

Unknown Target: Uncovering and Detecting Novel In-Flight Attacks to Collision Avoidance (TCAS)

Giacomo Longo*, Giacomo Ratto*, Alessio Merlo*, and Enrico Russo†

*CASP - University School of Advanced Defense Studies, Rome, Italy

firstname.lastname@unicasd.it

†DIBRIS - University of Genova, Genova, Italy

enrico.russo@unige.it

Abstract—The Traffic alert and Collision Avoidance System (TCAS) is a mandatory last-resort safeguard against mid-air collisions. Despite its critical safety role, the system’s unauthenticated and unencrypted communication protocols present a long-identified security risk. Although researchers have previously demonstrated practical injection attacks, official advisories have assessed these vulnerabilities as confined to laboratory environments, also stating that no mitigation is currently available. In this paper, we challenge both assertions. We present compelling evidence suggesting that an in-flight cyber-attack targeting TCAS has already occurred. Through a detailed analysis of public flight and communications data from a series of anomalous events involving multiple aircraft, we identify a distinct signature consistent with a ghost plane injection attack. We detail how this novel attack exploits legacy protocol features and describe three strategies of increasing sophistication; the most aggressive of these can reduce a target’s perceived range by over 3.5 kilometers, sufficient to trigger collision avoidance advisories on victim aircraft from a significant standoff distance. We implement and experimentally evaluate the attack strategy most consistent with the observed incident, achieving a spoofed range reduction of 1.9 km, confirming its feasibility. Furthermore, to provide a basis for responding to such threats, we propose a novel, backward-compatible methodology to geographically localize the source of such attacks by repurposing the TCAS alert data broadcast by victims. In simulated scenarios of the most plausible attack variant, our approach achieves a median localization accuracy of 855 meters. Applying this technique to real-world incident data, we were able to identify the anomaly and the likely origin of the observed ghost plane injection attack.

I. INTRODUCTION

Modern aviation safety is based on a defense-in-depth architecture, where independent systems provide layers of protection against catastrophic failure.

The Traffic alert and Collision Avoidance System (TCAS), internationally standardized as ACAS II, functions as a system of last resort within this architecture. Operating independently

of ground-based Air Traffic Control (ATC), TCAS provides an autonomous safeguard against mid-air collisions. Its directives are not mere suggestions; pilots are required to follow their Resolution Advisories (RAs) immediately and without question, even if they contradict an ATC instruction [1], [2], [3], [4].

This authority places TCAS at the heart of not only operational safety, but also airspace continuity and economic stability. A single spurious TCAS alert, whether genuine or maliciously induced, forces abrupt, high-rate maneuvers across multiple aircraft, with documented cases of crew and passenger injuries [5], [6], [7], [8], [9]. If such an alert is not an isolated technical mishap, but rather the result of an intentional disruption, the ripples can multiply: simultaneous false RAs across a busy terminal airspace could saturate controller communications [10], [11], trigger cascading go-arounds or missed approaches [12], and force authorities to ground flights or shut down an entire airspace sector [13]. Recent headline-grabbing incidents, such as the 2018 Gatwick drone shutdown [14] or the 2023 failure of the UK’s NATS system [15], have demonstrated the severe economic and operational impact of even brief airspace disruptions.

Despite its critical role, TCAS was designed for safety and reliability rather than security. Its communication protocols lack authentication and encryption, creating vulnerabilities that security researchers have long identified as potential attack vectors. However, only after these vulnerabilities were demonstrated in practice by triggering fake TCAS alerts [16] did the U.S. Cybersecurity and Infrastructure Security Agency (CISA) issue an advisory [17]. Although acknowledging the exploits, the advisory sought to contextualize the risk by framing the vulnerabilities as *difficult to exploit outside a laboratory setting*, requiring very specific conditions and *being unlikely to occur in real-world scenarios*. In addition, it described them as *not remotely exploitable* and noted that *no mitigation is currently available*.

This paper directly challenges such assertions. Using publicly available flight and communication data from the anomalous events at Ronald Reagan Washington National Airport (DCA) on March 1, 2025 [18], we show that a live cyberattack

against TCAS has already occurred. We identify and characterize a ghost aircraft injection that exploits TCAS backward compatibility with legacy transponders. This attack represents a simpler variant than prior laboratory demonstrations, which we have studied and implemented as a novel airborne attack. We further address the claim that no mitigation exists by introducing a backward-compatible method to detect and localize rogue TCAS transmissions, providing the basis for a practical mitigation strategy based solely on existing operational data.

This paper makes the following contributions.

- We provide the first in-depth analysis of a real-world TCAS anomaly, showing that an incident initially dismissed as a technical malfunction or benign interference is, in fact, consistent with a deliberate cyberattack. This analysis challenges the official assessment that such exploits are purely theoretical or cannot succeed outside a controlled laboratory environment.
- Starting from the DCA scenario and associated data, we identify and formalize three novel attack variants that exploit the TCAS legacy Mode C behavior to inject ghost aircraft. The most aggressive of these can reduce a target's perceived range by over 3.5 kilometers. These attacks reduce overall complexity compared to known Mode S-based techniques.
- We reproduce the above attack variant in a controlled setting. With a certified ramp test set, we verify interaction with real transponders and demonstrate spoofing capabilities that reduce perceived range by over 1.9 kilometers. These results support the technical soundness of the DCA reconstruction.
- Inspired by existing radio localization techniques approaches that can help detect and locate rogue transmitters (see § VI), we design a novel methodology that operates solely on existing operational data, without any additional infrastructure. In simulation, the method achieves a median localization accuracy of 855 meters. When applied to the DCA incident, the system required only two aircraft encounters to constrain the estimated source area to approximately 4.3 km^2 , enabling actionable localization in less than an hour from the start of the anomaly.

We begin with a background in § II, before analyzing the incident at DCA in § III. § IV describes and validates the attack methodology. § V presents our detection approach. § VI reviews related work and § VII concludes the article.

II. BACKGROUND

A. Arrival Procedures

Arrival procedures are the standardized steps that guide an aircraft through its approach and landing. Aircraft typically follow visual approaches, guided primarily by pilot observation and air traffic control instructions, or precision-based Instrument Arrival Procedures (IAPs). IAPs are predefined instrument flight procedures that may rely on ground-based radio navigation aids or satellite-guided Area Navigation (RNAV) systems.

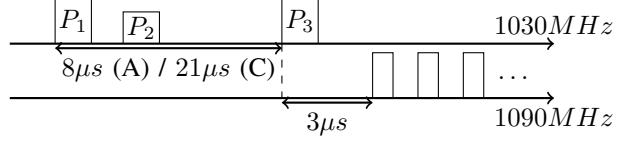


Fig. 1: Mode A/C interrogation and reply pulses.

RNAV procedures, increasingly prevalent [19, §12], are defined by a sequence of named waypoints, often referred to as *fixes*, interconnected to create a flight path. These procedures also incorporate mandatory altitude restrictions at each fix, ensuring obstacle clearance and a mandatory descent profile when approaching the runway. During approach, pilots report “final” when on the final segment near the runway waypoint, and short “final” immediately before landing.

B. Air Traffic Control

From the perspective of ATC, the approach phase involves a tiered system of monitoring, information provision, and procedural guidance [20]. Controllers inform arriving aircraft of weather conditions, active runway, and nearby traffic and direct them to follow established procedures. This control relies on a combination of technologies. Although direct visual observation is used when conditions allow, modern approach control is overwhelmingly based on secondary radar, which uses transponder signals to track the position, altitude, and identity of the aircraft. Central to this process is voice communication, enabling real-time exchange of instructions and situational awareness between controllers and pilots [21]. Such voice communications are kept in a standardized phraseology to ensure concise and unambiguous exchange [22, §12].

For example, both waypoints and runway designations employ short, mnemonic names, such as “FONVI” for a waypoint or “27” to indicate a magnetic heading. During their duties, air traffic controllers issue headings, altitude assignments, and clearances to proceed through IAP waypoints, ensuring a safe and orderly flow of traffic. Importantly, controllers are also responsible for managing potential traffic conflicts and coordinating evasive maneuvers. This involves issuing direct instructions for aircraft to change course or altitude, accommodating pilot-initiated maneuvers like go-arounds or collision avoidance responses, and managing unforeseen situations to maintain proper separation and prevent collisions.

C. Modes A and C

Secondary Surveillance Radar (SSR) Modes A and C are foundational cooperative surveillance technologies governed by ICAO Annex 10, Volume IV [23]. These protocols enable ATC to identify and track aircraft. The system operates on an interrogation-and-response basis, where a ground station or airborne interrogator transmits a query on 1030 MHz (also called *uplink* frequency), and the aircraft's transponder replies on 1090 MHz (*downlink* frequency). Communication over both uplink and downlink channels uses amplitude modulation.

In uplink, the interrogation mode is determined by the timing between two pulses, P_1 and P_3 , in the interrogation signal [24], as pictured in Figure 1. An $8\mu\text{s}$ spacing constitutes a Mode A interrogation, which elicits a 12-bit *squawk* code (a 4-digit octal number from 4096 possibilities) set by the pilot. A $21\mu\text{s}$ spacing constitutes a Mode C interrogation, which requests the aircraft's pressure altitude (PA), reported in 100-foot¹ increments. Both Mode A and Mode C interrogations utilize identical reply transmission formats on the downlink.

The reply comprises two framing pulses with 13 information pulses positioned in between them. The presence or absence of these pulses in between the framing pulses encodes either the identity code for Mode A or the altitude information for Mode C. Ground-based SSR systems, i.e., the ones employed by ATC, calculate an aircraft's range by measuring the signal's total round-trip time and subtracting a fixed processing delay of $3\mu\text{s}$; the aircraft's azimuth is determined by the rotational position of the directional antenna upon receiving the reply.

$$\rho = \frac{1}{2} \cdot c \cdot (\Delta T - 3\mu\text{s}) \quad (1)$$

The range calculation is performed according to Eq. 1, with the range ρ calculated as the response delay ΔT multiplied by the speed of light c and divided by 2.

D. Mode S and ADS-B

SSR Mode S (Select) evolved from Modes A and C to reduce interference in busy airspace while increasing the information capacity for both ground stations and aircraft [25]. Operating on the same 1030/1090 MHz frequency pair, Mode S introduces selective interrogation through globally unique 24-bit ICAO addresses, allowing targeted queries to individual aircraft. The system employs differential binary phase shift keying (D-BPSK) for uplink and amplitude modulation for downlink, with messages spanning 56 or 112 bits. Unlike earlier modes, Mode S includes error detection parity bits in every message.

To further reduce radio congestion, Mode S now includes the Automatic Dependent Surveillance-Broadcast (ADS-B) functionality [26]. Rather than waiting to be interrogated, an ADS-B equipped aircraft automatically and periodically broadcasts its state vector (identity, high-precision GNSS-derived position, altitude, and velocity) to all listeners.

E. TCAS

The TCAS, standardized by the ICAO as ACAS II, is an airborne system mandated as the final defense against mid-air collisions. Independent of ground-based ATC, TCAS directly interrogates nearby aircraft using SSR protocols, building a 3D map of surrounding traffic. Although TCAS relies on Mode S for air-to-air coordination, it also supports interacting with non-TCAS targets by actively interrogating such targets over Mode C [27, §2.2.3.8.1]. This process allows it to receive altitude replies from nearby aircraft that are not Mode S-equipped. TCAS transponders determine range like ground-based SSR

systems, with azimuth being estimated by a static radio direction finding array instead of a rotating antenna. According to the system specifications, this estimation method can result in bearing errors of up to 27 degrees [27, §2.2.4.6.4.2.1].

TCAS issues two alert levels: Traffic Advisory (TA), which alerts pilots of potential conflicts without commands, and RA, providing direct vertical maneuver commands (e.g., "Climb" or "Descend"). Pilots must comply immediately with the RA commands, even if they contradict the ATC instructions. When two TCAS-equipped aircraft detect a mutual threat, their systems communicate directly using Mode S to negotiate complementary RAs. This automated coordination is a safety feature that prevents two aircraft from making opposing, and thus disastrous [28], avoidance maneuvers.

However, if the intruder is a non-TCAS aircraft with only a Mode C transponder, the RA is determined unilaterally by the TCAS-equipped aircraft. During an RA, ADS-B broadcasts a dedicated message announcing the aircraft's involvement in a TCAS-commanded collision avoidance maneuver [23, §4.3.7.3]. This broadcast includes detailed threat information, which serves to alert the ATC during near-miss events. This message includes the TTI (Threat Type Indicator) field, which identifies whether the intruder has a Mode S transponder or legacy Mode A/C equipment. The TID (Threat Identity Data) field contains the Mode S address of the threat aircraft or, for threats not equipped with Mode S, their position data comprising (i) TIDA: altitude from Mode C, (ii) TIDR: range in 0.05 nautical mile² increments, and (iii) TIDB: bearing relative to the TCAS aircraft heading in increments of 6 degrees. When multiple threats are present, the broadcast contains only data for the most recently declared threat and indicates this condition in a dedicated MTE (Multiple Threat Encounter) field. The ARA (Active RAs) field indicates the characteristics of the RA, signalling the corrective actions, if any, the pilots must take.

III. AIRSPACE SURVEILLANCE FAILURES AT DCA

DCA is one of the most operationally constrained airports in the United States, primarily due to its location near restricted airspace and dense urban areas just kilometers from the neighboring city. As a result, the various IAPs used at DCA (see § II-A) are never simple "straight-in" paths. RNAV procedure Z, namely RNAV-Z, to Runway 19 exemplifies this challenge, weaving a path along the Potomac River for inbound traffic.

As shown in Figure 2, the lateral path of the RNAV-Z approach [29] is defined by a sequence of waypoints, DARIC, SUNEY, GREYZ, FONVI, JUBOL, WIRSO, FIROP, and finally RW19, that is, the threshold of Runway 19. The trajectory is designed to guide aircraft around restricted areas P-56A and P-56B [30, p.156] while aligning them for a safe approach to Runway 19. The RNAV-Z approach also includes a defined vertical profile, with mandatory crossing altitudes at each fix. The descent starts at 2700 ft above DARIC and

¹SI units: 1 foot (ft) is 0.3048 m.

²SI units: 1 nautical mile (nm) is 1852 m.

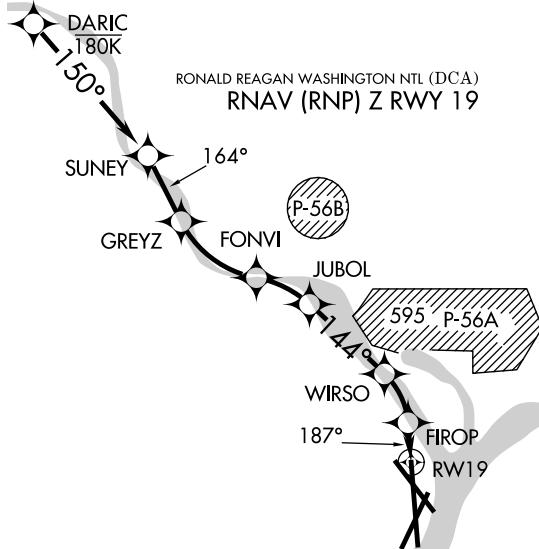


Fig. 2: Waypoints of RNAV-Z for Runway 19.

continues through successive constraints down to 241 ft at FIROP, following a three-degree path toward the runway.

In addition to RNAV-Z, Runway 19 at DCA offers various arrival procedures. We note two in particular: RNAV-Y, which adds the fix SLAKR just beyond JUBOL as an alternative lateral path along the Potomac, and Visual 19, a fully visual descent along the riverbank allowed only when the cloud ceiling and visibility exceed 3500 feet and three nautical miles.

On 1 March 2025, several aircraft operating along these arrival paths experienced a sequence of TCAS alerts that disrupted flight operations for more than three hours, from 11:09:40 to 14:10:25 UTC. At the time of writing, neither the Federal Aviation Administration (FAA) nor the National Transportation Safety Board (NTSB) had released an official explanation or preliminary report on the incident.

However, concerns have been formally raised by members of the U.S. House Committee on Transportation and Infrastructure (T&I) following an undisclosed briefing by FAA staff. In an official letter dated 14 April 2025 [31], the Committee reports “*ten resolution advisories and three aircraft go-arounds to commercial and U.S. Coast Guard aircraft in the vicinity of DCA*” and states that a spectrum analysis conducted by the FAA attributed the interference to counter-unmanned aircraft systems (c-UAS) deployed by the U.S. Secret Service (USSS) near the approach path. The following sections provide a factual description of events and our analysis, along with a threat model relevant to potential cyber-physical attacks.

A. Factual description of events

Data sources. The reconstruction relies on two publicly accessible sources: (i) ATC voice communications recorded on VHF channels, including raw audio [33] and curated transcripts of key events [34] (App. A), and (ii) ADS-B data, collected from ground receivers and made available by flight

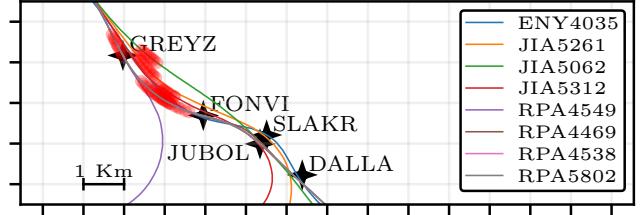


Fig. 3: Position of the aircraft at the moment of RA generation.

aggregators [35] due to the unencrypted and broadcast nature of the protocol.

The raw audio was used primarily to identify the affected flights. Ten distinct aircraft were involved: eight explicitly reported RAs: ENY4035, JIA5062, JIA5261, JIA5312, RPA4469, RPA4538, RPA4549, and RPA5802; while two only received TAs: JIA5146 and JIA5098.

ADS-B data includes 110 ACAS RA broadcast messages (Appendix A), emitted during the encounter. These messages refer to eight aircraft, consistent with those reporting RAs in voice communications. Each message includes information about the relative position and threat status of the intruder. TTI field uniformly reports a value of 2, indicating that TCAS classified intruders as Mode C transponder targets, allowing the inclusion of intruder-related fields in the broadcast (§ II-C). MTE flag is unset in 102 messages; the remaining 8, all from RPA4469, briefly report multiple intruders before reverting to a single threat, a fluctuation likely attributable to transient tracking errors by that plane. No radio exchanges mention multiple aircraft involved in a single encounter.

Observed events. We structured our analysis around a set of distinct event categories observed during the encounter. Table I presents one row per category, each describing: (i) supporting excerpts from voice transcriptions, (ii) the observation inferred from them, (iii) the corresponding ADS-B data fields, (iv) the interpretation derived from ADS-B, and (v) any relevant notes. We review each category in detail below.

Regarding RA, the location category (*RA loc.*) captures where the aircraft was when the RA was issued. Voice communications refer to the GREYZ and FONVI fixes as points where RAs occurred. ADS-B confirms this through GNSS-derived positions at RA time (Fig. 3), showing a spatial concentration of events along that corridor. The same figure reveals three go-arounds (JIA5261, JIA5312, and RPA4549), each deviating from the standard approach path immediately after their RA. Instructions (*RA instr.*) were reported as “descend” in the transcripts and match the sense encoded in the ADS-B message. Altitude values, that is, *Alt. (RA)*, show systematic differences: pilots referred to events around 1200 ft, based on a local reference (QNH), while RA messages placed the intruder at 2300 ft PA. A calibration using ADS-B data from aircraft on the DCA runway confirmed an offset of 325–375 ft, consistent with this difference.

Across all messages, the intruder’s altitude remained stable

TABLE I: Summary of key DCA event data from voice transcripts and ADS-B sources.

Category	Voice transcriptions	Voice (inferred)	ADS-B source	ADS-B (inferred)	Note
RA loc.	RPA4469 “...just about over FONVI..” RPA5802 “...RA over GREYZ...”	GREYZ-FONVI	Position report	GREYZ-FONVI	
RA instr.	RPA4469 “...we got a descending RA...” RPA5802 “...had an RA for descent ...”	descending	RA bcast status	descending	
Alt. (RA)	RPA4469 “...we were about 1200 feet ...” RPA5802 “...maybe about 1200–1100 ft ...”	1200 ft (QNH) + RA threshold	RA bcast status: TIDA	2300 ft (PA)	Voice reports use barometric (QNH) alt. ADS-B uses pressure alt. (PA).
TA loc.	JIA5197 “first one before SLAKR ...” “second one near DARIC ...” JIA5098 “one about short final ...” “the other one probably about 3 miles back ”	near DARIC, before SLAKR	N/A	N/A	“Short final” likely corresponds to SLAKR: “3 miles back” (from SLAKR) is consistent with DARIC.
Alt. (TA)	DCA TWR: “preceding arrival...TA...at 1200 feet ” RPA4538: “Yeah...” “target was 600 feet above us ”	1800 ft (QNH)	N/A	N/A	
Range	N/A	N/A	RA bcast status: TIDB	median 0.2nm, max 1nm	
Bearing	JIA5146 “...traffic right at our 11 o'clock ...”	11 o'clock	RA bcast status: TIDR	315°–345°	relative to own heading (head-up view).
Intruder visibility	RPA4469: “ Negative ” JIA5098 “one about short final ...” “the other one probably about 3 miles back ” RPA4538: “...did not see anything visually ” JIA5197: “...we didn't see anything either ” RPA5802: “We saw nothing out there though” DCA TWR: “ no known traffic between you and the field”	negative	N/A	N/A	DCA TWR “no known traffic” refers to absence of radar returns.
Visibility condition	JIA5098 “...on the River Visual 19 ” ENY3630 “...river visual 19 ”	3 miles	N/A	N/A	Arrival procedure requires 3500 ft ceiling and 3 miles visibility [32].

at 2300 ft. Only 4 of 110 messages reported 2500 ft, all from the same aircraft and likely attributable to undetected bit errors, a known possibility in Mode C transmissions, which lack parity checking [36]. In general, while voice reports are approximate and RA thresholds introduce some uncertainty, the data consistently depict a fixed-altitude intruder above the responding aircraft, aligning with the issuance of descending RAs.

Concerning TAs, they are not broadcast via ADS-B, and pilots are not required to report them to ATC. However, some TAs have been mentioned in radio communications. According to pilot reports, the locations (*TA loc.*) are both before and after the RA zone, with references to fixes near DARIC and before SLAKR. This distribution is consistent with the TCAS logic, which issues TAs under conditions less stringent than those for RAs. *Alt. (TA)* reports an altitude cited by a pilot of 1800 ft (QNH). Again, the discrepancy is within the expected range considering the pressure offset and the voice approximation.

The intruder geometry includes *range* and *bearing* information extracted from RA messages. Valid range estimates were available for 100 out of 110 events, based on the TIDR subfield. As shown in Figure 4a, all encounters occurred within 1 nm, with a median distance of 0.2 nm (0.37 km), confirming that RAs were consistently triggered at close horizontal range. No pilots reported precise range estimates, but one explicitly described the intruder as “at eleven o'clock”. Figure 4b shows the bearing distribution observed in the TIDB subfield, grouped in 30-degree intervals to reflect the 12-hour structure. The central value of 330° is consistent with the pilot's observation. The angular spread can be attributable to intrinsic

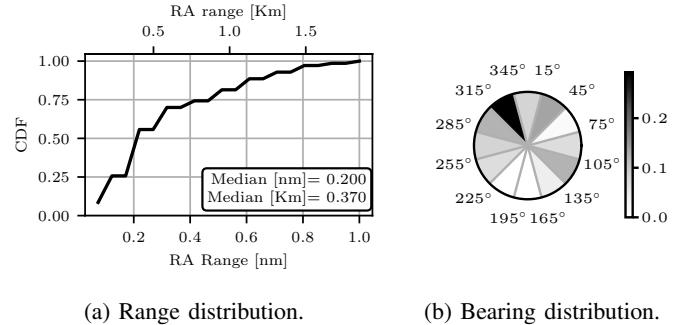


Fig. 4: RA broadcast status range (TIDR) and bearing (TIDB).

inaccuracy, which may deviate by up to 30 degrees [37], as well as to the challenges that different systems may face when estimating bearing at very short ranges. Regarding *intruder visibility*, the pilots never established visual contact, and the ATC did not report radar returns.

Finally, about *visibility condition*, at least two flights were cleared for a visual approach, which implies visibility conditions of at least 3 nm [32], as indicated in IAPs.

B. Analysis

We can first rule out equipment malfunction. The affected aircraft, manufactured by Bombardier and Embraer from 2004 to 2021 (Appendix B), all exhibited identical anomalous behavior. A simultaneous failure across multiple platforms, sustained over three hours and producing consistent Mode C signatures, is beyond the bounds of coincidental failure.

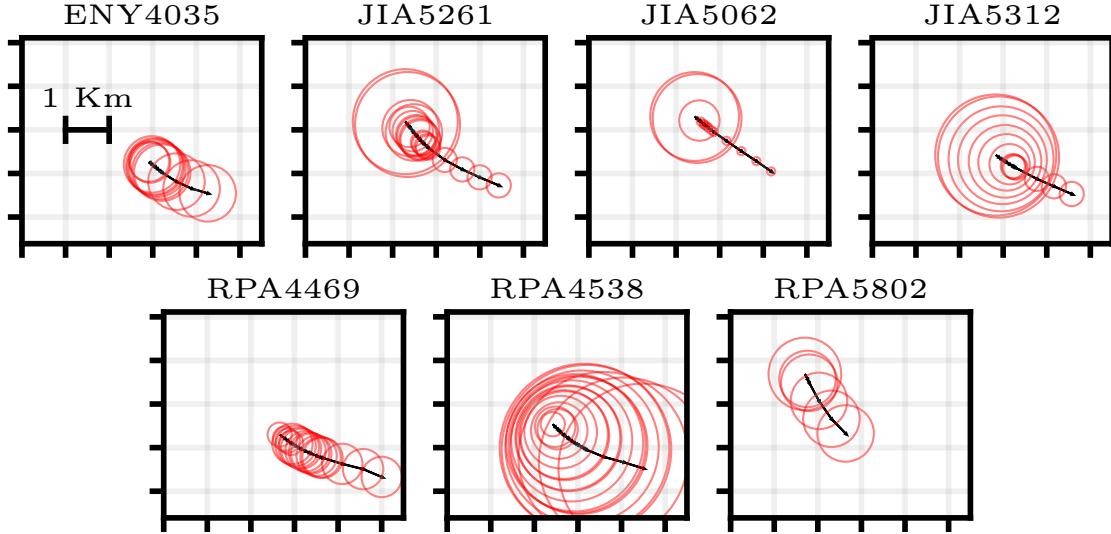


Fig. 5: Intruder range circles at aircraft positions during RAs.

These premises align with an independent reconstruction by Aireon [38], a provider of space-based ADS-B surveillance. Their analysis on Mode C data excluded TCAS ghosting or filtering issues with own-ship traffic. It concluded “*the data cannot say for certain, but it is possible the intruder was airborne or related to a ground-based transmitter used for testing or spoofing*”.

To further examine this conclusion, we turn to geometric evidence. Figure 5 shows the maximum range circles for the intruder, centered on the position of each aircraft during the RA events³. If the intruder were stationary, the circles would intersect at a fixed point. No such intersection exists. Thus, the intruder must have been moving alongside the aircraft, consistent with it being airborne.

However, “*the intruder was airborne*” appears difficult to reconcile with operational conditions: a conventional aircraft or helicopter flying at 2300 ft within 1 nm would have been visually observable given the reported visibility. Although a high-speed drone tracking aircraft at 80 m/s could theoretically produce such patterns and remain difficult to detect visually, several factors argue against this hypothesis. A drone would typically not carry TCAS equipment. The consistently reported altitude across all encounters suggests either exceptional altitude-holding capability or a fixed transmission parameter, which would be unjustified for a genuine platform.

Mode C is unaddressed, and the intruder, reported to be above the affected aircraft, was never seen by ATC, unlike the other traffic involved.

These factors, combined with the payload capacity, energy requirements, and precise tracking capability needed to maintain such encounters from a moving platform, strongly favor the explanation of a stationary transmitter, or “*a ground-based transmitter*”, employing range spoofing techniques.

³RPA4549 is not shown, as its transponder signaled a range estimation fault.

This hypothesis is also consistent with the interpretation of the official letter of the T&I (§ III), suggesting a c-UAS system near restricted airspace. In principle, such systems could implement area-denial strategies by injecting non-existent traffic from a ground-based transmitter into TCAS and triggering alerts, thereby discouraging airborne platforms from entering protected zones without physical engagement.

In § IV, we provide technical evidence that supports a ground-based transmitter as the root cause. We demonstrate that a ground-based attacker can exploit Mode C via range spoofing to inject an airborne intruder, resulting in effects compatible with the DCA events.

C. Threat Model

Considering the plausibility of a cyber-physical attack, we adopt the following threat model.

Actor capabilities. The actor is capable of designing or replicating attacks targeting standard-compliant ACAS/TCAS II. The actor is assumed to be capable of transmitting and receiving radio signals compatible with aviation surveillance protocols, using Commercial Off-The-Shelf (COTS) hardware, including SDRs, coupled with an amplifier and an antenna capable of transmitting on 1030/1090 MHz. Effective transmission requires an operational context that provides an unobstructed radio line-of-sight to aircraft along their flight path. This condition can be satisfied by positioning near low-altitude arrival or departure corridors, which follow predictable and well-documented procedures [39]. Civil aviation authorities publish these routes or can be inferred from public ADS-B data available on online platforms [40]. Moreover, to avoid detection by ground-based receivers (e.g., ATC radars), the actor may employ directional antennas or select a location that minimizes signal propagation toward those sensors.

Actor goals. The purpose of the malicious actor is to inject fabricated aircraft into the victim’s TCAS display. Beyond

triggering abrupt avoidance maneuvers, this compromises pilot situational awareness and trust in the collision avoidance system, potentially leading to incorrect or delayed decisions. This disruption can result in missed approaches, saturated ATC communications, or temporary airspace closures, ultimately degrading overall flight safety and causing significant operational and economic impact. Following the classification by Lykou et al. [41], we consider cybercriminals, cyber-terrorists, and nation-state actors as relevant threat agents. Cybercriminals are motivated by financial gain; cyberterrorists aim to cause large-scale disruption and erode public trust in aviation systems; nation-state actors may develop such capabilities for strategic or military purposes.

IV. ATTACK DESCRIPTION

A. Reference Practical Attack

Longo et al. [16] presented the first practical RF-based attack on TCAS, exploiting Mode S to inject false aircraft and induce TAs and TAs under controlled conditions. The attack relied on precise emulation of Mode S transponder behavior. The attacker generated spontaneous squitters to initiate surveillance, responded to incoming TCAS interrogations with valid replies, and maintained stateful interactions over time. The attack required consistent bidirectional communication over both Mode S frequencies, strict compliance with the TCAS timing, and complete handling of the surveillance logic. To meet the round-trip delay constraint required for range spoofing, the system achieved sub-128 μ s reply latency using optimized signal processing and real-time operating system configurations. This capability enabled dynamic control of the reported altitude and range of the spoofed aircraft, making it possible to meet the internal thresholds required to trigger TCAS alerts. The attacker could simulate both unilateral and coordinated encounters, using additional TCAS-specific message types to participate in RA negotiation sequences.

B. Attack for the DCA Scenario

The events at DCA point to an attack vector that exploits the legacy compatibility with Mode C transponder targets that are not TCAS-equipped. The underlying principles remain consistent with the Mode S attack. Because TCAS RAs depend on relative range, range-rate, and vertical separation, and can succeed at any bearing, the attack only manipulates altitude and range via malicious transponder replies.

Executing a similar attack against a Mode C target requires an adversary to solve a more demanding variant of the range spoofing problem, as presented in the following. A Mode A/C transponder replies 3 μ s after receiving the final pulse (P_3) of an interrogation (§ II-C). Increasing the perceived range is trivial, as it only requires the attacker to introduce an arbitrary delay longer than that of a genuine transponder. Conversely, to reduce the perceived range and make a ghost aircraft appear closer than the attackers' physical location, their reply must be transmitted earlier than a legitimate transponder. Specifically,

if the attacker replies τ seconds earlier than the standard processing delay, the spoofed range ρ_s becomes:

$$\rho_s = \frac{1}{2} \cdot c \cdot (\Delta T - 3\mu s - \tau) = \rho - \frac{1}{2} \cdot c \cdot \tau \quad (2)$$

This relationship, derived from Eq. 1, shows that the attacker can reduce the perceived range by $\frac{c \cdot \tau}{2}$, i.e., approximately 150 meters per microsecond in τ . To maximize this reduction, the spoofed reply must be sent as soon as possible. We describe three strategies to achieve this, with increasing levels of aggressiveness, hereafter referred to as S1, S2, and S3.

S1 - Immediate reply. This strategy represents the most conservative approach. In the context of Mode A/C interrogations and replies (Figure 6), an attacker responds immediately after receiving P_3 , violating the standard 3 μ s delay. This reduces the perceived range by approximately 500 meters (Eq. 2).

S2 - Mode discrimination. To achieve a greater reduction in perceived range, the attacker must transmit a forged reply before receiving the P_3 pulse, as illustrated in Figure 7. This strategy requires predicting the type of interrogation, which can be inferred by exploiting the timing difference between the Mode A and Mode C interrogations. Both start with a P_1 pulse, but in Mode A, the final P_3 pulse follows after 8 μ s, whereas in Mode C it arrives after 21 μ s. An attacker can listen for P_1 , wait slightly beyond the 8 μ s mark, and if no P_3 is observed, confidently conclude that the interrogation is Mode C. With this information, the attacker can transmit a forged reply immediately, without waiting for the actual P_3 at 21 μ s. This strategy anticipates the standard reply by an additional 13 μ s compared to S1, yielding $\tau = 16\mu$ s and enabling a range reduction of nearly 2.4 km.

S3 - Preemptive reply. A third, even more aggressive strategy, illustrated in Figure 8, involves replying immediately after detecting the first P_1 pulse, without waiting for P_2 or P_3 . This maximizes the spoofed range to over 3.5 kilometers but prevents the attacker from distinguishing between Mode A and Mode C interrogations. Since both modes use the same reply format, a single transmission is interpreted as a valid response to both, causing the victim system to display two targets: one showing an undefined squawk code (Mode A) and another at the spoofed altitude (Mode C), each at a different range.

C. Experimental Evaluation

To validate the hypothesis of a cyber-physical attack in the DCA scenario, we focused on the Mode Discrimination spoofing technique (S2) for implementation and experimental evaluation. Strategy S2 is more effective than S1 and is the only one consistent with the incident. It provides a significant range reduction without generating multiple tracks, unlike S3, and matches the single-target DCA scenario.

We organized our experimental session into three phases by evaluating (i) hardware latency, (ii) protocol compliance, and (iii) range spoofing capabilities.

Hardware latency. This phase tested whether the SDR setup from the Mode S attack could meet Mode C tighter timing constraints. We replicated a similar configuration: an Ettus

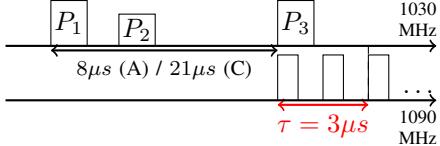


Fig. 6: Immediate reply (S1).

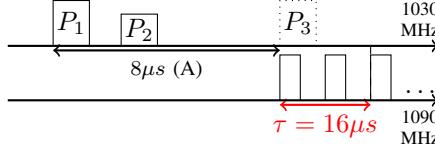


Fig. 7: Mode discrimination (S2).

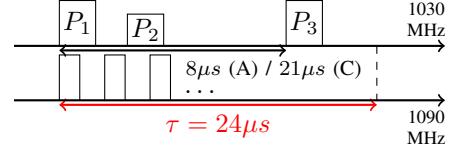


Fig. 8: Preemptive reply (S3).

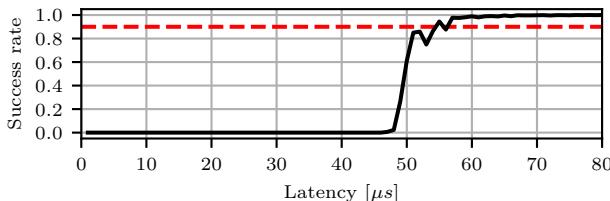


Fig. 9: USRP X300 latency test result.

USRP X300 SDR connected via PCIe to a PC (Intel i9-14900, 5.8 GHz, 32 GiB RAM) running real-time Ubuntu 24.04 LTS.

We conducted the tests using the SDR built-in latency test, which executes a no-op command to verify that data can be successfully looped back from the SDR at a given latency. In particular, we tested the range of latencies from $1\mu\text{s}$ up to $80\mu\text{s}$ at $1\mu\text{s}$ steps, recording for each test how many packets met the requested latency. Figure 9 presents the cumulative distribution function of the latency test results. A 90% reply rate, as required by the standards [23, §3.1.1.7.1], is achievable with $\tau = 52\mu\text{s}$ and $\rho_s - \rho = 7.79$ km. This spoofed range places the ghost aircraft nearly 8 km farther than intended. Notably, even the optimized software pipeline from the original Mode S implementation achieved zero reply rate beyond the $50\mu\text{s}$ mark, confirming that SDR-based platforms cannot meet the Mode C timing requirements. To meet this limitation, we use an FPGA with integrated RF transceivers (AMD Zynq UltraScale+ RFSoC, 92.16 MHz), programmed with Vivado 2024.2.

We implemented a Mode A/C transponder compliant with [23, §3.1.1], excluding requirements 3.1.1.{4,5,8,9,10} (ground station), and 3.1.1.7.11 (output power). It also includes attacker logic that replies to interrogations with attacker-controlled altitude, squawk, and τ (see § IV-B).

TABLE II: FPGA resource utilization summary.

Resource	Utilization	Available	Utilization %
Lookup Tables (LUT)	19145	425280	4.50
LUT-based RAM	706	213600	0.33
Flip-Flops	17801	850560	2.09
Block RAM	8	1080	0.74
Digital Signal Processors	3	4272	0.07
Input/Output Pins	10	408	2.45

Table II summarizes the FPGA resources consumed by our implementation, compared to the total available on the device.

Protocol compliance. In this phase, we test if our implementation is standard-compliant. We set up an RF testbed inside

a shielded enclosure containing five antennas: one RX/TX antenna for a certified Aeroflex IFR 6000 tester, two antennas for RX/TX of the FPGA, and two for recording tests with an Ettus USRP X300 SDR. The enclosure is required, since over-the-air transmissions are prohibited. While this controlled setup limits realism, we observed no engineering constraints that would prevent addressing potential performance issues in real-world conditions. We configured our implementation to respond with a squawk code of 1234, an altitude of 4000 ft, and disabled range spoofing ($\tau = 0$ in Eq. 2).

TABLE III: Aeroflex IFR6000 Test Results Summary.

Test	Metric	Measured Value	Res.
A/C Decoder & SLS	Decoded Squawk	1234	Pass
	Decoded Altitude	4000 ft	Pass
	Reply Delay (A/C)	$2.95\mu\text{s} / 2.95\mu\text{s}$	Pass
	Reply Jitter (A/C)	$0.010\mu\text{s} / 0.009\mu\text{s}$	Pass
	Reply Ratio (A/C)	100% / 100%	Pass
	-81 dBm Reply Ratio (A/C)	0% / 0%	Pass
A/C Spacing & Width	ATCRBS ALL-CALL (A/C)	—	Pass
	Pulse Amp Var. (A/C)	0.7 dB / 0.8 dB	Pass
A/C Duration & Amp	Max. Pulse Amp Var. (A/C)	0.7 dB / 0.8 dB	Pass
	Power & Frequency	Transmit Frequency A-C Diff. Effective Radiated Power	Pass Pass Fail
		1089.99 MHz -0.5 dB 0 W	

Table III lists test group, metric, measured value, and pass/fail result against standard limits. Appendix C shows the full set of Aeroflex test screens. Briefly, the decoder and Side Lobe Suppression (SLS) test confirms that the tester correctly decodes the configured squawk code and altitude. In addition, they include the A/C Spacing and Width test, which verifies the structure and timing of reply pulses; the Mode A/C Duration and Amplitude test, which checks the consistency and strength of the pulses; and the Power & Frequency test, which measures the carrier frequency and the effective radiated power. All values meet performance standards, except for transmitted power, which was low due to the absence of an external amplifier in the test setup.

Range spoofing evaluation. In this phase, we evaluated the S2 range spoofing technique using our FPGA implementation. The experiment involved the same SDR setup as before transmitting 3079 Mode A and 3022 Mode C interrogations to the FPGA device while recording reply timing. As detailed in Table IV, our implementation achieved a mean τ of $12.85\mu\text{s}$, with an overall performance between 78.7% to 82.5% of the ideal attack performance $\hat{\tau}$, and with correct discrimination between mode A and C interrogations. Figure 10 shows the CDF of the achieved spoofed range and its precision, with the

TABLE IV: Range spoofing results.

Metric	μ	σ	min	max
τ [μs]	12.85	0.122	12.6	13.2
$\tau \div \hat{\tau}$ [%]	80.3	—	78.7	82.5
$\rho_s - \rho$ [m]	-1925.513	18.342	-1972.634	-1888.692
Reply rate (A/C) [%]	0.0/100.0	-	0.0/100.0	0.0/100.0

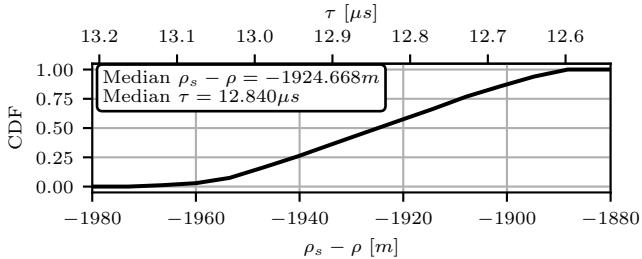


Fig. 10: CDF of range spoofing in our implementation.

entire distribution confined to less than 50 meters around the median.

D. Discussion

Our implementation, validated with a certified avionics tester, can emit replies that appear at zero range from a distance of approximately 2 km. At DCA, RAs were triggered at a median distance of 370 m. An attacker could realistically operate from more than 2 km away, causing short-range alerts by exploiting line-of-sight to approaching aircraft and without requiring physical access to the airspace, as observed during the DCA incident. Compared to previously demonstrated Mode S attacks, this variant is less capable in terms of spoofed range and cannot target specific aircraft with selective interrogation (§ II-D). However, placing the transmitter near a strategic point, such as the final approach corridor of DCA, makes the novel attack equally effective, although it lacks the above capability. A key advantage of the Mode C variant is its simpler modulation and protocol structure, which makes it easier to implement. Although FPGAs are typically more complicated to manage than SDRs, the low resource usage of our design, with logic usage peaking at 4.5% and all other categories well below 3%, indicates a modest deployment effort. The FPGA used in our implementation is a high-end platform with a cost of approximately \$25,000, compared to \$10,000 reported for Mode S attacks. However, given the low resource usage, this cost difference is unlikely to be operationally significant. Intuitively, the lightweight nature of the design also suggests that it could be ported to a substantially cheaper FPGA.

V. DETECTION METHODOLOGY

The DCA incident illustrates a key limitation in the current surveillance architecture. Although ATC is trained and equipped to handle airborne emergencies, it lacks the technical means to assess whether anomalies such as phantom TCAS

alerts are caused by system malfunctions or deliberate interference. This diagnostic gap persists despite the CISA advisory on the potential for spoofing and other cyber threats. In the DCA case, the anomaly was diagnosed through spectrum analysis, looking for the presence of a fixed transmitter located in a position inconsistent with any expected activity of the aircraft. However, this detection solution requires time, specialized equipment, and dedicated personnel, and building such infrastructure is not feasible on short timescales. In this context, we argue that ATC already has access to the data needed to determine whether a single transmitter is responsible and to estimate its location. TCAS resolution advisories include structured threat descriptors such as TIDB and TIDR, which encode relative bearing and range. Voice reports provide a complementary data source. Pilots are required to report RAs, and these communications often include additional cues such as headings, relative position, and observed intruder behavior.

However, this data presents considerable challenges. These data sources, while available, come with significant challenges. Information is indirect, ephemeral, and noisy. Bearing information has low precision, range measurements are subject to spoofing, and the alerts themselves only occur for brief moments when an aircraft passes through the attacker's operational area. Pilot reports are often approximate.

Under these conditions, the aim is to design a framework that can process this sparse and unreliable data to infer whether a single non-cooperative transmitter is present and estimate its location. The goal is to fuse fragmented in-flight observations over time into a coherent and actionable position estimate, enabling both post-incident forensics and near-real-time response.

Localization as a Sequential Monte Carlo (SMC) task. We formulate the attacker localization problem as an SMC task, employing a particle filter to estimate the posterior distribution of the attacker's position. This allows us to represent our belief about the attacker's state as a large set of weighted hypotheses (particles) and to iteratively refine these beliefs as each piece of evidence from the incident timeline is processed. Over successive iterations, particles in regions of the state space that are inconsistent with the observations will see their weights diminish. In contrast, particles in regions that better explain the evidence will be assigned higher weights. The final output of this process is a geospatial probability map, which provides a quantitative and actionable estimate of the attacker's location, suitable for directing investigative resources. Importantly, the method exhibits graceful failure characteristics. When the underlying assumptions, such as the presence of a single stationary transmitter, are invalidated, the filter fails to converge, producing a diffuse probability map. Rather than yielding incorrect results, this behavior provides an explicit signal of model mismatch and limits the risk of misleading conclusions.

A. Assumptions

To constrain the problem space and develop a tractable model, we make the following simplifying assumptions: (i)

We model the attacker as a stationary transmitter, consistent with an adversary positioned at a fixed ground location to intercept aircraft on an approach path. (ii) All anomalous TCAS events originate from a single geographical point. (iii) All encounters are assumed to be with the same malicious actor. (iv) All non-attacker equipment, including victim aircraft navigation systems and TCAS hardware, functions correctly and according to established standards. The anomalies are attributed solely to malicious external transmissions.

B. State representation and inputs

The state of our system is represented by a set of N particles, where each particle represents a hypothesis about the location of the attacker. We define the i -th particle state as $\mathbf{x}_i = \langle x, y \rangle$, representing the attacker's position in a 2D Cartesian coordinate system. Each particle carries an associated weight w_i , indicating the likelihood that the position of that particle is the actual location of the attacker, with normalized weights such that $\sum_{i=1}^N w_i = 1$.

The dataset for the technique consists of information found within the ADS-B data: the periodic position broadcasts from the affected aircraft and the ACAS RA reports transmitted during the encounters. In particular: (i) The victim aircraft's own periodically broadcast state vectors, which provide a time-stamped history of its latitude ϕ , longitude λ , and track angle h . (ii) The ACAS RA broadcasts triggered during an encounter. From these, we extract from TIDB and TIDR the reported bearing intervals $\theta_{min}, \theta_{max}$ and range ρ_{min}, ρ_{max} of the ghost aircraft.

To perform the necessary Euclidean distance and bearing computations, we project all geographical coordinates (ϕ, λ) into a 2D Cartesian space (x, y) . We employ a transverse Mercator projection centered on the area of interest [42, p.174]. This creates a local flat-earth tangent plane in which distances and bearings can be accurately calculated without resorting to computationally intensive geodesic calculations [43, p.48]. Thus, each observation in the dataset z_k is represented as a tuple $\langle id, x, y, h, \theta_{min}, \theta_{max}, \rho_{min}, \rho_{max} \rangle$, where id is the unique identifier of the aircraft.

C. Workflow overview

Our methodology Sequential Monte Carlo for Rogue ACAS Transmitters (SMC-RAT), operates through a structured sequence of stages that progressively refine hypotheses about the transmitter's location. This procedure is described in Algorithm 1. The algorithm begins with the input processing stage (§ V-D). Line 1 initializes an empty set for the filtered observations. The procedure then iterates through each aircraft identified in the dataset (Line 2), isolates its associated observations, and applies a per-aircraft consistency check to retain only inlier measurements (Line 3). After processing all aircraft, a global inlier filter is applied to the aggregated set of measurements (Line 5).

Next, the filter initialization stage (Lines 6–8) prepares the particle set. The algorithm samples an initial position for each

Algorithm 1 SMC-RAT

```

Input: Observations  $z$ , number of particles  $N$ 
Output: Particle positions  $\mathbf{x}$  and weights  $w$ 
1:  $Z_{filt} \leftarrow \emptyset$                                  $\triangleright$  Input processing
2: for each aircraft  $a \in \{z.id \mid z\}$  do
3:    $Z_{filt} \leftarrow Z_{filt} \cup \text{KEEPINLIERS}(z|_{id=a})$ 
4: end for
5:  $z \leftarrow \text{KEEPINLIERS}(Z_{filt})$ 
6: for  $i \in 1, \dots, N$  do                          $\triangleright$  Filter initialization
7:    $\mathbf{x}_i \leftarrow \text{INITPOSITION}(z_1)$ 
8:    $w_i \leftarrow \frac{1}{N}$ 
9: end for
10: for  $j \in 2, \dots, |z|$  do       $\triangleright$  Motion-update-resample loop
11:   for  $i \in 1, \dots, N$  do
12:      $\mathbf{x}_i \leftarrow \text{MOTIONMODEL}(\mathbf{x}_i)$ 
13:      $w_i \leftarrow \text{UPDATEMODEL}(\mathbf{x}_i, w_i, z_j)$ 
14:   end for
15:    $W \leftarrow \sum_{i=1}^N w_i$                        $\triangleright$  Normalization
16:   for  $i \in 1, \dots, N$   $w_i \leftarrow \frac{w_i}{W}$  end for
17:   if NEEDSRESAMPLING( $w$ ) then            $\triangleright$  Resampling
18:     for  $i \in 1, \dots, N$  do
19:        $\mathbf{x}_i \leftarrow \text{RESAMPLE}(\mathbf{x}_i, w_i)$ 
20:        $w_i \leftarrow \frac{1}{N}$ 
21:        $\mathbf{x}_i \leftarrow \text{MOTIONMODEL}(\mathbf{x}_i)$ 
22:     end for
23:   end if
24: end for

```

particle (Line 7) as detailed in § V-E. Then, each particle is assigned a uniform initial weight (line 8).

The core of the methodology is the motion-update-resample loop (Lines 10–21), which iteratively refines the location estimate by processing each subsequent observation. For every particle, the algorithm first applies a motion model to its position (Line 12; see § V-F) and then updates its weight based on how well this hypothesized position explains the current measurement (Line 13; see § V-G). After updating all particles for a given observation, their weights are normalized (Lines 15–16). The algorithm then checks if resampling is needed (Line 17; see § V-H). If the particle weights become too uneven, a resampling step is performed (Line 19). After resampling, the motion model is applied again to the particles (Line 21).

This entire process repeats for all filtered observations, culminating in a weighted particle distribution that represents a probabilistic estimate of the attacker's location, whose usage we describe in § V-I.

D. Input processing

Raw ACAS RA data is noisy, especially the bearing information. To improve data quality, we applied a two-stage filtering process to remove outlier measurements.

First, we perform a consistency check on a per-aircraft basis. For each aircraft that experienced RAs, we analyze the set

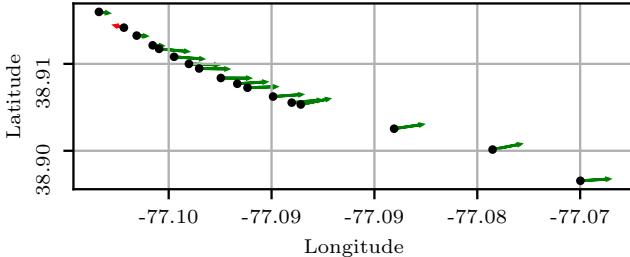


Fig. 11: Outlier filtering applied to RPA4469 ($\kappa = 0.5$).

of true bearing versors $\{\mathbf{b}_1, \dots, \mathbf{b}_M\}$ associated with its M encounters.

$$\mathbf{b}_k = [\cos \theta_c \quad \sin \theta_c]^T \quad \theta_c = (\hat{\theta}_c + h) \bmod 360^\circ \quad (3)$$

We calculate \mathbf{b}_k as in Eq. 3, by taking the center $\hat{\theta}_c$ of the relative bearings θ_{min} and θ_{max} , and adding the angle of the aircraft track h to obtain the true angle of the bearing θ_c .

$$\sum_{i=1}^M s_{ij} > \kappa \cdot M \rightarrow \mathbf{B}_j \text{ inlier} \quad s_{ij} = \frac{\mathbf{b}_i \cdot \mathbf{b}_j}{\|\mathbf{b}_i\| \cdot \|\mathbf{b}_j\|} \quad (4)$$

As detailed in Eq. 4, a verson \mathbf{b}_j is considered an inlier if the sum of its cosine similarities s_{ij} with all other versors (including itself) exceeds a threshold $\kappa \in (0, 1)$ proportional to the number of encounters M , with κ setting the required level of agreement.

As an example, Figure 11 illustrates this process for the RPA4469 flight. Most data points, marked as inliers, point towards the east. In contrast, the single outlier, depicted in red, points in the opposite direction.

In the second stage, the same filtering logic is applied to the entire collection of inlier versors identified in the first step. We consolidate all measurements that passed the per-aircraft check into a single set \mathbf{B}_{filt} . The measurement corresponding to the verson $\mathbf{b}_j \in \mathbf{B}_{filt}$ is retained in the final dataset if it satisfies the same condition from Eq. 4, but evaluated over the complete set \mathbf{B}_{filt} .

E. Filter initialization

For particle initialization, we leverage bearing and range data from the first valid encounter.

$$\theta_i^{(1)} = [\mathcal{U}(\theta_{min}, \theta_{max}) + h] \bmod 360^\circ \quad (5)$$

$$\rho_i^{(1)} \sim \mathcal{U}(\rho_{min}, \rho_{max} + \frac{c \cdot \tau_{max}}{2}) \quad (6)$$

$$\mathbf{x}_i^{(1)} = [x + \rho_i \cos \theta_i \quad y + \rho_i \sin \theta_i]^T \quad (7)$$

Each particle initial position $\mathbf{x}_i^{(1)}$ is sampled as in Eq. 7, adding to the aircraft initial position a polar offset $\rho_i^{(1)}, \theta_i^{(1)}$. $\theta_i^{(1)}$ is computed in Eq. 5 similar to θ_c in Eq. 3, with the initial relative bearing being sampled from a uniform distribution \mathcal{U} between θ_{min} and θ_{max} ⁴. Eq. 6 calculates $\rho_i^{(1)}$, drawn from

⁴All angular calculations implicitly handle the wrap-around at 360° throughout this article.

the range between ρ_{min} and ρ_{max} plus the maximum attacker range spoof capability τ_{max} .

F. Motion model

Given our assumption of a stationary attacker, the motion model primarily serves to maintain particle diversity and explore the solution space rather than tracking the actual target motion. We implement a simple Gaussian perturbation model in which the position of each particle \mathbf{x}_i evolves according to the following:

$$\mathbf{x}_i^{(k)} = \mathbf{x}_i^{(k-1)} + \eta \quad (8)$$

where $\eta \sim \mathcal{N}(0, \sigma)$ represents isotropic 2D Gaussian noise with standard deviation σ .

G. Measurement model

The measurement model forms the core of the particle weight update, evaluating how well each hypothesized attacker position explains the observed data.

$$w_i^{(k)} = w_i^{(k-1)} \cdot \mathcal{L}_\theta(\mathbf{x}_i, z_k) \cdot \mathcal{L}_\rho(\mathbf{x}_i, z_k) \quad (9)$$

For each particle i and observation z_k , we update its likelihood at step k as shown in Eq. 9, where \mathcal{L}_θ and \mathcal{L}_ρ represent the bearing and range likelihood components, respectively.

$$\mathcal{L}_\theta = \begin{cases} 1.0 & \text{if } |\delta_\theta| \leq 6^\circ \\ \exp \left[-\frac{1}{2} \left(\frac{\delta_\theta}{\sigma_\theta} \right)^2 \right] & \text{if } 6^\circ < |\delta_\theta| \leq 27^\circ \\ \epsilon & \text{otherwise} \end{cases} \quad (10)$$

The bearing likelihood \mathcal{L}_θ is calculated as in Eq. 10, with δ_θ being the bearing error between θ_c (see Eq. 3) and the direction to the particle, which is determined by $\text{atan2}(\mathbf{x}_{i_y} - y, \mathbf{x}_{i_x} - x)$. Here, TIDR's inherent 6-degree quantization is assigned the maximum likelihood, while values lying within the maximum allowable bearing error of 27 degrees are subject to exponential decay⁵. Otherwise, a small floor likelihood ϵ is assigned to ensure numerical stability.

$$\mathcal{L}_\rho = \begin{cases} 1.0 & \text{if } \delta_\rho \in [\rho_{min}, \rho_{max} + \frac{c \cdot \tau_{max}}{2}] \\ \epsilon & \text{otherwise} \end{cases} \quad (11)$$

The range likelihood component \mathcal{L}_ρ , defined in Eq. 11, is considered fully consistent (likelihood of 1.0) if the distance to the particle $\delta_\rho = \|\mathbf{x}_i - \langle x, y \rangle\|$ lies between ρ_{min} and ρ_{max} indicated within the TIDR field, with the latter being augmented by the maximum attacker range spoof capability.

H. Resampling

The particle filter requires resampling when the distribution of particle weights becomes highly uneven, indicating that most particles have negligible importance. We quantify this degeneracy using the Effective Sample Size (ESS), defined as:

$$\text{ESS} = \frac{1}{\sum_{i=1}^N (w_i^{(k)})^2} \quad (12)$$

⁵ $\sigma_\theta \approx 9$

When ESS falls below a threshold N/κ_{ESS} , the filter performs resampling to redistribute particles to regions of higher probability. This prevents the filter from degenerating to a state in which only a few particles carry significant weight [44]. We employ stratified resampling: the algorithm divides the cumulative distribution function of particle weights into N equal strata and draws one sample from each stratum [45], [46]. The set of resampled particles $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ then replaces the original particles, with all weights reset to uniform values $w_i = 1/N$. Following resampling, we apply jittering to restore particle diversity and prevent sample impoverishment, that is, each resampled particle undergoes perturbation in Eq. 8.

I. Output

The particle filter produces a probabilistic representation of the attacker's location. The primary output is a probability density estimate over the 2D position space, constructed through Gaussian kernel density estimation.

$$\hat{p}(\mathbf{x}_q) = \sum_{i=1}^N w_i \cdot \mathcal{K}(\mathbf{x}_q - \mathbf{x}_i) \quad (13)$$

Eq. 13 presents how to perform a likelihood query for a given point \mathbf{x}_q , with \mathcal{K} representing a Gaussian kernel. From this density, we compute the most likely attacker position $\hat{\mathbf{x}}$ (centroid) as:

$$\hat{\mathbf{x}} = \sum_{i=1}^N w_i \cdot \mathbf{x}_i \quad (14)$$

To quantify the uncertainty of this estimate, we compute a confidence ellipse for a given confidence percentage $\kappa\%$. This begins with calculating the weighted covariance matrix e_{Σ} of the particle distribution. We then find a scaling factor e_s using the inverse cumulative distribution function of the χ^2 distribution with two degrees of freedom, such that $e_s = \chi_{\text{ppf}}^2(\kappa\%, 2)$. An eigendecomposition of e_{Σ} yields its eigenvalues, $e_{\lambda_1}, e_{\lambda_2}$, and corresponding eigenvectors. The lengths of the semi-major and semi-minor axes e_a, e_b are then $e_a, e_b = \sqrt{e_s \cdot \lambda_{\max}}, \sqrt{e_s \cdot \lambda_{\min}}$ [47]. The orientation of the ellipse e_{ϕ} is the angle of the eigenvector associated with the major axis [48].

Finally, to evaluate specific locations of interest, such as known restricted areas, we implement a circle likelihood query. This computes the probability that the attacker lies within a circular region of radius r centered at $\langle x_c, y_c \rangle$.

$$\tilde{p}(r, x_c, y_c) = \frac{1}{\pi r^2} \iint_{||\langle x, y \rangle - \langle x_c, y_c \rangle|| \leq r} \hat{p}(\langle x, y \rangle) dx dy \quad (15)$$

Eq. 15 provides the exact continuous formulation for this query, that is, the integration of Eq. 13 over the desired circle.

$$p^*(r, x_c, y_c) = \frac{\tilde{p}(r, x_c, y_c)}{\tilde{p}(r, \hat{\mathbf{x}}_x, \hat{\mathbf{x}}_y)} \quad (16)$$

An absolute density value can be difficult to interpret on its own, so we normalize it to create a more intuitive relative score. This normalized likelihood, defined in Eq. 16, compares

the average density in the query circle to that of an identical one centered on the location of the maximum probability.

$$\tilde{p}(r, x_c, y_c) \approx \frac{1}{\kappa_m} \sum_{i=1}^{\kappa_m} \hat{p} \left(\langle x_c, y_c \rangle + r \cdot [\cos \theta_i \quad \sin \theta_i]^T \right) \quad (17)$$

Rather than computing \tilde{p} analytically, we approximate its value using Monte Carlo integration [49, §7.7]. Eq. 17 defines this procedure, with each θ_i drawn from $\mathcal{U}(0, 2\pi)$, and κ_m denotes the number of sample points.

Experimental evaluation. We validate the SMC-RAT methodology through a two-part experimental evaluation. First, we assess its accuracy and convergence properties using a large set of controlled simulations. Second, we apply the methodology to the real-world data from the DCA incident to demonstrate its practical utility in a forensic context.

J. Simulations

We evaluated our methodology by conducting a series of experiments on simulated data. The goal is to assess the localization accuracy of our approach under different attack conditions corresponding to the three range-spoofing strategies (S1, S2, and S3) detailed in § IV. We execute 300 thousand simulation trials, varying the attacker's τ_{\max} , with values of $3\mu\text{s}$, $16\mu\text{s}$, and $24\mu\text{s}$. In each trial, we generate a random scenario with a single stationary attacker performing a range-spoofing attack that targets a reported distance of zero from the victim and between two and ten aircraft encounters. We require at least two aircraft encounters to rule out chance coincidences. For each encounter, we simulate an aircraft with a random initial position, heading, and a constant speed drawn uniformly from 50 to 100 m/s. The aircraft follows a linear trajectory, generating between three and ten TCAS RA events at random intervals of between one and five seconds. To model realistic uncertainty, we added errors to the quantized TIDB and TIDR fields following ICAO standards and our dataset. For range, we apply a Gaussian error ($\sigma = 10$), modeling a worst-case system where $\geq 90\%$ of ranges fall within the 14.5m ICAO-specified resolution [23, §4.3.2.1.3.1]. For bearing, we modeled noise with a Gaussian error ($\sigma = 9$), a similarly conservative model where roughly 85% of estimates lie within the 10-degree RMS ICAO recommendation [23, §4.3.2.1.3.2]. We also introduced a 10% probability of catastrophic bearing failure (replacing TIDB with a random value), a conservative rate, as our DCA data analysis showed a $\approx 33\%$ bearing outlier rate. We apply our localization algorithm using the following hyperparameters across all simulations: the particle filter population is $N = 1000$, its motion model σ is 25, the inlier consensus threshold κ is 0.5, and the bearing likelihood standard deviation σ_{θ} is 9. We trigger resampling when the ESS drops below half of the total particle population, i.e., $\kappa_{ESS} = 2$. Our performance metric is the localization error, defined as the Euclidean distance between the ground-truth attacker position and $\hat{\mathbf{x}}$, as estimated by our filter.

Figure 12 depicts the CDF of these localization errors for each of the three attack scenarios. For S1, the methodology

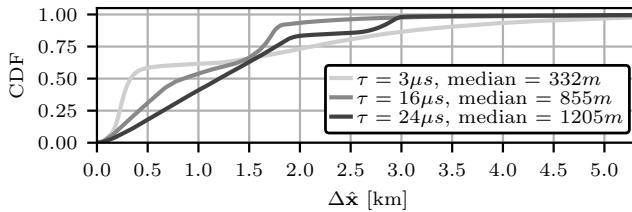


Fig. 12: CDF of localization accuracy in our detection.

achieves the highest precision with a median localization error of 332 m. Strategy S2 shows an increase in error with a median of 855 m. S3 exhibits the most significant localization errors with a median of 1205 m.

Since the method's assumptions cannot be proven in advance for an unknown case, we assessed whether their violation produces evident, detectable failures. We conducted 200 thousand additional simulations with mobile attackers moving at speeds between 1-100 m/s with inter-encounter times ranging from 10 to 180 s. For each scenario, we computed the 33% confidence ellipsoid and measured its area. Compared to the stationary case, whose median area was 5.16 km², the moving scenarios had a median area of 155.21 km².

We performed 100 thousand additional trials on the testbed workstation to assess computational performance. The SMC-RAT algorithm achieved a per-trial average execution time of 7.4s, 99th percentile of 10.7s, and a maximum of 17.6s.

K. Application to DCA incident

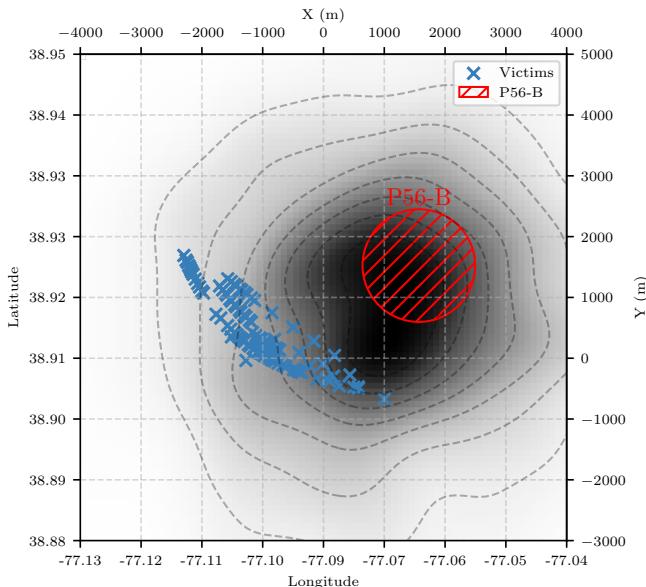


Fig. 13: Plausible attacker locations for the DCA incident.

Figure 13 depicts the final geospatial probability distribution for the attacker's location. The map highlights a concentrated area of high probability. The centroid of this distribution

(Eq. 14), representing the single most likely location, is located approximately 891.7 meters from the center of the restricted airspace of P-56B.

To further evaluate P-56B as a potential source, we calculate its relative likelihood using the procedure in Eq. 17 using 1000 as κ_m . This analysis indicates that the probability of the transmitter being located within the P-56B area is 94.1% relative to the most probable location identified by our filter.

L. Discussion

Detection performance and limitations. The SMC-RAT methodology enables operators to determine whether an anomalous transmission event plausibly originates from a single rogue transmitter. This is assessed by inspecting the size and shape of the confidence region. When the underlying assumptions are violated, such as the presence of a mobile rather than a stationary transmitter, the ellipse area increases by orders of magnitude, providing a quantitative signal of model mismatch.

Once validated, localization achieves sub-kilometer precision for practical attacks, significantly reducing the search space compared to the full anomaly extent. This computation takes a few seconds, enabling operators to use SMC-RAT both to refine estimates as an incident unfolds and for post-incident forensic analysis. However, the achievable accuracy remains bounded by limitations in the underlying data, particularly the inherent unreliability of the TCAS bearing estimates and the attacker's ability to spoof range. To mitigate the impact of bearing noise, we applied a filtering step to discard outliers, although the resolution of the TCAS ultimately restricts the accuracy. For range spoofing, the localization errors in our simulations scale with the used τ . To model a worst-case scenario, we conservatively set τ to its theoretical upper bound. However, our experimental validation of attacks shows that the real-world τ values are significantly lower due to processing delays, which in turn reduces the spoofing capability and improves the localization accuracy.

Another limitation lies in the availability of positional data required for automatic localization. In the TCAS protocol, the bearing and range data are only available during Mode C encounters. Mode S encounters provide only the intruder's identifier, without positional information. This allows spoofed Mode S targets to evade localization if the attacker ignores ATC interrogations and omits ADS-B transmission, both of which are optional for triggering a TCAS alert. ATC may still retrieve this information by requesting that pilots include the relevant threat data displayed on their cockpit instruments when reporting RA. In such scenarios, the effectiveness of our methodology depends on the systematic adoption of this reporting procedure by ATC, the manual entry of pilot-provided data into the system, and the accuracy with which pilots observe and communicate the TCAS parameters.

Recommendations for protocol enhancement. To ensure that critical threat data is available in the event of a cyberattack, and to improve the accuracy and automation of our localization method, we propose a minor modification to the ADS-

B standard [26]. First, to address the evasion potential of rogue Mode S targets, we propose the introduction of a new broadcast message for both RAs and TAs, which carries the threat data currently available only for Mode C encounters. This message would include the intruder’s range, bearing, and altitude, allowing for localization even when the target does not transmit ADS-B. Second, we recommend the addition of a dedicated *ACAS TA Broadcast Message*. Broadcasting threat parameters upon TAs, not only RAs, would increase both the frequency and timeliness of positional data, significantly improving downstream localization performance. The current standard reserves five unallocated message types for ACAS broadcasts, offering a backward-compatible and straightforward path to implementation.

DCA scenario. The FAA spectrum analysis identified the activation of USSS c-UAS in a DoD facility near the approach path to DCA (§ III). Continually, our independent analysis places the likely source of the anomalous transmissions near the restricted airspace P-56B. This area is subject to stringent protection policies, and the use of USSS-operated active interdiction systems (§ IV-B) is both legally allowed and technically feasible [50].

However, detection alone does not constitute attribution [51]. As an illustrative example, a nation-state actor might deliberately position the source of an attack near sensitive airspace, such as P-56B, to induce suspicion toward domestic government systems. This form of misdirection aligns with the threat model discussed in § III-C.

Whether the source was benign or malicious, timely localization would have been critical. Our method could have enabled operators to rapidly identify the transmission origin and take appropriate action. In the DCA case, after just two aircraft encounters, the particle filter had already constrained the estimated source area⁶ to approximately 4.3 km². Since the second encounter occurred within 40 minutes of the incident onset, actionable localization could have been achieved in under an hour. This stands in contrast to the three hours during which the anomaly persisted. In the case of a malicious actor, such rapid convergence would have substantially narrowed the search space for law enforcement.

VI. RELATED WORK

Protocol-level attacks. Several studies examining aviation cybersecurity and its foundational surveillance protocols [41], [52], [53], such as Mode A/C, Mode S, and ADS-B, highlight that these systems were designed without modern security mechanisms. This lack of authentication and encryption exposes them to a broad range of attacks, including false injection, and may help explain the feasibility of scenarios like the DCA incident. Strohmeier et al. [54] show that spoofed Mode S replies can create ghost targets, triggering RAs to which pilots must respond. They note that mitigation is complex since ATC cannot override RAs, a dynamic reflected both in the likely cause and the operational handling of the DCA incident.

Smith et al.[39] simulate collision avoidance attacks against CAS using an SDR-based adversary that transmits Mode S messages. They highlight conditions similar to those at DCA, particularly in stacks near busy airports, where aircraft follow tightly constrained vertical profiles. In such settings, a spoofed RA could lead to cascading alerts and ATC interventions, increasing the risk of level busts and system saturation. However, their analysis does not consider Mode A/C as an attack vector and does not assess real-world feasibility. Similarly, previous research [55], [56], [57], [58], [59] has investigated the use of SDR and forged Mode S messages to carry out false injection attacks against TCAS. Proof-of-concept implementations were developed within simulated environments. In particular, some works [55], [58] emphasize the challenges in achieving range spoofing capabilities required for practical attacks such as the DCA incident. Longo et al. [16] demonstrated a practical TCAS attack using Mode S replies, overcoming range constraints and showing that ghost aircraft can trigger alerts. In contrast, the novel method we study and implement in this paper, motivated by the DCA incident, relies on Mode C replies.

Mitigations. Mitigation strategies in surveillance protocols, mainly ADS-B, have been extensively surveyed and classified by Strohmeier et al. [60], and include approaches potentially extendable to TCAS. The first category relies on cryptography [61], [62], [63], [64] and the extension of the protocol [65]. However, such solutions are acknowledged to be challenging to deploy and susceptible to downgrade attacks in mixed-equipment environments [53], [16], [64], [66].

An alternative category relies on position verification, detecting anomalies between the claimed and estimated emitter locations. Applied techniques often rely on multilateration (MLAT), which estimates the location of the emitter through the time difference of arrival (TDOA) [67], [68]. A step beyond MLAT, Strohmeier et al. [60] propose a lightweight location verification system based on grid-based and statistical estimators and achieve localization accuracies around 145 m for ground-based ADS-B spoofers. As with all ground-based verification systems, coverage and accuracy remain inherently limited to areas where suitable receiver infrastructure is deployed and properly calibrated.

A third category uses physical-layer fingerprinting techniques to distinguish legitimate aircraft from spoofers based on unique signal characteristics [69], [70], [71]. These methods require specialized hardware and training, and their sensitivity to channel conditions limits real-world deployment. Our work resembles techniques in the localization approaches, but does not rely on any dedicated infrastructure, nor on physical-layer fingerprinting. The solution operates on aircraft state information that is either broadcast (as in the Mode C case at DCA), passively acquired, or derivable from standard radio communications, particularly in Mode S environments. While the current implementation assumes access to this data, full automation would only require a minor, non-substantial update to existing protocols, unlike approaches in the cryptography and protocol extension category.

⁶ $\kappa\% = 33\%$

VII. CONCLUSION

Using publicly available data on the DCA incident, this paper analyzes its root cause, demonstrates the plausibility of a cyberattack, and infers and validates a novel variant of the TCAS injection attack that has never been reported. Theoretical analysis bounds the maximum spoofing range to 3.5 km. Our implementation, validated using a certified avionics tester, achieves zero-range replies from nearly 2 km, enough to justify the DCA attack.

We also propose a detection methodology that provides comparable insight to spectrum analysis to identify transmitter locations, using only existing TCAS data without protocol updates or additional infrastructure. In the DCA case, our method localized the source within 4.3 km² in less than 40 minutes, allowing actionable localization well before the anomaly ended. In general, simulations show that it detects single-transmitter events and achieves 855 m median accuracy.

ETHICS CONSIDERATIONS

The general class of attacks against TCAS is publicly known and has been the subject of prior research and official advisories. However, the specific techniques and real-world evidence presented in this work constitute a significant evolution of the publicly understood threat. Following responsible disclosure principles, we have provided our detailed findings, including the technical analysis of the attack variant and our proposed detection methodology, to relevant national and international aviation authorities and cybersecurity agencies before publication. This includes briefing the U.S. Cybersecurity and Infrastructure Security Agency (CISA), which is updating its public TCAS advisory to address the vulnerabilities and operational risks described in this work. To ensure that our experimental validation did not pose a risk to operational airspace, we conducted all tests involving radio transmissions within a controlled laboratory environment using an RF-shielded enclosure. This practice prevented any leakage of signals into the open air and ensured full compliance with all applicable laws and telecommunication regulations. Finally, while this paper provides sufficient detail for the security community to understand and verify our findings, we deliberately withhold specific implementation details or source code that would lower the barrier to replicating the attack.

ACKNOWLEDGMENT

The authors wish to thank Baykar Piaggio Aerospace S.p.A. for their generous provision of access to specialized test instruments and valuable technical assistance during the experimental evaluation phase of this research. This work was partially funded by the NextGenerationEU project “Security and Rights in CyberSpace” (SERICS) and NATO Science for Peace and Security project n. G5633 “AI4CUAV”.

REFERENCES

- [1] European Commission, “Commission Regulation (EU) No 1332/2011 of 16 December 2011 laying down common airspace usage requirements and operating procedures for airborne collision avoidance (Text with EEA relevance),” Official Journal of the European Union, L 336, 20 December 2011, pp. 20–22, 2011. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2011/1332/oj>
- [2] F. A. Administration, “Ac 90-120 - operational use of airborne collision avoidance systems,” 11 2024. [Online]. Available: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_90-120.pdf
- [3] U.S. Department of Transportation, Federal Aviation Administration, “TCAS Transition Program (TTP) Industry Alert Bulletin,” Federal Aviation Administration, Industry Alert Bulletin, August 2002. [Online]. Available: https://www.faa.gov/sites/faa.gov/files/2022-11/TTP_Industry_Alert_Bulletin.pdf
- [4] EUROCONTROL, *Performance Assessment of Pilot Compliance with Traffic Collision Avoidance System Advisories Using Flight Data Monitoring Guidance Material*, 2nd ed., Brussels, Belgium, apr 2017.
- [5] Aircraft and R. A. I. Commission, “Aircraft accident investigation report – japan airlines flight 907 (boeing 747-400d, ja8904) / japan airlines flight 958 (douglas dc-10-40) – a near midair collision over the sea off yaizu city, shizuoka prefecture, japan at about 15:55 jst, january 31, 2001,” Aircraft and Railway Accidents Investigation Commission, Ministry of Land, Infrastructure and Transport, Japan, Investigation Report, July 2002. [Online]. Available: https://jtsb.mlit.go.jp/eng-air_report/JA8904.pdf
- [6] Aviation Safety Council, “Aviation Occurrence Report: China Airlines Flight CI611, Boeing 747-200, B-18255, Near Penghu, Taiwan Strait, May 25, 2002,” Aviation Safety Council, Taiwan, Republic of China, Final Report ASC-AOR-05-02-001, Feb. 2005. [Online]. Available: https://www.tsb.gov.tw/media/4429/ef306_final_report.pdf
- [7] M. Lacagnina, “Easy Does It: TCAS resolution advisories require rapid but not radical response,” *AeroSafety World*, pp. 44–47, oct 2008.
- [8] National Transportation Safety Board, “Loss of Control on Approach, Colgan Air, Inc., Operating as Continental Connection Flight 3407, Bombardier DHC-8-400, N200WQ, Clarence Center, New York, February 12, 2009,” National Transportation Safety Board, Tech. Rep. NTSB/AAR-11/03, May 2011, accident No. DCA10MA029. [Online]. Available: <https://data.ntsb.gov/carol-repgen/api/Aviation/ReportMain/GenerateNewestReport/74981/pdf>
- [9] —, “In-flight Upset, Wheels Up, LLC, operating as a Part 121 scheduled passenger flight, Boeing 737-800, N271LV,” National Transportation Safety Board, Tech. Rep. NTSB/AAR-24/01, May 2024, accident No. DCA21FA066. [Online]. Available: <https://data.ntsb.gov/carol-repgen/api/Aviation/ReportMain/GenerateNewestReport/195189/pdf>
- [10] V. Mellone and S. Frank, “The U.S. Air Traffic Control System Wrestles with the Influence of TCAS,” *Flight Safety Digest*, pp. 1–8, November 1993.
- [11] “ASRS Directline: Issue 4,” *Aviation Safety Reporting System (ASRS) Directline*, no. 4, 1996, accessed 18 Jul 2025. [Online]. Available: <https://asrs.arc.nasa.gov/publications/directline.html>
- [12] N. A. S. R. System, “CALLBACK Issue 463: Those Go-Arounds We Wish Were Better,” *CALLBACK*, no. 463, August 2018. [Online]. Available: https://asrs.arc.nasa.gov/docs/cb/cb_463.pdf
- [13] P. Fleurquin, J. J. Ramasco, and V. M. Eguiluz, “Systemic delay propagation in the US airport network,” *Scientific Reports*, vol. 3, no. 1, Jan. 2013. [Online]. Available: <http://dx.doi.org/10.1038/srep01159>
- [14] S. Calder, “Gatwick disruption cost over £50m,” *The Independent*, jan 2019. [Online]. Available: <https://www.independent.co.uk/travel/news-and-advice/gatwick-drone-airport-cost-easyjet-runway-security-passenger-cancellation-a8739841.html>
- [15] N. Lancefield, “Airport chaos which hit 700,000 passengers and cost £100m caused by engineer’s password glitch,” *The Independent*, nov 2024. [Online]. Available: <https://www.the-independent.com/travel/news-and-advice/nats-airport-chaos-password-wfh-b2647651.html>
- [16] G. Longo, M. Strohmeier, E. Russo, A. Merlo, and V. Lenders, “On a Collision Course: Unveiling Wireless Attacks to the Aircraft Traffic Collision Avoidance System (TCAS),” in *33rd USENIX Security Symposium (USENIX Security 24)*. Philadelphia, PA: USENIX Association, Aug. 2024, pp. 6131–6147. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity24/presentation/longo>

[17] Cybersecurity and Infrastructure Security Agency, “Traffic Alert and Collision Avoidance System (TCAS) II,” ICS Advisory, January 2025, ICSA-25-021-01. [Online]. Available: <https://www.cisa.gov/news-events/ics-advisories/icsa-25-021-01>

[18] CBS News, “D.C. airport plagued by false collision alerts from rogue signals, FAA says,” 2025, accessed on July 2025. [Online]. Available: <https://www.cbsnews.com/news/dc-airport-false-collision-alerts/>

[19] International Civil Aviation Organization, *Procedures for Air Navigation Services — Aircraft Operations, Volume III — Aircraft Operating Procedures*, International Civil Aviation Organization, Montreal, 2018, doc 8168.

[20] —, “Annex 11 to the Convention on International Civil Aviation: Air Traffic Services,” Montréal, QC, Canada, July 2018.

[21] —, “Annex 10 to the Convention on International Civil Aviation, Aeronautical Telecommunications, Volume III: Communication Systems (Part I — Digital Data Communication Systems, Part II — Voice Communication Systems),” Montreal, Canada, July 2007.

[22] International Civil Aviation Organization (ICAO), “Procedures for Air Navigation Services – Air Traffic Management,” International Civil Aviation Organization, Montréal, Canada, Tech. Rep. Doc 4444, 2016.

[23] International Civil Aviation Organization, “Annex 10 to the Convention on International Civil Aviation - Aeronautical Telecommunications - Volume IV: Surveillance and Collision Avoidance Systems, Fifth Edition ,” <https://store.icao.int/en/annex-10-aeronautical-telecommunications-volume-iv-surveillance-radar-and-collision-avoidance-systems>, July 2014, accessed on May 2025.

[24] RTCA, Inc., “Minimum Operational Performance Standards for Air Traffic Control Radar Beacon System/Mode Select (ATCRBS/Mode S) Airborne Equipment with Change 1,” RTCA, Inc., Washington, D.C., Standard D0-181F, jan 2022.

[25] International Civil Aviation Organization, *Technical Provisions for Mode S Services and Extended Squitter*, 2nd ed. Montreal, Canada: ICAO, 2012, no. Doc 9871.

[26] RTCA Inc., “Minimum Operational Performance Standards for 1090 MHz Extended Squitter Automatic Dependent Surveillance — Broadcast (ADS-B) and Traffic Information Services — Broadcast (TIS-B),” RTCA, Washington, DC, Technical Standard DO-260C, 2016.

[27] RTCA, Inc., “DO-185B, Minimum Operational Performance Standards for Traffic Alert and Collision Avoidance System II (TCAS II),” RTCA, Inc., Washington, D.C., Tech. Rep., June 2008.

[28] Bundesstelle für Flugunfalluntersuchung (BFU), “Investigation Report: Accident with a Boeing B757-200 and a Tupolev TU154M near Ueberlingen/Lake of Constance/Germany on 1 July 2002,” German Federal Bureau of Aircraft Accidents Investigation, Braunschweig, Germany, Tech. Rep. AX001-1-2/02, May 2004. [Online]. Available: <http://www.bfu-web.de>

[29] Federal Aviation Administration, “RNAV (RNP) Z RWY 19 - Airport Charts for DCA - Airport Data and Information Portal,” <https://adip.faa.gov/agis/public/#/airportCharts/DCA>, 2025, Accessed on May 2025.

[30] —, “Special Use Airspace,” March 2025. [Online]. Available: https://www.faa.gov/regulations_policies/orders_notices/index.cfm/go/document.information/documentID/1043578

[31] H. C. on Transportation and Infrastructure, “Letter to the Department of Homeland Security Regarding Counter-UAS Activity Near DCA,” https://democrats-transportation.house.gov/imo/media/doc/ti_chs_letter_cuas_near_dca.pdf, Jul. 2025, accessed on July 2025.

[32] Federal Aviation Administration, “RIVER VISUAL RWY 19 - Airport Charts for DCA - Airport Data and Information Portal,” <https://adip.faa.gov/agis/public/#/airportCharts/DCA>, 2025, Accessed on May 2025.

[33] LiveATC.net, “KDCA 3/1/25 1100-1430Z TCAS TA/RA phantom alerts & go-arounds,” KDCA3/1/251100-1430ZTCASTA/Rphantomalerts&go-arounds, 2025, Accessed on July 2025.

[34] VASAviation, “Numerous COLLISION WARNINGS Near Washington DCA Airport!” <https://www.youtube.com/watch?v=pOXV3AjESVU>, Accessed on July 2025.

[35] The ADSB.lol Team, “Historical data for globe.adsb.lol,” https://github.com/adsblol/globe_history_2025, 2025, accessed on July 2025.

[36] D. B. Jenkins, B. A. Wyndham, and P. Banks, “Fine Resolution Errors in Secondary Surveillance Radar Altitude Reporting,” Royal Signals and Radar Establishment (RSRE), Ministry of Defence, Malvern, Worcestershire, UK, RSRE Report 87019, Jan 1988. [Online]. Available: <https://apps.dtic.mil/sti/tr/pdf/ADA196757.pdf>

[37] United Kingdom Overseas Territories Aviation Authority, “ACAS Training for Pilots,” Overseas Territories Aviation Circular, Tech. Rep. OTAC 91-5 / 119-8 / 121-6 / 125-6 / 135-6, 2024, issue 3, 17 May 2024. [Online]. Available: https://www.airsafety.aero/getmedia/02473800-b5c9-4669-a4e9-8765283c322a/20240517_ALPR06_OTAC_91-5_119-8_121-6_125-6_135-6_ACAS_Training_for_Pilots_Issue3.pdf

[38] A. Hoag and G. Sirigu, “DCA TCAS Anomalies Explained,” <https://aireon.com/dca-tcas-anomalies-explained/>, Aireon, March 2025. Accessed on April 2025.

[39] M. Smith, M. Strohmeier, V. Lenders, and I. Martinovic, “Understanding realistic attacks on airborne collision avoidance systems,” *Journal of Transportation Security*, vol. 15, no. 1-2, p. 87–118, Feb. 2022. [Online]. Available: <http://dx.doi.org/10.1007/s12198-021-00238-2>

[40] M. Schafer, M. Strohmeier, V. Lenders, I. Martinovic, and M. Wilhelm, “Bringing up OpenSky: A large-scale ADS-B sensor network for research,” in *IPSN-14 Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*. IEEE, Apr. 2014, p. 83–94. [Online]. Available: <http://dx.doi.org/10.1109/IPSN.2014.6846743>

[41] G. Lykou, G. Iakovakis, and D. Gritzalis, *Aviation Cybersecurity and Cyber-Resilience: Assessing Risk in Air Traffic Management*. Springer International Publishing, 2019, p. 245–260. [Online]. Available: http://dx.doi.org/10.1007/978-3-030-00024-0_13

[42] F. Pearson II, “Map projection equations,” Naval Surface Weapons Center, Dahlgren, VA, Tech. Rep. NSWC/DL-TR-3624 (AD-A037381), 1977.

[43] J. P. Snyder, *Map projections: A working manual*, 1987. [Online]. Available: <http://dx.doi.org/10.3133/pp1395>

[44] A. Johansen, “A tutorial on particle filtering and smoothing: Fifteen years later,” 2009.

[45] R. Douc and O. Cappe, “Comparison of resampling schemes for particle filtering,” in *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005.*, 2005, pp. 64–69.

[46] C. Kuptametee and N. Aunstri, “A review of resampling techniques in particle filtering framework,” *Measurement*, vol. 193, p. 110836, Apr. 2022. [Online]. Available: <http://dx.doi.org/10.1016/j.measurement.2022.110836>

[47] R. A. Johnson, D. W. Wichern *et al.*, “Applied multivariate statistical analysis,” 2002.

[48] W. E. Hoover, “Algorithms for confidence circles and ellipses,” U.S. Department of Commerce, National Oceanic and Atmospheric Administration, National Ocean Service, Charting and Geodetic Services, Rockville, MD, NOAA Technical Report NOS 107 C&GS 3, September 1984. [Online]. Available: https://www.ngs.noaa.gov/PUBS_LIB/AlgorithmsForConfidenceCirclesAndEllipses_TR_NOS107_CGS3.pdf

[49] W. H. Press, *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.

[50] Federal Aviation Administration, “Security Instructions - Washington, DC Special Flight Rules Area (SFRA) 4/9433,” Washington, DC. [Online]. Available: https://www.faa.gov/air_traffic/publications/us_restrictions/procedures/doc/SFRA_FDC_4-9433.pdf

[51] F. Skopik and T. Pahi, “Under false flag: using technical artifacts for cyber attack attribution,” *Cybersecurity*, vol. 3, no. 1, Mar. 2020. [Online]. Available: <http://dx.doi.org/10.1186/s42400-020-00048-4>

[52] E. Habler, R. Bitton, and A. Shabtai, “Assessing Aircraft Security: A Comprehensive Survey and Methodology for Evaluation,” *ACM Computing Surveys*, vol. 56, no. 4, p. 1–40, Nov. 2023. [Online]. Available: <http://dx.doi.org/10.1145/3610772>

[53] N. Mäurer, T. Guggemos, T. Ewert, T. Gräupl, C. Schmitt, and S. Grundner-Culemann, “Security in Digital Aeronautical Communications A Comprehensive Gap Analysis,” *International Journal of Critical Infrastructure Protection*, vol. 38, p. 100549, Sep. 2022. [Online]. Available: <http://dx.doi.org/10.1016/j.ijcip.2022.100549>

[54] M. Strohmeier, M. Schäfer, R. Pinheiro, V. Lenders, and I. Martinovic, “On Perception and Reality in Wireless Air Traffic Communication Security,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 6, pp. 1338–1357, 2017. [Online]. Available: <https://dx.doi.org/10.1109/TITS.2016.2612584>

[55] T. M. Graziano, “Establishment of a Cyber-Physical Systems (CPS) Test Bed to Explore Traffic Collision Avoidance System (TCAS) Vulnerabilities to Cyber Attacks,” Ph.D. dissertation, Virginia Tech, 2021. [Online]. Available: <http://hdl.handle.net/10919/104624>

[56] P. M. Berges, “Exploring the vulnerabilities of traffic collision avoidance systems (TCAS) through software defined radio (SDR)

exploitation,” Ph.D. dissertation, Virginia Tech, 2019. [Online]. Available: http://hdl.handle.net/10919/90165	ADS-B	Automatic Dependent Surveillance-Broadcast [23, §1].
[57] P. M. Berges, B. A. Shivakumar, T. Graziano, R. Gerdes, and Z. B. Celik, “On the feasibility of exploiting traffic collision avoidance system vulnerabilities,” in <i>2020 IEEE Conference on Communications and Network Security (CNS)</i> . IEEE, 2020, pp. 1–6. [Online]. Available: http://dx.doi.org/10.1109/CNS48642.2020.9162216	ARA	Active RAs [23, §4.3.8.4.2.2.1.1].
[58] A. Lomas, “DEFCON 28 Aerospace Village: ILS and TCAS Spoofing,” Tech. Rep., 2020, https://www.pentestpartners.com/security-blog/ils-and-tcas-spoofing/ , accessed on June 2025.	ATC	Air Traffic Control.
[59] J. Hannah, R. Mills, R. Dill, and D. Hodson, “Traffic collision avoidance system: false injection viability,” <i>The Journal of Supercomputing</i> , vol. 77, no. 11, p. 12666–12689, Apr. 2021. [Online]. Available: http://dx.doi.org/10.1007/s11227-021-03766-9	c-UAS	Counter-Unmanned Aircraft Systems.
[60] M. Strohmeier, V. Lenders, and I. Martinovic, “On the Security of the Automatic Dependent Surveillance-Broadcast Protocol,” <i>IEEE Communications Surveys & Tutorials</i> , vol. 17, no. 2, pp. 1066–1087, 2015. [Online]. Available: https://dx.doi.org/10.1109/COMST.2014.2365951	CAS	Collision Avoidance System.
[61] K. Sampigethaya and R. Poovendran, “Security and privacy of future aircraft wireless communications with offboard systems,” 02 2011, pp. 1 – 6. [Online]. Available: https://dx.doi.org/10.1109/COMSNETS.2011.5716527	CDF	Cumulative Distribution Function.
[62] A. Yang, X. Tan, J. Baek, and D. S. Wong, “A New ADS-B Authentication Framework Based on Efficient Hierarchical Identity-Based Signature with Batch Verification,” <i>IEEE Transactions on Services Computing</i> , vol. 10, no. 2, pp. 165–175, 2017. [Online]. Available: https://dx.doi.org/10.1109/TSC.2015.2459709	CISA	U.S. Cybersecurity and Infrastructure Security Agency.
[63] Z. Wu, A. Guo, M. Yue, and L. Liu, “An ADS-B Message Authentication Method Based on Certificateless Short Signature,” <i>IEEE Transactions on Aerospace and Electronic Systems</i> , vol. 56, no. 3, pp. 1742–1753, 2020. [Online]. Available: https://dx.doi.org/10.1109/TAES.2019.2933957	COTS	Commercial Off-The-Shelf.
[64] M. Ngamboé, X. Niu, B. Joly, S. P. Biegler, P. Berthier, R. Benito, G. Rice, J. M. Fernandez, and G. Niculescu, “CABBA: Compatible Authenticated Bandwidth-efficient Broadcast protocol for ADS-B,” <i>International Journal of Critical Infrastructure Protection</i> , vol. 48, p. 100728, Mar. 2025. [Online]. Available: http://dx.doi.org/10.1016/j.ijcip.2024.100728	D-BPSK	Differential binary phase shift keying.
[65] B. Nuseibeh, C. B. Haley, and C. Foster, “Securing the Skies: In Requirements We Trust,” <i>Computer</i> , vol. 42, no. 9, p. 64–72, Sep. 2009. [Online]. Available: http://dx.doi.org/10.1109/MC.2009.299	DCA	Ronald Reagan Washington National Airport.
[66] K. Varner, W. Zaeske, S. Friedrich, A. Kaiser, and A. Bowman, “Agile, post-quantum secure cryptography in avionics,” <i>CEAS Aeronautical Journal</i> , May 2025. [Online]. Available: http://dx.doi.org/10.1007/s13272-025-00806-5	DoD	U.S. Department of Defense.
[67] A. Smith, R. Cassell, T. Breen, R. Hulstrom, and C. Evers, “Methods to Provide System-Wide ADS-B Back-Up, Validation and Security,” in <i>2006 ieee/aiaa 25TH Digital Avionics Systems Conference</i> , 2006, pp. 1–7. [Online]. Available: https://dx.doi.org/10.1109/DASC.2006.313681	ESS	Effective Sample Size.
[68] M. Monteiro, A. Barreto, R. Division, T. Kacem, J. Carvalho, D. Wijesekera, and P. Costa, “Detecting malicious ADS-B broadcasts using wide area multilateration,” in <i>2015 IEEE/AIAA 34th Digital Avionics Systems Conference (DASC)</i> . IEEE, 2015, pp. 4A3–1.	FAA	Federal Aviation Administration.
[69] M. Leonardi and F. Gerardi, “Aircraft Mode S Transponder Fingerprinting for Intrusion Detection,” <i>Aerospace</i> , vol. 7, no. 3, p. 30, Mar. 2020. [Online]. Available: http://dx.doi.org/10.3390/aerospace7030030	FPGA	Field Programmable Gate Array.
[70] G. Gurur, Y. Dalveren, A. Kara, and M. Derawi, “A Radio Frequency Fingerprinting-Based Aircraft Identification Method Using ADS-B Transmissions,” <i>Aerospace</i> , vol. 11, no. 3, p. 235, Mar. 2024. [Online]. Available: http://dx.doi.org/10.3390/aerospace11030235	GNSS	Global Navigation Satellite System.
[71] J. Zhang, F. Ardizzon, M. Piana, G. Shen, and S. Tomasin, “Physical Layer-Based Device Fingerprinting for Wireless Security: From Theory to Practice,” <i>IEEE Transactions on Information Forensics and Security</i> , vol. 20, p. 5296–5325, 2025. [Online]. Available: http://dx.doi.org/10.1109/TIFS.2025.3570118	IAP	Instrument Arrival Procedure.
[72] International Civil Aviation Organization, <i>Doc 8400 - Procedures for Air Navigation Services — ICAO Abbreviations and Codes</i> , 9th ed. Montréal, Quebec, Canada: International Civil Aviation Organization, 2016.	ICAO	International Civil Aviation Organization.
	MLAT	Multilateration [23, §6.1].
	MTE	Multiple Threat Encounter [23, §4.3.8.4.2.2.1.4].
	NTSB	National Transportation Safety Board.
	PA	Pressure Altitude referred to a standard altimeter setting of 1013.25 hectopascals..
	QNH	Altimeter sub-scale setting to obtain elevation when on the ground [72, §p.12].
	RA	Resolution Advisory [23, §4.1].
	RF	Radio Frequency.
	RMS	Root Mean Square.
	RNAV	Area Navigation.
	SDR	Software Defined Radio.
	SLS	Side Lobe Suppression.
	SMC	Sequential Monte Carlo.
	SMC-RAT	Sequential Monte Carlo for Rogue ACAS Transmitters.
	SSR	Secondary Surveillance Radar [23, §1].
	T&I	U.S. House Committee on Transportation and Infrastructure.
	TA	Traffic Advisory [23, §4.1].
	TCAS	Traffic alert and Collision Avoidance System.
	TDOA	Time Difference of Arrival.
	TID	Threat Identity Data subfield [23, §4.3.8.4.2.2.1.6].
	TIDA	Threat Identity Data Altitude subfield [23, §4.3.8.4.2.2.1.6.1].
	TIDB	Threat Identity Data Bearing subfield [23, §4.3.8.4.2.2.1.6.3].
	TIDR	Threat Identity Data Range subfield [23, §4.3.8.4.2.2.1.6.2].
	TTI	Threat Type Indicator subfield [23, §4.3.8.4.2.2.1.5].
	USRP	Universal Software Radio Peripheral.
	USSS	U.S. Secret Service.

APPENDIX

A. Availability and artifacts

In this appendix we provide the means to repeat the analysis performed in the article and reproduce all of the analysis and figures found within it. Furthermore, it provides additional visual proof to some claims found in the manuscript.

1) Description & Requirements

a) How to access

The dataset, scripts, and the source code associated with this article can be found hosted on Zenodo at URL <https://doi.org/10.5281/zenodo.17428590> with DOI 10.5281/zenodo.17428590.

Code	AGPL-3.0-or-later + Commons Clause license condition
Data	CC-BY-4.0

b) Hardware dependencies

None

GLOSSARY

ACAS Airborne Collision Avoidance System [23, §1].

c) Software dependencies

In order to ensure reproducibility, all of the artifacts can be produced by leveraging a containerized environment. As such, the only software dependencies are zstd to extract the archive, a POSIX shell, GNU Make, and Podman version 4 or greater.

d) Benchmarks

None

2) Artifact Installation & Configuration

We have included additional instructions in `README.md` files scattered throughout the artifact directories. Below we describe just the minimal steps.

We recommend using a Fedora Linux 42 system with `make` and `podman` installed from the default repositories.

Please ensure that at least 150 GiB of storage is available.

a) Set-up

- 1) Download, extract, and enter the artifact directory.
- 2) Open a shell inside of the artifact directory.
- 3) Run the test from Section A2b.
- 4) Run `make prepare`.

b) Basic functional test

- 1) Open a shell inside of the artifact directory.
- 2) Run `make check`.

An output with all OKs or SKIPS indicates success.

3) Experiment Workflow

All of the experiments are performed either via manual inspection or by invoking a set of commands contained within the dataset root directory's `Makefile`. Such commands are run within a container in order to ensure a consistent environment w.r.t. the installed dependencies.

The exact metadata for the execution environment used to produce the results found in the article can be found under `output_paper/container-os-release.txt`, `output_paper/container-pip-freeze.txt`, and `output_paper/container-python-version.txt`.

4) Major Claims

a) Section III-A

(C₁) There are 110 ACAS RA messages in our dataset.
(C₂) All ACAS RA messages indicate a Mode C intruder.
(C₃) 102 ACAS RA messages signal a single-threat encounter.
(C₄) The runway at DCA is at 325-375 ft in the ADS-B data.
(C₅) 106 ACAS RA messages indicate the intruder at 2300 ft.
(C₆) 100 out of 110 events ACAS RA messages have a range.
(C₇) TIDR is within 1 nm, at a median distance of 0.2 nm.
(C₈) TIDBs are mostly around 315 to 345° w.r.t. plane heading.

b) Section IV-C

(C₉) A purely SDR-based setup like in Longo et al. is incapable of the latencies needed for a Mode C attack.
(C₁₀) We implemented a Mode A/C transponder compliant with ICAO Annex 10, Volume IV §3.1.1, excluding requirements 3.1.1.{4,5,8,9,10} (ground station), and 3.1.1.7.11 (output power).
(C₁₁) The FPGA implementation has low resource utilization.
(C₁₂) The achievable practical τ with attack S2 is roughly 78% to 82% of its theoretical maximum.

(C₁₃) The achievable practical τ has a precision of around 50 meters around its median value.

c) Section V-J

(C₁₄) "... data analysis showed $a \approx 33\%$ bearing outlier rate".
(C₁₅) The methodology achieves median localization errors for strategies S1, S2, and S3 of 332 m, 855 m, and 1205 m.

d) Section V-K

(C₁₆) The most likely location for the DCA attacks is at 891.7 m from P-56B's center, with $\approx 94\%$ probability.

e) Section V-L

(C₁₇) SMC-RAT computation is efficient, taking on average 7.4 seconds on our benchmark system.

(C₁₈) "When the underlying assumptions are violated, such as ... a mobile rather than a stationary transmitter, the ellipse area increases by orders of magnitude ...".

5) Evaluation

(E₁) SDR latency testing analysis [10 machine-minutes]

Preparation: Steps from Section A2a.

Execution: From the artifact root directory run `make analyze-sdr-latency`.

Results: The analysis should generate the following files under `output`:

sdr-latency.pdf Figure 9 from the article.

sdr-latency.txt Containing the latency at which a reply rate $\geq 90\%$ is achieved (C₉).

(E₂) Transponder test visual inspection [10 man-minutes]

Preparation: Step 1 from Section A2a

Execution: From the artifact root directory, go inside `dataset/photos`, watch the video `tester.mkv`, and the pictures found in the subdirectories.

Results: This inspection should assert that the protocol compliance test was executed as described, passing all tests besides the one related to transmission power (C₁₀). From this content we derive Table III.

(E₃) FPGA utilization report analysis [10 man-minutes]

Preparation: Step 1 from Section A2a.

Execution: From the artifact root directory, go inside `code/fpga`, and read `utilization-report.txt` and `utilization-summary-report.txt`.

Results: This inspection should assert that the FPGA resource usage is well below the utilized platform limits, proving (C₁₁). Those files are the source for Table II.

(E₄) Automated dataset analysis [10 machine-minutes]

Preparation: Steps from Section A2a.

Execution: From the artifact root directory, run `make analyze`.

Results: The analysis should generate the following files under `output`:

acas_bearing_distribution.pdf Figure 4b from the article. Shows claim (C₈).

acas_ra_analysis.md A textual analysis of fields found in the ACAS RA broadcast messages.

(C₁) is under "Number of events".

(C₂) is under "acas_ra_ME_TTI_value", which is equal to 2 in all instances.

(C_3) is under “*acas_ra_ME_MTE_value*”, which is equal to 0 in 102 instances.

(C_5) is under “*acas_ra_ME_TIDA_altitude_ctr*”, which is equal to 2300 in 106 instances.

(C_6) is under “*acas_ra_ME_TIDR_parsed*”, contains 89 “Range estimate”s and 11 “Range estimate is less than 0.05 NM”s.

Part of (C_7) under “*acas_ra_ME_TIDR_ctr*”, the maximum value is 1.0 nm.

acas_ra_range_cdf.pdf Part of (C_7) , showing the median at 0.2 nm. Figure 4a from the article.

acas_ra_range_ENY4035.pdf,

acas_ra_range_JIA5261.pdf,

acas_ra_range_JIA5062.pdf,

acas_ra_range_JIA5312.pdf,

acas_ra_range_RPA4469.pdf,

acas_ra_range_RPA4538.pdf,

acas_ra_range_RPA5802.pdf Figure 5 from the article.

acas_ra_region_plot.pdf Figure 3 from the article.

adsb_min_altitude.md (C_4) follows from the observed values.

spoofing-cdf.pdf Figure 10 from the article.

spoofing.md (C_{12}) and (C_{13}) , by comparing the found values with the theoretical maximum of $16\mu s$. Table IV from the article.

(E_5) *Running SMC-RAT on DCA data [15 machine-minutes]*

Preparation: Steps from Section A2a.

Execution: From the artifact root directory run `make analyze-washington`.

Results: The analysis should generate the following files under `output`:

acas_inliers.json (C_{14}) is under the “`invalid_percentage`” key.

acas_plane_RPA4469.pdf Figure 11 from the article.

acas_localization_partial_7.pdf Figure 13 from the article.

washington_partial_7.json (C_{16}) is found under the “`p56b_distance`” and “`p56b_likelihood`” keys.

(E_6) *Paper’s SMC-RAT analysis [10 machine-minutes]*

Preparation: Steps from Section A2a.

Execution: From the artifact root directory, run `make analyze-paper-simulations`.

Results: The analysis should generate the following files under `output`, copied verbatim from prior runs of Experiment (E_7) or their reanalysis:

benchmark_simulated_trial.json (C_{17}) is found under the “`mean_time_s`”, “`percentiles_s.99`”, and “`max_time_s`” entries.

detection-simulations.txt (C_{15}) is found under the three “median centroid delta” entries. (C_{18}) is found under the “ratio” entry.

simulations_cdf_by_tau_still.pdf Figure 12 from the article.

(E_7) *Evaluating SMC-RAT performance [4 machine-hours]*

This test repeats Experiment (E_6) on a smaller scale.

Preparation: Steps from Section A2a.

Execution: From the artifact root directory, run `make analyze-repeat-simulations`.

Results: The analysis should reproduce on a smaller scale the same files as Experiment (E_6) by generating and performing 1K moving attacker runs, 1K still attacker runs, and 1K SMC-RAT benchmark runs. Observe that (C_{15}) , (C_{17}) and (C_{18}) still hold on this restricted subset by checking the following:

- **For (C_{15}) :** Assert that the median centroid delta entries in `detection-simulations.txt` roughly match the figures in the paper. The distances should scale with τ and be ≤ 400 m, ≤ 1000 m, and ≤ 1500 m for strategies S1, S2, and S3, respectively. Values may vary slightly based on the run; for instance, results around 320 m, 893 m, and 1251 m are acceptable and confirm the claim.
- **For (C_{17}) :** Assert that the `mean_time_s` in `benchmark_simulated_trial.json` is reasonably low. This value is hardware-dependent. For reference, an Intel i9-13900H took ≈ 9 s (`output_e7/output/benchmark_simulated_trial.json`). Any value under 30 seconds confirms the claim.
- **For (C_{18}) :** Assert that the `ratio` entries in `detection-simulations.txt` indicates at least an order of magnitude (i.e., $> 10\times$) of difference between the still and moving trials.

6) Customization

Experiment (E_7) can be customized by changing the number of trials to be performed for each category. To do so, open the `Makefile` and change the variables `N_STILL_TRIALS`, `N_MOVING_TRIALS`, and `N_BENCHMARK_TRIALS`.

B. Planes involved in the events at DCA.

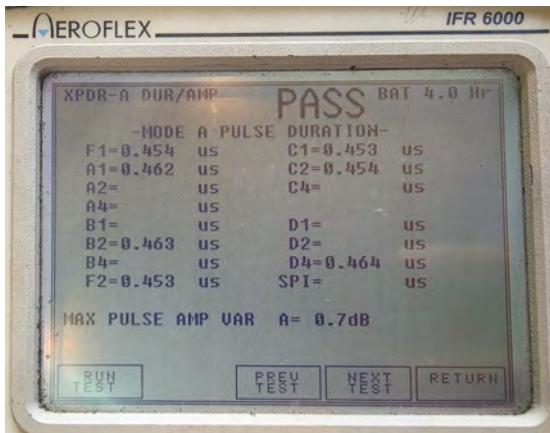
Table V shows the full listing of planes involved in the DCA incidents reported in section III, grouped by commercial operator. Each listing contains the plane flight number, the outcome (TA/RA), Mode-S ICAO code, plane manufacturing year and model.

TABLE V: Planes involved in the events at DCA.

Flight	TA/RA	ICAO	Plane year	Plane model
<i>Flights operated by ENVOY AIR INC</i>				
ENY4035	RA	a33e08	2021	EMBRAER ERJ-170-200 (long wing)
<i>Flights operated by AMERICAN AIRLINES INC</i>				
JIA5062	RA	a6bb97	2010	BOMBARDIER Regional Jet CRJ-700
JIA5098	TA	a6b48f	2010	BOMBARDIER Regional Jet CRJ-700
JIA5146	TA	a95d52	2004	BOMBARDIER Regional Jet CRJ-700
JIA5197	RA	a6b840	2010	BOMBARDIER Regional Jet CRJ-700
JIA5261	RA	a9739c	2004	BOMBARDIER Regional Jet CRJ-700
JIA5312	RA	a6eeb9	2011	BOMBARDIER Regional Jet CRJ-700
<i>Flights operated by REPUBLIC AIRWAYS INC</i>				
RPA4469	RA	a01097	2007	EMBRAER ERJ-170-200 (long wing)
RPA4538	RA	a0835d	2008	EMBRAER ERJ-170-200 (long wing)
RPA4549	RA	a4d04c	2013	EMBRAER ERJ-170-200 (long wing)
RPA5802	RA	a1ccce	2009	EMBRAER ERJ-170-200 (long wing)

C. Aeroflex IFR6000 test results.

Figure 14 shows photographs of the tests performed using the Aeroflex IFR6000.



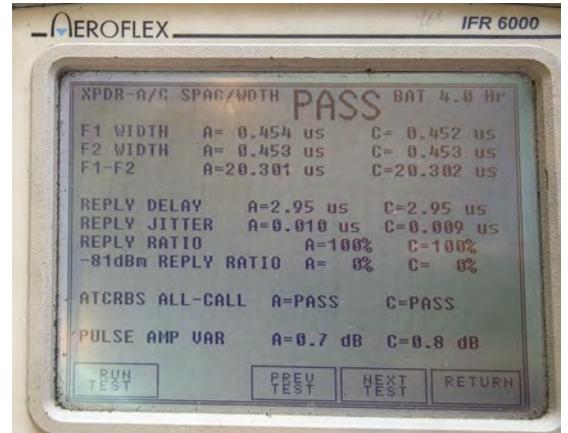
(a) Test for Mode A.



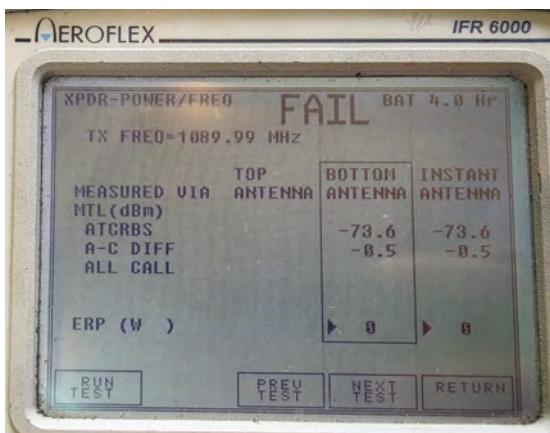
(b) Test for Mode C.



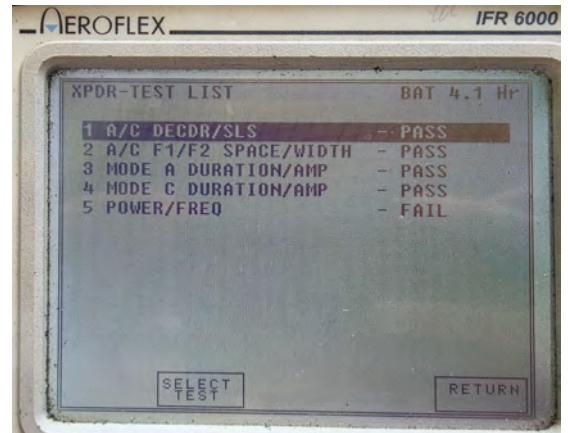
(c) Test of encoding and SLS.



(d) Test of reply rate and pulse spacing.



(e) Test of power and frequency.



(f) Test summary.

Fig. 14: Aeroflex IFR6000 test results: photographs of the performed tests.