. Accessed on 2018-12-14. 2018. URL: https://www.faa.gov/nextgen/equipadsb/.

[198] Eurocontrol. *Aircraft Equipage Requirements in the European Commission IRs 1207/2011 and 1028/2014*. Online. Accessed on 2018-12-04. 2017. URL: https://www.eurocontrol.int/spi-ir.

[199] Martin Strohmeier et al. "Assessing the Impact of Aviation Security on Cyber Power". In: *Cyber Conflict (CYCON), 8th International Conference on*. IEEE. 2016.

[200] European Commission. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*. Regulation 2016/679. European Parliament, May 2018. URL: https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1532348683434&uri=CELEX:02016R0679-20160504.

[201] Matthew Smith et al. "Analyzing Privacy Breaches in the Aircraft Communications Addressing and Reporting System (ACARS)". In: *ArXiv e-prints* (May 2017). URL: https://arxiv.org/pdf/1705.07065.pdf.

[202] International Civil Aviation Organization. "Global Air Navigation Plan 2016-2030". In: 5th ed. Doc 9750-AN/963. 2016, p. 97. URL: https://www.icao.int/airnavigation/Documents/GANP-2016-interactive.pdf.

[203] Honeywell, Inc. *Primus Epic Integrated Avionics*. Online. Accessed on 2018-12-12. 2018. URL: https://aerospace.honeywell.com/en/products/cockpit-systems/primus-epic.

[204] Vinay M. Igure, Sean A. Laughter, and Ronald D. Williams. "Security issues in SCADA networks". In: *Computers & Security* 25.7 (2006), pp. 498 –506. URL: http://www.sciencedirect.com/science/article/pii/S0167404806000514.

[205] Y. Mo et al. "Cyber–Physical Security of a Smart Grid Infrastructure". In: *Proceedings of the IEEE* 100.1 (Jan. 2012), pp. 195–209.

[206] Aloke Roy. *Secure Aircraft Communications Addressing and Reporting System (ACARS)*. US Patent 6677888. 2004.

[207] Nikita Borisov, Ian Goldberg, and David Wagner. "Intercepting Mobile Communications: The Insecurity of 802.11". In: *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom '01)*. 2001, pp. 180–189.

[208] Aircraft Owners and Pilots Association. *Learn to Fly*. Online. Accessed on 2018-12-13. 2018. URL: https://www.aopa.org/training-and-safety/learn-to-fly.

[209] Aviation Insider. *Type Rating*. Online. Accessed on 2018-12-13. 2018. URL: https://www.aviationinsider.co.uk/pilots/type-rating/.

[210] Civil Aviation Authority. *Multi pilot type rating for aeroplanes*. Online. Accessed on 2018-12-13. 2018. URL: https://www.caa.co.uk/Commercial-industry/Pilot-licences/Aeroplanes/Multi-pilot-type-rating-for-aeroplanes/.

[211] *Final Report A-00XCENIPA2008*. Tech. rep. Aeronautical Accident Investigation and Prevention Center, Sept. 2006. URL: https://skybrary.aero/bookshelf/books/546.pdf.

[212] *Air Accident Investigation Commission - Prelimiary Report Boeing 717-412F*. Tech. rep. Interstate Aviation Committee, Jan. 2017.

[213] *Aircraft Accident Investigation Report*. Tech. rep. National Transportation Safety Committee, Ministry of Transportation, Republic of Indonesia, May 2012.

[214] German Federal Bureau of Aircraft Accidents Investigation. *Investigation Report AX001-1-2/02*. Tech. rep. May 2004. URL: http://www.bfu-web.de/EN/Publications/Investigation%20Report/2002/Report_02_AX001-1-2_Ueberlingen_Report.pdf.

[215] Robert T. Hays et al. "Flight Simulator Training Effectiveness: A Meta-Analysis". In: *Military Psychology* 4.2 (1992), pp. 63–74.

[216] Eduardo Salas, Clint A Bowers, and Lori Rhodenizer. "It is not how much you have but how you use it: Toward a rational use of simulation to support aviation training". In: *The International Journal of Aviation Psychology* 8.3 (1998), pp. 197–208.

[217] N Dahlstrom et al. "Fidelity and validity of simulator training". In: *Theoretical Issues in Ergonomics Science* 10.4 (2009), pp. 305–314.

[218] Janeen A Kochan, Eyal G Breiter, and Florian Jentsch. "Surprise and unexpectedness in flying: Database reviews and analyses". In: *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 48. 3. SAGE Publications Sage CA: Los Angeles, CA. 2004, pp. 335–339.

[219] Annemarie Landman et al. "The influence of surprise on upset recovery performance in airline pilots". In: *The International Journal of Aerospace Psychology* 27.1-2 (2017), pp. 2–14.

[220] Stephen M Casner, Richard W Geven, and Kent T Williams. "The effectiveness of airline pilot training for abnormal events". In: *Human Factors* 55.3 (2013), pp. 477–485.

[221] Annemarie Landman et al. "Training Pilots for Unexpected Events: A Simulator Study on the Advantage of Unpredictable and Variable Scenarios". In: *Human Factors* 0.0 (2018). PMID: 29913086.

[222] Patrick Gontar et al. "Are pilots prepared for a cyber-attack? A human factors approach to the experimental evaluation of pilots' behavior". In: *Journal of Air Transport Management* vol. 69 (June 2018), pp. 26–37. URL: https://hal.archives-ouvertes.fr/hal-01739475.

[223] Jan-Philipp Buch et al. "What the Hack Happened to the Flight Deck: Analyzing the Impact of Cyberattacks on Commercial Flight Crews". In: *AIAA SciTech 2019*. Jan. 2019.

[224] European Aviation Safety Agency. *Impact Assessment of Cybersecurity Threats*. Tech. rep. EASA_REP_RESEA_2016_1. European Union, 2018.

[225] Sylvia Pfeifer. *Gatwick flights suspended again after another drone sighting*. Online. Accessed on 2019-03-25. Dec. 2018. URL: https://www.ft.com/content/b6fbacdc-0547-11e9-99df-6183d3002ee1.

[226] Julia Kollewe and Gwyn Topham. *EasyJet says Gatwick drone chaos cost it £15m*. https://www.theguardian.com/business/2019/jan/22/easyjet-gatwick-drone-cost-brexit-flights. Accessed on 2019-02-15. Jan. 2019.

[227] BBC. *Gatwick drone policing costs 'shocking'*. Online. Accessed on 2019-03-25. Mar. 2019. URL: https://www.bbc.co.uk/news/uk-england-47696499.

[228] Laminar Research. *X-Plane 11*. Online. Accessed on 2018-11-21. Aug. 2018. URL: https://www.x-plane.com/.

[229] Barry C. Breen. "Digital Avionics Handbook". In: ed. by U. Ferrell C. Spitzer and T. Ferrell. Third. CRC Press, 2015. Chap. 21, 21.1–21.12.

[230] Christian Wolff. *Frequency-Modulated Continuous-Wave Radar (FMCW Radar)*. Online. Accessed on 2018-11-25. 2018. URL: http://www.radartutorial.eu/02.basics/Frequency%20Modulated%20Continuous%20Wave%20Radar.en.html#abs8.

[231] Ron Hovav. *Airbus Erroneus Radio Altitudes*. Tech. rep. FMG/15 – WP/08. International Civil Aviation Organization, 2011. URL: https://www.icao.int/safety/acp/ACPWGF/ACP-WG-F-25/ACP-WGF25-IP07_Appendix1_FMG15%20WP08%20-%20Airbus%20Erroneous%20Radio%20Altitudes.pdf.

[232] International Civil Aviation Organization. *Radio Altimeter Spectrum*. Reference Number RPGITUWRC2019-P08, Accessed on 2018-11-21. Feb. 2018. URL: https://www.icao.int/NACC/Documents/Meetings/2018/RPG/RPGITUWRC2019-P08.pdf.

[233] SKYbrary. *Response to a "PULL UP" Warning*. Acessed on 2018-08-30. Sept. 2017. URL: https://www.skybrary.aero/index.php/Response_to_a_%22PULL_UP%22_Warning.

[234] Eurocontrol. "Flying without a transponder—10 minutes is all it can take". In: *NetAlert* 19 (May 2014), p. 5. URL: https://www.eurocontrol.int/sites/default/files/publication/files/NetAlert-19.pdf.

[235] Federal Aviation Administration. "Introduction to TCAS II Version 7.1". In: U.S. Department of Transport, 2011. Chap. 4, p. 23.

[236] Federal Aviation Administration. "Introduction to TCAS II Version 7.1". In: U.S. Department of Transport, 2011. Chap. 1, 17.

[237] Federal Aviation Administration. "Introduction to TCAS II Version 7.1". In: U.S. Department of Transport, 2011. Chap. 1, 17–19.

[238] Federal Aviation Administration. "Introduction to TCAS II Version 7.1". In: U.S. Department of Transport, 2011. Chap. 1, 22–24.

[239] ARINC. *ARINC Characteristic 735B-2: Traffic Computer TCAS and ADS-B Functionality*. Tech. rep. 735B-2. 2015.

[240] Nick Foster. *gr-air-modes*. Online. Accessed on 2018-11-23. Sept. 2017. URL: https://github.com/bistromath/gr-air-modes.

[241] Salvatore Sanfilippo. *dump1090*. Online. Accessed on 2018-11-23. Jan. 2017. URL: https://github.com/antirez/dump1090.

[242] OpenSky Network. *dump1090 Decoder with High Precision Timestamping*. Online. Accessed on 2018-11-23. Nov. 2018. URL: https://github.com/openskynetwork/dump1090-hptoa.

[243] Linar Yusupov. *ADSB-Out*. Online. Accessed on 2018-11-23. Dec. 2017. URL: https://github.com/lyusupov/ADSB-Out.

[244] C. Young. *Stratux*. Online. Accessed on 2018-12-13. 2018. URL: http://stratux.me/.

[245] Eurocontrol. "Transponder failure is not always total". In: *NetAlert* 19 (May 2014), pp. 6–7. URL: https://www.eurocontrol.int/sites/default/files/publication/files/NetAlert-19.pdf.

[246] SKYbrary. *Airborne Collision Avoidance System (ACAS)*. Accessed on 2018-12-04. Dec. 2017. URL: https://www.skybrary.aero/index.php/Airborne_Collision_Avoidance_System_(ACAS)#Complying_with_RAs.

[247] Federal Aviation Administration. "Instrument Flying Handbook". In: FAA-H-8083-15B. U.S. Department of Transport, 2012. Chap. 9, 9.35–9.38.

[248] Flight Operations Suppoer & Line Assistance. *Getting to Grips with CAT II/CAT III Operations*. Tech. rep. Airbus, 2001. URL: https://www.skybrary.aero/bookshelf/books/1480.pdf.

[249] Federal Aviation Administration. *Lighting Systems – Precision Approach Path Indicators (PAPI)*. Online. Accessed on 2018-11-23. June 2015. URL: https://www.faa.gov/about/office_org/headquarters_offices/ato/service_units/techops/navservices/lsg/papi/.

[250] Ofcom. *UK Amateur Radio License – Terms, Conditions and Limitations*. Accessed on 2018-08-31. 2018. URL: https://www.ofcom.org.uk/__data/assets/pdf_file/0027/62991/amateur-terms.pdf.

[251] International Civil Aviation Organization. "Annex 10 to the Convention on International Civil Aviation—Aeronuatical Telecommunications". In: vol. 1. 2006. Chap. 3, 3.19–3.20.

[252] Federal Aviation Administration. "Instrument Flying Handbook". In: FAA-H-8083-15B. U.S. Department of Transport, 2012. Chap. 9, 9.40.

[253] Chad Trautvetter. *FAA Reminds Pilots of Possible ILS Interference*. Online. Accessed on 2018-11-23. Apr. 2012. URL: https://www.ainonline.com/aviation-news/business-aviation/2012-04-24/faa-reminds-pilots-possible-ils-interference.

[254] Jordan Miller. "Handling ILS anomalies". In: *IFR Magazine* (Sept. 2012). Accessed on 2018-12-13. URL: http://www.ifr-magazine.com/issues/1_24/features/Handling-ILS-anomalies_240-1.html.

[255] International Civil Aviation Organization. "Airborne Collision Avoidance System (ACAS) Manual". In: Doc 9863 AN/461. 2006. Chap. 6, 6.3–6.4.

[256] William Roberson and James A. Johns. "Fuel Conservation Strategies: Descent and Approach". In: *AERO* 38 (2010), pp. 25–28. URL: https://www.boeing.com/commercial/aeromagazine/articles/qtr_02_10/5/.

[257] International Air Transport Association (IATA). *Jet Fuel Price Monitor*. Accessed on 2018-11-20. Nov. 2018. URL: https://www.iata.org/publications/economics/fuel-monitor/Pages/index.aspx.

[258] Civil Aviation Authority. *Disruptive Passengers*. Online. Accessed on 2018-11-20. 2018. URL: https://www.caa.co.uk/Passengers/On-board/Disruptive-passengers/.

[259] Brendan Dorsey. *Hawaiian Airlines Passenger Fined $100,000 for Bad Behavior*. Online. Accessed on 2018-11-20. Aug. 2017. URL: https://thepointsguy.com/2017/08/hawaiian-airlines-passenger-fined/.

[260] *Limerick court fines man €1,000 after disrupting flight to tune of €100k*. Online. Accessed on 2018-11-20. Apr. 2015. URL: https://www.irishexaminer.com/breakingnews/ireland/limerick-court-fines-man-1000-after-disrupting-flight-to-tune-of-100k-674943.html.

[261] Eurocontrol. "European Action Plan for Air Ground Communications Safety". In: 2006. Chap. AGC4, pp. 51–52.

[262] Rohde & Schwarz. *Automatic monitoring of radio signals at airports interfering with radio services*. Online. Accessed on 2019-03-25. Mar. 2019. URL: https://www.rohde-schwarz.com/ca/applications/automatic-monitoring-of-radio-signals-at-airports-interfering-with-radio-services-application-card_56279-4816.html.

[263] CRFS. *Resolving RF interference at airports*. Online. Accessed on 2019-03-25. Mar. 2019. URL: https://www.crfs.com/applicationstory/resolving-rf-interference-airports/.

[264] S. Rajendran et al. "Electrosense: Open and Big Spectrum Data". In: *IEEE Communications Magazine* (2018).

[265] FlightAware. *Frequently Asked Questions*. Accessed on 2018-03-15. 2018. URL: https://uk.flightaware.com/about/faq#military.

[266] WiMAX Forum. *AeroMACS*. Accessed on 2018-08-31. 2018. URL: http://wimaxforum.org/Page/AeroMACS.

[267] IEEE Computer Society. *IEEE Standard for Local and Metropolitan Area Networks Part 16: Air Interface for Fixed and Mobile Broadband Wireless Access Systems*. Tech. rep. 802.16e-2005. 2015.

[268] INDRA. *AEROMACS – Security Analysis*. Tech. rep. 15.02.17. SESAR Joint Undertaking, 2014. URL: https://www.eurocontrol.int/sites/default/files/article/content/documents/communications/d08-aeromacs-safety-security-analysis-p15.02.07.pdf.

[269] D. Zeng et al. "DME Potential for Data Capability". In: *2018 Integrated Communications, Navigation, Surveillance Conference (ICNS)*. Apr. 2018, 4D3–1–4D3–16.

[270] Martin Strohmeier. "Security in Next Generation Air Traffic Communication Networks". PhD thesis. University of Oxford, 2016.

[271] Donald L. McCallie. “Exploring Potential ADS-B Vulnerabilities in the FAA’s NextGenAir Transportation System”. MA thesis. Air Force Institude of Technology, 2011.

[272] Airbus S.A.S. *EASA certifies new "Autopilot/Flight Director" TCAS mode for A380*. Online. Accessed on 2018-12-16. Aug. 2009. URL: https://www.airbus.com/newsroom/press-releases/en/2009/08/easa-certifies-new-autopilot-flight-director-tcas-mode-for-a380.html.