Flying in Private Mode: Understanding and Improving the Privacy ICAO Address Program

Guillaume Michel*
Swiss Federal Institute of Technology in Lausanne, 1015 Lausanne, Switzerland

Martin Strohmeier†
armasuisse Science and Technology, 3603 Thun, Switzerland

Location privacy of aircraft has recently gained attention as air traffic management was modernized using novel surveillance technologies. Business aviation circles and various military and government entities voiced serious concerns about automated and ubiquitous tracking. Consequently, some flight authorities have started addressing these operational privacy issues with novel programs.

We conduct a first analysis of the Privacy ICAO Address (PIA) program launched by the Federal Aviation Administration on 1 January 2020, aiming to increase privacy of General Aviation in the United States. A methodology using air traffic communication data gathered from crowdsourced networks is demonstrated to identify and track aircraft enrolled in this program, demonstrating the privacy performance of the program does not meet its goals. Using this method, we identify 14 exemplary aircraft enrolled in the program, this number is expected to grow significantly in the future.

We further predict the future efficacy of the PIA program with a novel aircraft privacy simulator. We show that after 100 days, on average 69.2% of a fleet of 100 aircraft enrolled in the program, can be tracked. We suggest two improvements to the program, which would significantly decrease aircraft traceability to 44.0% after 100 days, and 0.89% after 42 days, respectively.

I. Introduction

In recent years, Air Traffic Management (ATM) systems around the world have been updating their infrastructure in order to integrate new surveillance technologies. At the core of this modernization, Automatic Dependent Surveillance – Broadcast (ADS-B) has now become the standard for air traffic surveillance in many countries [1]. For example, the US Federal Aviation Administration (FAA) made ADS-B equipment mandatory in the US airspace from 1 January 2020, as part of the NextGen program.

*Cybersecurity student, guillaume.michel@epfl.ch
†Cyber Defence Campus (CYD) scientist, martin.strohmeier@armasuisse.ch

ADS-B is a step change away from independent radar surveillance towards all aircraft deriving their own position via GNSS and broadcasting it towards other aircraft and ground stations. Incidentally, affordable Software Defined Radios (SDR) are able to receive such ADS-B transmissions within a radius of up to 600 km. Crowdsourced web services such as the OpenSky Network [2] and FlightRadar24 [3] exploit this fact and distribute ADS-B receivers to volunteers around the world in order to obtain global coverage of the airspace. These crowdsourced networks publicly display aircraft location and their identity in real time, enabling anybody to trivially track transponder-equipped aircraft.

Aircraft are usually identified by their call sign and a globally unique 24-bit identifier assigned by the International Civil Aviation Organization (ICAO). Both identifiers are transmitted via ADS-B, thus are publicly accessible through crowdsourced networks. For commercial flights, such tracking is generally not an concern as their flight plans are public and known in advance. However, researchers have shown that it can cause operational privacy issues for military [4], government and business aircraft [5]. The Aircraft Owners and Pilots Association (AOPA), concerned by such privacy issues, petitioned the FAA to allow aircraft to fly anonymously [6]. In response, the FAA proposed a concept of periodically-assigned random ICAO addresses, with the end goal of increased anonymity for general aviation operators.

The FAA launched this so-called Privacy ICAO Address (PIA) program on 1 January 2020 [7]. This program allows enrolled aircraft to request a new ICAO address every 60 calendar days in its first phase, and every 20 business days in the second phase. While more aircraft are expected to join this program, at launch only a relatively small number of aircraft were using it. We show in this paper that the efficacy of the PIA program depends strongly on the number of aircraft using it concurrently and suffers from incomplete implementation. Consequently, our work illustrates that the PIA program does not provide the anonymity level expected by aircraft operators in the first phase of the program nor in its second phase.

Our paper contains the following contributions:

1) We define the aircraft traceability index, a novel privacy metric appropriate for the aircraft tracking scenario.
2) We provide a first real-world privacy analysis of the FAA PIA program. We introduce the program and the related procedures and illustrate how it has been used in the United States from January to May 2020. We evaluate the privacy offered to enrolled aircraft by conducting automated tracking of aircraft movements using publicly accessible crowdsourced networks.
3) We build an open-source simulator for the aircraft tracking scenario. Using this simulator, we examine aircraft tracking efficacy under different constraints and predict how the system could scale in the future.
4) Finally, we suggest two possible improvements to the PIA program enhancing privacy in General Aviation; changing aircraft addresses for each flight or for all aircraft at the same time. We use our simulator to evaluate the effect of these changes on the participants' privacy.

The remainder of this paper is organized as follows: Section II provides the background concerning ADS-B, the PIA program and metrics used in this paper. Section III.C describes the threat model. Section IV contains the description of

an automated aircraft tracking method for aircraft in the PIA program. Section V details our real-world observations and gives predictions on how privacy would scale with additional aircraft. Section VI proposes two privacy improvements for the program before Section VII concludes the analysis.

II. Background

In this section we briefly introduce basic modern aircraft communication devices and protocols. This covers the ADS-B surveillance technology and the PIA program in the United States, along with some other aircraft privacy measures. It introduces the privacy metrics used in this analysis and briefly discusses related work.

A. Aircraft Identifiers

An aircraft in flight needs to be uniquely identifiable for obvious safety reasons. All aircraft must be registered in a country and are assigned a globally unique registration code, also know as *Tail Number*. In the United States, most registration codes start with the prefix `N`, except for military aircraft. Hence, this code is sometimes called *N-Number*.

Each aircraft also has a unique 24-bit transponder address. The ICAO assigns ranges of these 24-bit addresses to states, the state aviation authority is then responsible for assigning them to registered aircraft. In the United States, there is a mapping from the registration code to the ICAO address of any aircraft, hence it is possible to determine the registration code from the ICAO address and vice versa [8] but this is not generally true around the world. Both registration codes and ICAO addresses are usually meant to remain static unless an aircraft is sold but only the ICAO address is generally broadcast by the aircraft transponder.

For each flight, an aircraft must further use an identifier, which in most cases is a *call sign*. This mutable identifier can be set separately through the flight deck for each flight. Commercial aircraft mostly use the flight number as call sign while private aircraft usually use their registration code or a private call sign [9]. Tracking an aircraft by its ICAO address thus is more accurate than tracking it by its call sign, as the ICAO address is meant to be immutable whereas a call sign can easily be changed.

B. ADS-B

ADS-B is a modern air traffic surveillance technology. It allows aircraft to broadcast their position to air traffic control (ATC) and in turn receive traffic data. It aspires to replace Secondary Surveillance Radar (SSR) as the main surveillance technology for controlling aircraft worldwide. Multiple crowdsourced networks gather this unencrypted ADS-B data worldwide and make it publicly available online. ADS-B equipment consists of two different components: *ADS-B Out*, and *ADS-B In*. ADS-B Out is made of two elements: a certified GNSS position source and a datalink radio. ADS-B Out will periodically, twice a second, broadcast information gathered by the GNSS, comprising its current location, altitude and velocity on the 1090 Extended Squitter (ES)/ Universal Access Transceiver (UAT) data links [11].

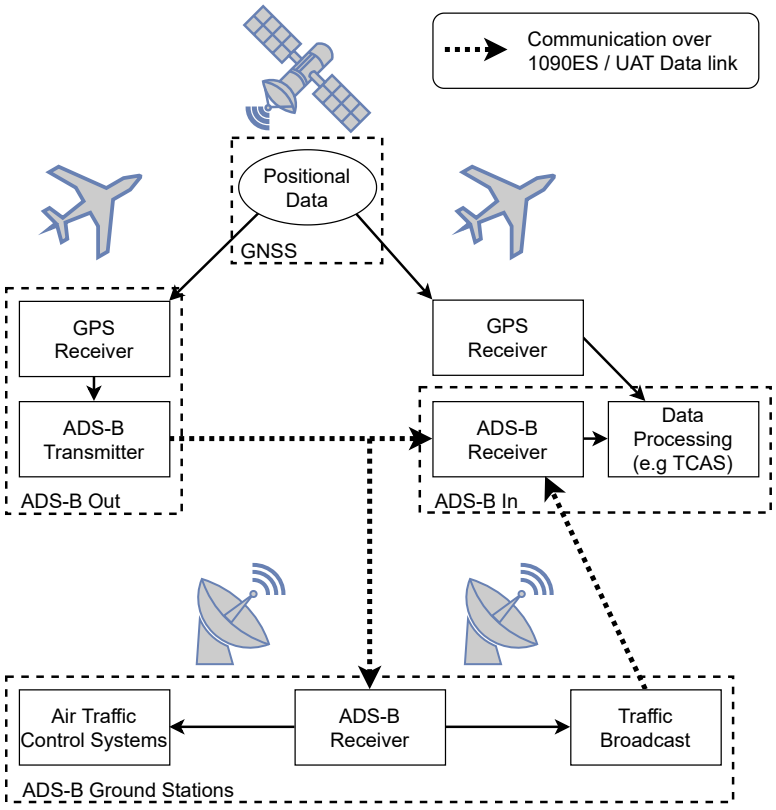

Fig. 1 Illustration of the ADS-B ecosystem modeled after [10].

ADS-B In systems receive traffic data sent by ground stations and other aircraft on the data link as illustrated in Figure 1. ADS-B broadcasts the ICAO 24-bit transponder code specifically assigned to the aircraft as well as the flight's chosen call sign.

Multiple countries, such as Australia, Canada, European countries and the United States have a mandate requiring aircraft flying in their airspace to use ADS-B systems [1]. In the United States, ADS-B as a new standard is defined by the FAA's NextGen program. This program forces all aircraft flying in the US airspace (with a few specific exemptions) to communicate their position using ADS-B, from 1 January 2020 [12]. For each flight, the flight plan's call sign must match the call sign transmitted via ADS-B.

C. Previous ADS-B Privacy Policies

ADS-B is known to have multiple security and privacy issues stemming from its lack of encryption [10] [13]. As ADS-B transmissions contain the call sign and the ICAO 24-bit transponder address of the aircraft, tracking an aircraft by looking up its ICAO address or call sign on a global flight tracking service such as OpenSky Network [2] or FlightRadar24 [3] is easy. That is usually not an issue for commercial flights, but it may be problematic for private or government aircraft operators, as demonstrated by [5]. To address such privacy issues caused by tracking, third-party

call sign providers such as FltPlan* or ForeFlight† have emerged in the US. They provide their customers with private call signs registered under the companies' airline labels. Premium services are also available, which allow aircraft operators to change the private call sign for each flight. This approach mitigates private aircraft tracking by call sign, but tracking an aircraft by its ICAO address remains trivial.

Further to this, the FAA has long had its own privacy program called Limiting Aircraft Data Displayed (LADD) which allows aircraft operators to request to limit flight tracking information transmitted over the internet by the FAA. This program also prevents all participating vendors to publicly display information about any aircraft on the blocked list [14]. Large commercial provides such as Flightradar24 or FlightAware participate in this agreement and withhold information on blocked aircraft on their websites. This measure does not solve the issue of aircraft privacy as non-participating crowdsourced networks such as the OpenSky Network are still able to track the aircraft and make their data publicly available online. This measure may even have a negative impact on privacy as blocked aircraft can easily be identified when comparing Flightradar24 and OpenSky Network data. Once identified as potentially sensitive, these aircraft can then be tracked freely on the OpenSky Network.

Beyond the standard ADS-B system, the Universal Access Transceiver (UAT) datalink offers an *Anonymous Mode* for general aviation aircraft in the US. Aircraft using this Anonymous Mode can self-generate and use a pseudonymous identifier instead of their true ICAO address. Air Traffic Controllers can monitor aircraft flying with these pseudonyms as they are provided out of band with the true ICAO addresses of the aircraft and have access to their ADS-B transmissions. [15] demonstrates how to efficiently track when these pseudonyms change. An observer may link two pseudonyms or a pseudonym with an true ICAO address, based on temporal and spatial correlation between consecutive locations of aircraft. As a consequence, the *Anonymous Mode* offers no tracking privacy when considering a stronger adversary, who has access to a crowdsourced network.

Finally, the aircraft registrations are publicly available in the Civil Aviation Registry (CAR) of several countries. This implies that aircraft ownership is a public information, and it is theoretically possible to identify the owner of a given aircraft. However, as described in [5], some aircraft owners obfuscate their aircraft ownership and, for example, register their aircraft to Limited Liability Companies (LLC). Even though establishing a link between the LLC and the actual aircraft owner is not trivial, this obfuscation does not mitigate the aircraft tracking issue discussed in the present paper.

D. Privacy ICAO Address (PIA) program

The FAA introduced the PIA program starting on 1 January 2020 [7]. This program is limited to (1) U.S registered aircraft, (2) equipped with a 1090 MHz ADS-B transponder, (3) using a third-party call sign and (4) flying in domestic

*https://www.fltplan.com/
†https://www.foreflight.com/

U.S. airspace. This program allows aircraft operators to request an alternate, temporary ICAO Address, which will not appear in the CAR. Hence, aircraft operators willing to enroll in this program must be able to change the transponder identifier. The ICAO address change must be performed on ground, it is forbidden to perform the change during flight. Only one PIA can be issued at a time for an aircraft. These PIAs are valid indefinitely, until a new change is requested by the aircraft operator.

As of May 2020, there were two authorized private flight id providers, FltPlan and ForeFlight, providing call signs starting with the prefixes `DCM` and `FFL`, respectively. Aircraft operators willing to register into the PIA program should use one of these two flight ID providers (a third, FlightAware, was added recently). The program is planned to be made available in two phases. In the first phase, the service is operated, monitored and maintained by the FAA. Aircraft operators are able to request a new PIA at any time following a 60-calendar day period from a previous PIA assignment. In the second phase, the program will be transitioned to third party service providers that will operate, monitor and maintain it. Once phase two is in place, the FAA service will be removed. Aircraft operators will be able to request a new PIA at any time following a 20-business day period from a previous PIA assignment. The system will fundamentally stay the same, the only observable difference is the PIA update frequency. It is not permitted to use a PIA outside of the US airspace. For international flights, aircraft enrolled in the PIA must use their old, permanently assigned ICAO address [16].

As of May 2020, in order to enroll into the PIA program, an aircraft operator has to (1) provide a Public ADS-B Performance Report (PAPR) using the aircraft's permanently assigned ICAO, to the FAA. This enables the FAA to assess the ADS-B Out equipment. (2) The aircraft operator has to submit a PIA request to the FAA. Upon confirmation, a new PIA will be assigned to the aircraft. (3) Aircraft enrollment to a third-party flight ID provider needs to be proven to the FAA. (4) Upon success of the previous steps, the new temporary PIA must be programmed into the aircraft's transmitter. (5) Within 30 days, the aircraft operator has to perform a test flight from at least 30 minutes to provide a new PAPR, and validate the PIA on the FAA website. When requesting a new PIA, steps (2), (4) and (5) are necessary before the validation of this new PIA. Thus, we assume that the process of requesting a new PIA involves manual administrative work and may take several business days.

The PIA can be used as a pseudonym replacing the ICAO address for a given period of time. This policy does not in any case provide anonymity but arguably pseudonimity to enrolled aircraft for a given period of time, and assumes the hardness of establishing a link between an old and a new PIA when a change occurs. Intuitively, to optimize the privacy level, aircraft operators should request a new PIA as often as allowed to minimize the time they use the same pseudonym, as they can be tracked by their pseudonym.

III. Privacy Metrics for Aircraft Tracking

In this section we apply the traditional mixnets approach to the aircraft domain and develop a novel traceability index that considers the special requirements of the aviation scenario. Finally, we model the threat to aircraft privacy as it applies to this scenario.

A. Aircraft Mixnets

When measuring privacy, common metrics in the literature are *k-anonymity*, *l-diversity* and *t-closeness*. In the present case, it makes sense to apply these metrics when multiple aircraft change simultaneously both their PIA and call sign during the same time period. An aircraft is called *k-anonymous* if it cannot be distinguished among k other aircraft with a probability better than random. *k-anonymity* is usually extended with *l-diversity* and *t-closeness*. For example, for multiple distinguishable types of aircraft, one can say that one aircraft a is *k-anonymous*, *l-diverse* and *t-close* according to its type in a given set of aircraft, if it cannot be distinguished from $(k-1)$ other aircraft, there are l well represented types of aircraft, and the distance between the distribution of the type of a and all other types is no more than a threshold t.

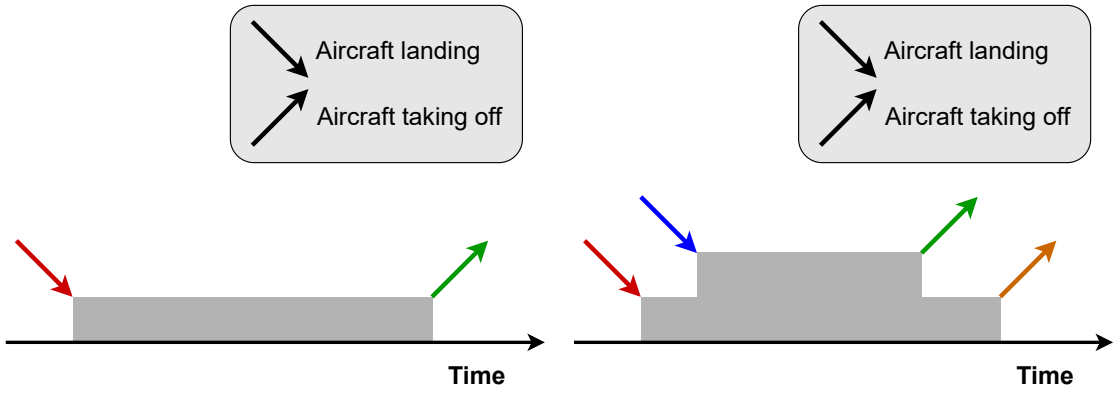

(a) Airport in the presence of a single aircraft **(b) Airport in the presence of two distinct aircraft**

Fig. 2 The gray bars represent the count of aircraft on ground at the given airport over time. The colors in the arrows represent the aircraft identifier.

Figure 2 represents aircraft arriving to and leaving from a given airport. An aircraft changes its PIA and call sign, represented by colors, between its arrival and its departure. An external observer can observe an aircraft identified by its PIA and call sign arriving at an airport, and another aircraft leaving the airport. Figure 2a represents a single aircraft landing and leaving the airport with a different PIA and call sign. In this case, it is trivial for an adversary to assert that an aircraft first landed and then the same aircraft departed from the airport using a different identifier. In Figure 2b, as two aircraft land at the airport and change their PIA and call sign before departing, the adversary has no better choice than a random guess to associate the aircraft leaving the airport with those who arrived before. In this case, we say that

both aircraft are *2-anonymous*. The airport serve as a mix station for aircraft changing their PIA and call sign during the same time period, while they are staying at the airport. An adversary cannot accurately associate incoming and outgoing aircraft using only their ADS-B transmissions.

We now consider a *set of airport*s as a large, various latency, free route, mix network (mixnet) [17]. The aircraft are shuffled at airports when they change their identifiers (call sign and ICAO address), which makes them harder to track over time. Aircraft fly from one airport to another and are shuffled with the other aircraft on ground at the same airport, changing their identifiers simultaneously.

In classical mixnets, elements arrive one-by-one or by batch at a mix station, and leave by batch after the mix. Hence, the duration of an element stay at a mix station can be upper-bounded by the time interval between two batch releases. Elements usually follow random paths across the network for a better anonymity. In many classical mixnets such as Tor [18] elements go through the mixnet to become untraceable, thus the elements are shuffled between the mixnet multiple entry and exit points.

The mixnet we will use to quantify aircraft anonymity is different from classical mixnets in a few details. Our mixnet is considered to be of variable latency, in contrast with classical mixnets it is thus not latency dependent. Aircraft arrive one by one, stay for a variable duration at an airport, serving as mix station, and leave one by one independently. Unlike in classical mixnets, no flight trajectory is random, which facilitates correlational analysis. Furthermore, the mixnet of airports has no entry nor exit points, all aircraft stay in the mixnet from their first until their last flight.

B. Traceability Index

We would like to measure the overall privacy level of the global system over time and not only during a single PIA change. *k-anonmity*, *l-diversity* and *t-closeness* metrics can measure the privacy level of the set of aircraft changing their PIA simultaneously at a given airport at a point in time. But these metrics make less sense when trying to retrace the flight history of an aircraft using data provided by a crowdsourced network over a long period of time. The unit we use to measure privacy of the PIA program throughout this analysis is the *traceability index*, represented as the expected ratio of successfully tracked aircraft over time. This unit is the measure of the variable latency free route mixnet efficiency. It is computed as the combination of numerous *k-anonymous*, *l-diverse* and *t-close* mix steps over time. Each aircraft follows its own flight frequency distribution. Thus the traceability index of the system depends on the time, represented as the expected number of flights for each aircraft. *k-anonmity*, *l-diversity* and *t-closeness* provide the probability of a single PIA change identification. Whereas the *traceabilty index* represent the number of aircraft successfully followed over time, which is a combination of multiple *k-anonmity* steps. For example, if we are tracking aircraft in a system containing 30 aircraft updating regularly their identifiers, and after 10 days we have lost 9 of them, the system's traceability index after 10 days is $(30 - 9)/30 = 0.7$. To be able to measure the traceability index of a system, we assume that all flights data is recorded and accessible.

Some aircraft repeat the same flying pattern, periodically flying between the two same airports. Hence, they are easy to identify. Some airports are weakly connected between each other in terms of flights, sometimes only a single aircraft is flying between such a pair of airports. Thus, if there is a flight between any such pair of airports, the aircraft's identity can easily be guessed. Identifying obvious aircraft increases the tracking accuracy for other aircraft in the system, as it reduces the number of *anonymous* aircraft. As imposing random, diverse trajectories to aircraft operators in order to increase the global privacy level would be nonsense, aircraft privacy is upper bounded by the diversity of actual flights. If an aircraft is the only one using PIA in an airport, it is trivial to link its old and new PIA after a change, it can thus be identified at all times.

C. Threat Model

We consider an adversary who has complete access to the data of a crowdsourced network covering most of the US airspace. A single SDR can receive the signal emitted by aircraft within a radius of 600 km. Crowdsourced networks of such receivers, for example the OpenSky Network [2] or Flightradar24 [3], have a good global coverage, particularly sufficient in the US airspace.

We mainly consider aircraft tracking on historical data but also discuss live tracking. In the first case, the adversary has access to all trajectory data for a given time window and tries to track aircraft's movements over time. An example would be an investigation whether a CEO used their corporate aircraft for personal business. The difficulty of this operation lies in making correct assumptions when multiple aircraft update both their PIA and call sign in a small time interval at the same airport.

The live tracking case focuses on associating a newly observed PIA with a known aircraft at the very moment when a fresh PIA is broadcasted. This would, for instance, allow paparazzi, to learn where a celebrity is headed in real time. Live tracking is harder to realize as no future data is available for use. An adversary able to perform live tracking can also perform tracking on historical data. However the reverse is not true.

As of May 2020, two third-party call sign providers were members of the PIA program. Since their call sign is public, we can consider each as a distinct set of aircraft, even though they use the same ICAO addresses range. An adversary can trivially split aircraft using `DCM` and `FFL` call signs in two distinct sets. For simplicity, we assume that aircraft do not switch from one third-party call sign provider to another.

IV. Methodology

In order to track aircraft enrolled in the FAA PIA program efficiently, an adversary needs to first detect an aircraft using a PIA, then to associate this PIA with a FAA registration in the CAR, and finally to track any PIA change for the aircraft.

A. Detecting Aircraft using PIAs

In order to track aircraft enrolled in the PIA program even though they changed their ICAO address and possibly their call sign, it is necessary to identify the aircraft's registration in the CAR. We discovered that all aircraft enrolled in the PIA progam are assigned an N-Number in the range between N41000 and N42. When looking up these aircraft in the CAR, they appear as *Reserved* with *No fee* by *SBS PROGRAM OFFICE* on the *10/03/2019*. This N-Number range is not fully reserved for the PIA program, some 'normal' registrations appear in the same range. In total, 1062 N-Numbers in this range were reserved by the *SBS PROGRAM OFFICE* in May 2020. This number is still limited but we assume that the FAA can easily extend this range when more aircraft enroll in the program. The range of ICAO addresses associated with this N-Number range is from *a4d691* to *a4f945* [8]. We note that these reservations appear on the online public FAA database but not on the database available for download on the FAA website. As aircraft enrolled in the PIA must further have either a third-party call sign, we can now systematically search for ADS-B transmissions from aircraft whose call sign is either `DCM` or `FFL`, whose ICAO address is between a4d691 to a4f945, and whose associated N-Number registration in the CAR is *Reserved* by *SBS PROGRAM OFFICE*.

B. Associating a PIA with a FAA registration

Once an aircraft with these properties is found, we need to look at its first flight f_0 using the PIA, which should be after 01/01/2020. We note that some aircraft keep using the same call sign when switching their ICAO address to a PIA, forfeiting any privacy gain from the get-go. To identify the old ICAO address (associated with the true FAA registration indicating the aircraft's type and owner), we simply need to find the last flight f^* using the same call sign (DCM or FFL) as f_0 before f_0 happens. Intuitively, if f^* lands at airport A, and f_0 depart from airport A it is an almost certain match.

A few aircraft in our observed set changed their call sign when starting to use a PIA for the first time. For these aircraft, the strategy is to identify the departure airport A of their first flight f_0 while using a PIA. From there, with future data, we can observe future flights using the same PIA to get an approximate flight distribution. Then, we look for all aircraft using an ICAO whose very last flight was observed landing at A before the departure of f_0, and select the closest flight distribution, based on past data.

C. Tracking PIA changes

An update of the PIA is always conducted while the aircraft is on ground at an airport. If the aircraft is the only one to get a new PIA at airport A during its stay, we can trivially associate the old PIA arriving at A and never departing again, with the newly observed PIA departing from A.

As aircraft will be allowed to request a new PIA 20 business days after the last PIA assignment, the likelihood increases that multiple aircraft change their PIA during the same time frame at the same airport. We can use several methods to increase our chance of correct association of the candidates c under these circumstances. First, it is

possible to identify an aircraft category from its flight data simply by looking at its trajectory data [9]. Using this meta information, the candidates can be filtered and only those with matching characteristics kept. Second, if an aircraft using a newly-assigned PIA is going to the origin of one of the candidates last flight, then there is a high probability that one of c is going back to its origin, as return flights are highly typical. Other flight pattern analysis can be applied in order to increase the assignment accuracy in the rare cases with multiple candidates.

Finally, whenever an aircraft is flying outside of the US airspace, it must use its permanently assigned true ICAO address associated with its N-Number registered in the CAR. Hence, flying abroad reveals the identity of aircraft using a PIA and resets the operational security gains.

V. Results

This section presents our results from the real-world observations and experiments we conducted. It also introduces our purpose-built simulator, which we use to demonstrate the effect of scaling several relevant parameters, notably the number of aircraft in the PIA scheme.

A. Real-World Analysis

We performed real-world observations based on OpenSky Network [2] data between January and May 2020. Our data and scripts are available on Github [19]. We detected that 16 aircraft in total have been using a PIA during this period, 9 with a `DCM` call sign, and 7 with a `FFL` call sign. We could reliably identify their true registrations using the method described above. Some aircraft did not change their call sign when they switched from their permanently assigned ICAO address to their new PIA and were thus trivial to identify. For all other aircraft, we successfully applied the strategy of analyzing flight patterns as described in Section IV. To preserve the anonymity of specific aircraft, we do not discuss the particular details of their cases.

We further found that none of the aircraft updated its PIA during the observed time period, which was longer than the minimal update frequency during phase 1 of the PIA program (60 calendar days). Hence, we can speculate that aircraft operators either do not seek to update their PIA as often as possible or that the FAA is unable to process their requests fast enough.

Furthermore, we observe that aircraft with a PIA do not change their call sign over time. As we could not observe a single PIA change, we cannot yet say whether aircraft might update their call sign at the same time, which would be desirable. If the call sign update is not observed when changing to a new PIA, this change would be useless as anyone could still track the aircraft by its call sign. Historically, we found multiple aircraft using FltPlan as third-party call sign provider, which were changing their call sign with every new flight in 2019 (previous to the PIA deployment).[‡]

Further analysis uncovered that PIA users with a `DCM` call sign are generally hidden on Flightradar24 and FlightAware,

[‡]This feature is available for FltPlan premium users and aims to increase their privacy although it does not protect against ICAO address tracking.

i.e., their flights do not appear on the live map nor in the flights history. By doing a differential analysis with data from Flightradar24 and OpenSky, we can detect which aircraft are members of the LADD program, increasing their probability of also being part of the PIA program. We note that PIA users with a `FFL` call sign are not hidden by Flightradar24.

B. Simulator-based Analysis

To extend our analysis of the PIA program beyond the scale of the still small real-world data, we developed a simulator [20]. This section discusses our simulator's design and assumptions and provides the results of varying the main parameters influencing the privacy guarantees of the program.

1. Simulator overview

In order to obtain realistic inputs (e.g., flight patterns and frequency) for our simulator, we first analyzed OpenSky's business and general aviation flights data for the year 2019. As each real-world business aircraft has its own unique flying pattern, it is difficult to simulate new aircraft with truly realistic flights patterns. We eventually chose to simulate random flights as this option can provide a true lower bound for the traceability index of the PIA program. With a random flight mode, the entropy is maximized as there is no possible correlation between flights. In the real world, typical correlation patterns make aircraft identification easier. For example, many aircraft perform round trip flights: if an aircraft goes from airport A_0 to A_1, there is a non-negligible probability that its next flight will be the return flight from A_1 to A_0. Our simulation foregoes this additional information in order to obtain the traceability lower bound.

The performance is measured by the system's average traceability index, which is the percentage of aircraft successfully tracked over time. When an aircraft is not tracked successfully anymore, it is considered as lost by the simulator. In the real world, if a lost aircraft flies abroad, it will trivially be identified again, as it has to use its publicly assigned ICAO address. The same occurs when someone takes a photo of an aircraft and uploads it with the time and location of the picture. However, we did not attempt to model these techniques in the simulator in order to obtain the lower bound of the traceability index.

Table 1 Simulation parameters and value ranges.

Parameter name	Values [default]
Aircraft number	2-1000 [200]
Airport number	2-1000 [100]
Number of aircraft types	1-20 [1]
PIA change frequency	0-60 days (simultaneous), each flight [28 days]
Future data	none, 1 day-10 years [30 days]
Average flight frequency	0.1-10 flights/day/aircraft [0.31]
Simulation duration	1 day-10 years [1 year]

The simulation parameters are displayed in Table 1. The system takes an arbitrary number of aircraft, which represent the number of aircraft enrolled in the program, all using the same call sign provider and strictly following the PIA update policy. For example, if 300 aircraft are enrolled in the program, 200 using a `DCM` call sign and 100 using a `FFL` call sign, and 50% of them update their PIA and call sign every 28 days, the relevant aircraft number for our analysis should be $200 * 0.5 = 100$ or $100 * 0.5 = 50$ depending on the targeted call sign.

The simulator first creates a set of aircraft, each with a defined type selected at random from the given set of aircraft types. Each aircraft is placed at a random airport taken from the given set of airports of real US airports. Each airport from the set has latitude and longitude coordinates. Then, the simulator simulates flights from random aircraft being on ground to random airports according to the flight frequency for the simulation duration. The flight duration is defined as the distance between two airports divided by the aircraft's type speed. The arrival airport can be the same as the departure airport. The simulator further takes a policy as input, which defines when aircraft change their PIA and call sign.

We implemented two policies, corresponding to the two phases of the PIA program. The first policy means aircraft update the PIA along with the call sign every 60 calendar days. The second policy updates the PIA along with the call sign every 28 calendar days, which correspond to 20 business days, as described by the FAA [7]. The first PIA change date follows a uniform distribution for all aircraft, and then the PIA is updated at the appropriate frequency. Additionally, we also implemented a baseline policy with no PIA change, which resulted in a 100% traceability index. To track the aircraft, we implemented the methodology described in Section IV. In this methodology, for each PIA guess, we are using data up to 30 days after the flight's date.

2. Fleet size

In order to provide *k-anonymity* to an aircraft when it updates its PIA, there must be at least $(k-1)$ other aircraft changing their PIA at the same airport during the same time frame. Hence, the traceability index in the PIA program depends on the number of enrolled aircraft as well as the number of airports they frequent. Consequently, to maximize the privacy level, one should maximize the number of aircraft per airport changing their PIA and call sign during the same time period.

Figure 3 represents the traceability index over a year for different sizes of fleets. We used a set of 100 airports, all simulated aircraft are of a single aircraft type, and both PIA and call signs are updated every 28 days. As expected, we can observe that a larger fleet provides a lower traceability index, for a constant number of frequented airports. After 150 days, the traceability index of the 50 aircraft fleet simulation is 87.9%, 57.8% for the 100 aircraft fleet, 15.3% for the 200 aircraft fleet, and 1.5% for the 500 aircraft fleet. Thus, increasing the number of aircraft for a fixed number of frequented airports decreases the traceability index significantly.

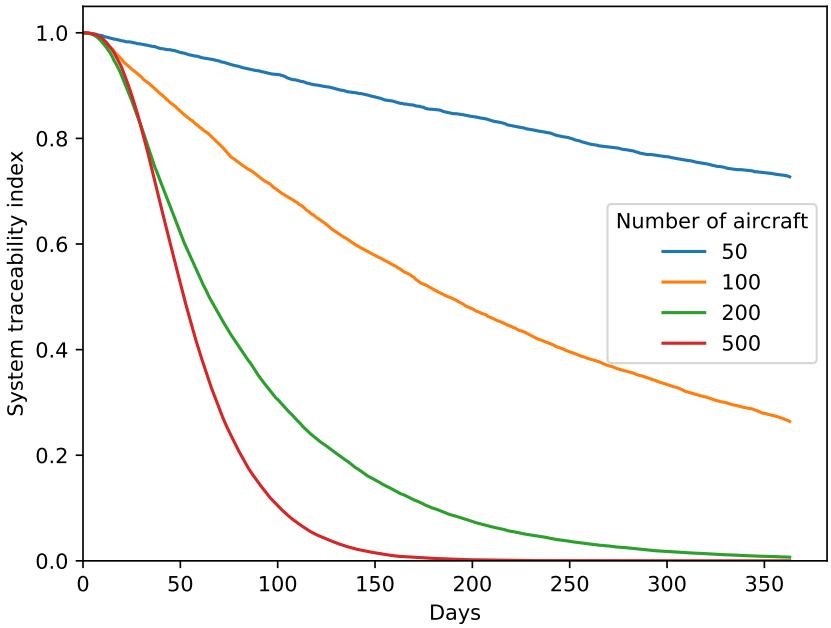

Fig. 3 System traceability index over time depending on the number of aircraft in the system. 100 airports, 1 aircraft type, 28 days PIA update frequency, 100 simulations.

3. Frequented airports

The number of involved airports, too, has a strong impact on the traceability index as shown in Figure 4. In this simulation, we simulated a set of 200 aircraft of the same type frequenting a variable number of airports at random for a year. After one year, the traceability index is 7.80% (corresponding to 15.60 aircraft) when aircraft frequent 200 airports against 0.785% (1.57 aircraft) for aircraft frequenting 100 airports.

Simply increasing the number of aircraft enrolled in the PIA program would effectively also increase the number of airports frequented by these diverse aircraft, so the overall result would not scale exactly as in Figure 3. Limiting the use of the PIA program to aircraft frequenting only a small set of airports would exclude a lot of aircraft, and may thus decrease the privacy of the system. Once a large number of aircraft are enrolled in this program, new aircraft joining it would most probably frequent the same airports as aircraft already enrolled.

We note that in the simulation all airports are statistically equally frequented, which is not the case in real world. Busy airports are likely to host many aircraft enrolled in the PIA program, which is good for aircraft privacy. Quiet airports are likely to host only a few or no aircraft enrolled in the PIA program, thus a change of PIA and call sign is rare. If an aircraft is the only one to change its PIA and call sign during its stay at the airport, the change is trivially traceable. Hence, busy airports offer a better privacy protection compared with quiet airports, but this is not considered by the simulator at this point.

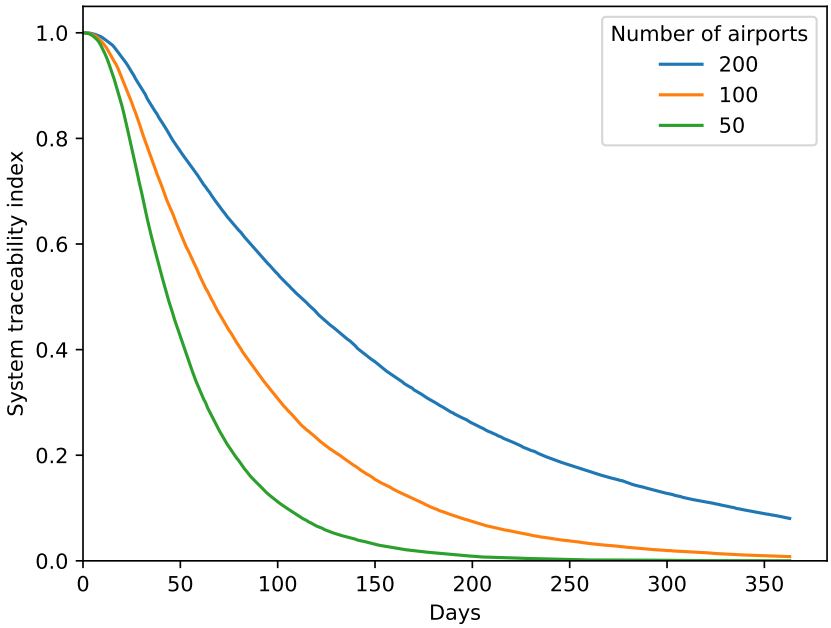

Fig. 4 System traceability index over time depending on the number of airports frequented. 200 aircraft, 1 aircraft type, 28 days PIA update frequency, 100 simulations.

4. PIA change frequency

Figure 5 shows the results of our simulation with 200 aircraft of a single type frequenting 100 airports, and variable PIA update frequency. It demonstrates that it is preferable to have a short PIA update frequency to minimize the system's traceability index. In the first phase of the PIA program, aircraft operators are allowed to request a new PIA 60 days after their previous assignment (we neglect the processing time). The graph shows that after one year, the traceability index is 30.7% (61.4 aircraft out of 200 are still being tracked) for this case, which is far from desirable. In contrast, the same result is reached after 101 days when the PIA update frequency is lowered to 28 days, corresponding to phase two. The same result is reached after only 37 days with a frequency of 10 days. This confirms our initial assumption that PIA and call sign should be changed as often as possible.

5. Types of aircraft

Using the aircraft type information in the methodology reduces the number of candidates after a new PIA is observed and increases the global traceability index of the system.

Each aircraft has its type chosen uniformly at random from the set of aircraft types. We suppose that the adversary is able to distinguish the type of aircraft by looking at the flight ADS-B transmissions as described in [9]. Figure 6 illustrates our simulation with 200 aircraft of 1, 3 and 7 different types, respectively. These aircraft frequent 100 airports and update their PIA every 28 days. It shows that an adversary being able to distinguish 7 types of aircraft is much more

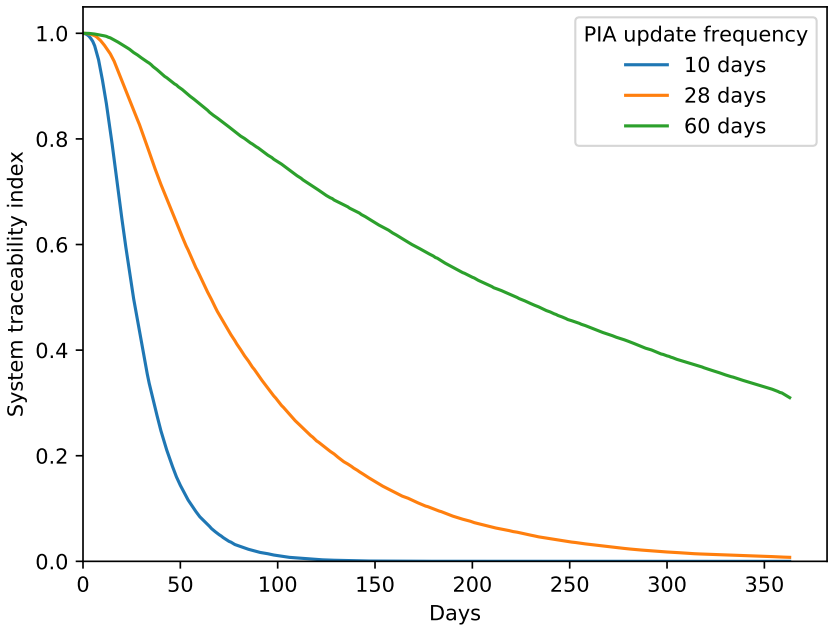

Fig. 5 System traceability index over time depending on the PIA and call sign simultaneous update frequency. 200 aircraft, 100 airports, 1 aircraft type, 100 simulations.

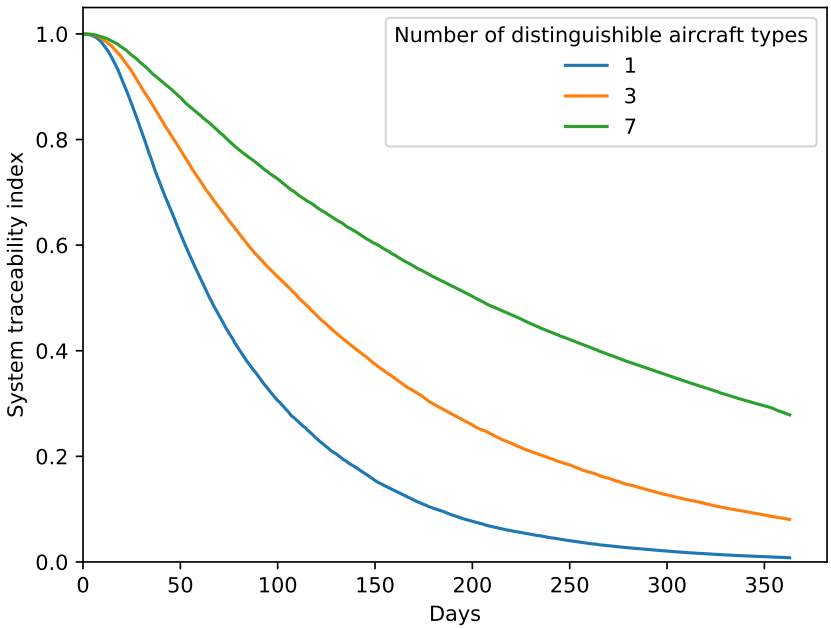

Fig. 6 System traceability index over time depending on the number of aircraft types an adversary is able to distinguish. 200 aircraft, 100 airports, 28 days PIA update frequency, 100 simulations.

powerful compared to an adversary unable to distinguish aircraft types at all. After 200 days, an adversary unable to distinguish aircraft types can track 7.7% of the aircraft, against 26.0%, 50.3% for adversaries able to distinguish 3 and 7 aircraft types, respectively.

C. Live tracking

Live tracking, the identification of an aircraft leaving an airport, with only historical data, is trivial if an aircraft keeps the same call sign when moving to a PIA. Live tracking is harder to perform if an aircraft changes both their PIA and call sign simultaneously.

Simulating live tracking from a simulators not feasible as the aim of the simulator is to track aircraft over time. The only parameters influencing a live tracking is the number of aircraft on ground at a given airport, when an aircraft using a new PIA and call sign departs from this airport. If we consider types of aircraft, we take into account only the aircraft of the same type as the one which departed on ground, before the aircraft's departure. Let this number of aircraft be k. Then, as mentioned above, as this mix provides *k-anonymity*, the probability of successful aircraft identification is $\frac{1}{k}$.

D. Limitations of the PIA Scheme

Best privacy practices for free route mixnets involve random routes inside the mixnet [17]. If the path taken inside the mixnet is not random, an adversary could establish correlations and follow elements in the mixnet easily. However, for real-world aircraft, flying to a random airport is nonsensical in terms of cost, time and pollution. Thus, it is impossible to get the same privacy level for General Aviation as for example in TOR [18]. The best theoretic privacy bound that can be reached by General Aviation would be to change both aircraft's PIA and call sign for each flight. This allows aircraft to be shuffled at each airport they stay at.

As each aircraft has its own flying patterns, each aircraft will have a different traceability index. Consider the extreme where an aircraft is the only one using a PIA at an airport and flies only from and to this single airport; it can never by anonymous. On the other extreme, an aircraft using a PIA and frequenting busy airports with lots of simultaneous PIA users, will have a much lower traceability index. If the operator of an aircraft using a PIA decides never to request a new PIA, it will obviously become traceable, and other aircraft traceability index will rise, as there will be less aircraft updating regularly their PIA address. For this reason, it is very difficult to provide strong anonymity guarantees for any specific aircraft.

Furthermore, the PIA privacy scheme is limited to General Aviation in the United States and to a certain type of ADS-B transponder, a single worldwide privacy scheme including all aircraft would be best for aircraft privacy.

Another limitation to the privacy General Aviation can aim for, is physical plane spotting. Anyone can take pictures of aircraft and upload them to a public platform, allowing an a wide audience to know aircraft's position. But this method cannot be automated yet, and is thus expensive. Aircraft spotting was already a threat to aircraft privacy before

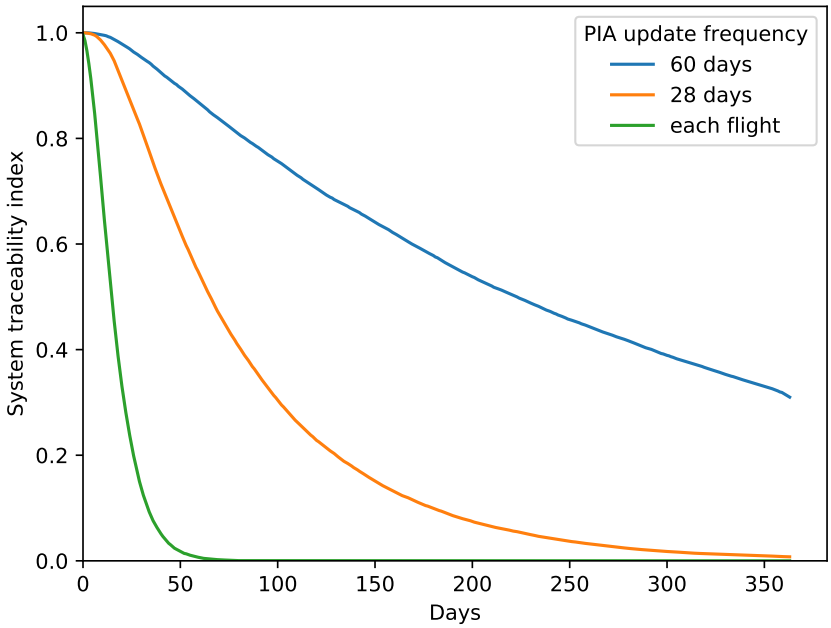

Fig. 7 System traceability index over time depending the PIA change policy. 200 aircraft, 100 airports, 1 aircraft type, 100 simulations.

ADS-B, and cannot be mitigated efficiently.

VI. Mitigations & Improvements

Our evaluation shows that the traceability performance of the PIA program is limited by the PIA change frequency imposed by the FAA. Furthermore, requesting a new PIA requires effort from the aircraft operators discouraging frequent change. Thus, an obvious improvement would be to automate the PIA change operation as much as possible. Another obvious approach to obtain maximum privacy is to broaden the user base and extend the system worldwide. As an added benefit, aircraft could fly abroad and would not become re-traceable with each international flight. Finally, all aircraft enrolled in such a program should use the same call sign space, preventing the possibility to split aircraft into smaller groups.

Based on these possible mitigations and the constraints imposed by the real-world PIA, we can analyse two straightforward solutions: changing the PIA with every new flight and changing all PIAs simultaneously. In the following, we gauge how much these tweaks could improve the privacy of the existing system in the US.

A. Change PIA for each flight

As stated before, the highest level of privacy can be reached only when aircraft change their PIA and call sign as often as possible. Hence, the policy providing the lowest traceability index to the aircraft is to make them change their

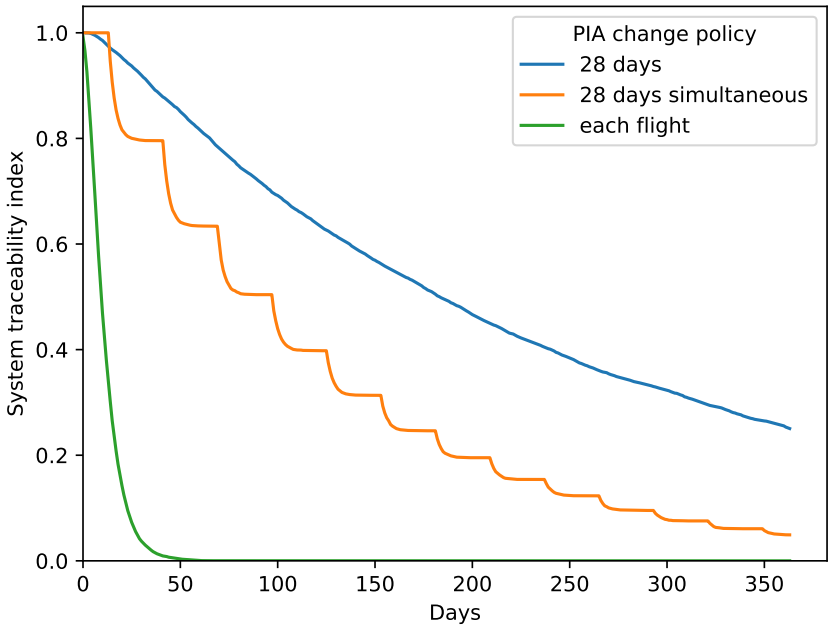

Fig. 8 System traceability index over time depending the PIA update policy. 100 aircraft, 100 airports, 1 aircraft type, 100 simulations.

PIA and call sign for each new flight. This way, all aircraft of the same type will be shuffled with other aircraft present at the same airport. Figure 7 shows the performance of this scheme compared with different PIA change frequencies. The simulation was run with 200 aircraft of a single type frequenting 100 different airports for one year. The green line is the maximal average privacy level that can be reached.

This method provides the best possible privacy performance without modifying aircraft trajectories, although it may yet be impractical for the administration to issue and track a new PIA for every flight. Indeed, the current PIA request procedure requires a test flight before the new PIA can successfully be used. While the FAA could delegate the responsibility of issuing new PIAs to the third-party call sign providers, this solution would require additional effort from aircraft operators, who would still need to manually perform the change before each flight.

B. Change all PIAs simultaneously

The second proposed scheme consists of periodically changing the PIAs of all enrolled aircraft simultaneously. This would provide a better privacy level for aircraft staying at an airport with other PIA aircraft at the time of the change. We note that Aircraft operators do not need to update the PIA at the exact same time but simply before their first flight after the deadline. As some aircraft may be in-flight at the time of change, they perform the change on ground before their next flight. The changeover can be scheduled at night in order to minimize the number of in-flight aircraft.

Figure 8 shows the expected performance of this policy with a 28 day simultaneous changeover. Although it is still

far from the best achievable privacy level, it provides a much lower traceability index compared with non-simultaneous random PIA changes. After 150 days, the 28 days policy traceability index is 56.9% whereas it is significantly lower at 31.3% if all aircraft change their PIA simultaneously.

This policy could be imposed by the FAA, but it would force all aircraft operators to periodically update their PIA and call sign, unless this process can be automated. This policy would also require the program to be modified, for instance removing the test flight, or considering the first flight of an aircraft after a change as its test flight. There would be no additional administrative work compared to the current policy, as there would be the same total number of PIA changes. If the FAA does not want to impose this policy, aircraft operator desiring more privacy could independently form a group and agree upon the dates they will change their PIA. Although this provides less privacy compared to the case where all aircraft are forced to participate, it is still an improvement over the current state of the program.

VII. Conclusion

This paper gives an assessment on the Privacy ICAO Address (PIA) program introduced by the FAA in NextGen, in its state of May 2020. We demonstrate that air traffic can be modelled as a variable-latency, free-route mixnet. In this mixnet, airports serve as mix stations, shuffling the aircraft identifiers. We then introduce a novel privacy metric for aircraft privacy: the *traceability index*. It assesses the privacy guarantees for aircraft and their passengers in the considered mixnets. We further illustrate how an adversary is able to track flights using publicly available data, even if an aircraft is enrolled in the PIA program.

As of May 2020, we observed only 16 aircraft enrolled in this program. We demonstrate that this number of aircraft is considerably too low to provide any privacy to the program's users and it thus does not improve General Aviation privacy. Further analysis demonstrates that the program requires significantly more enrolled aircraft to increase the global privacy level. But even if a large number of aircraft were to join the program, the privacy performance of the PIA program would remain weak compared with the maximum achievable performance. We finally suggested improvements to the PIA program, providing better privacy guarantees in General Aviation under real-world constrains and hope that our analysis serves to inform future discussion in this area.

References

[1] Aircraft Owners and Pilots Association, "WHERE IS ADS-B OUT REQUIRED?" `https://www.aopa.org/go-fly/aircraft-and-ownership/ads-b/where-is-ads-b-out-required`, 2020.

[2] OpenSky Network, `https://opensky-network.org`, 2020.

[3] FligthRadar24, `https://flightradar24.com`, 2020.

[4] Schäfer, M., Strohmeier, M., Smith, M., Fuchs, M., Lenders, V., Liechti, M., and Martinovic, I., "OpenSky report 2017: Mode

S and ADS-B usage of military and other state aircraft," *2017 IEEE/AIAA 36th Digital Avionics Systems Conference (DASC)*, IEEE, 2017, pp. 1–10.

[5] Strohmeier, M., Smith, M., Lenders, V., and Martinovic, I., "The Real First Class? Inferring Confidential Corporate Mergers and Government Relations from Air Traffic Communication," *2018 IEEE European Symposium on Security and Privacy (EuroS P)*, 2018, pp. 107–121. https://doi.org/10.1109/EuroSP.2018.00016.

[6] AOPA, "AOPA pushing ADS-B privacy forward potential 1090ES solution to test next year," https://www.aopa.org/news-and-media/all-news/2018/december/13/aopa-pushing-ads-b-privacy-forward, 2018.

[7] Federal Aviation Administration, "ADS-B Privacy," https://www.faa.gov/nextgen/equipadsb/privacy/, 2020.

[8] Michel, G., "ICAO – N-Number Converter," https://github.com/guillaumemichel/icao-nnumber_converter, 2020.

[9] Strohmeier, M., Smith, M., Lenders, V., and Martinovic, I., "Classi-Fly: Inferring Aircraft Categories from Open Data," *arXiv preprint arXiv:1908.01061*, 2019. URL https://arxiv.org/abs/1908.01061.

[10] Strohmeier, M., Lenders, V., and Martinovic, I., "On the Security of the Automatic Dependent Surveillance-Broadcast Protocol," *IEEE Communications Surveys & Tutorials*, Vol. 17, 2014, pp. 1066–1087. https://doi.org/10.1109/COMST.2014.2365951.

[11] "Freeflight Systems NextGen Avionics - ADS-B Solutions," https://web.archive.org/web/20141122083331/http://freeflightadsb.com/products/ads-b, 2014.

[12] Federal Aviation Administration, "Final Rule for ADS-B Out," https://www.govinfo.gov/content/pkg/FR-2010-05-28/pdf/2010-12645.pdf, 2010.

[13] Smith, M., Moser, D., Strohmeier, M., Lenders, V., and Martinovic, I., "Undermining Privacy in the Aircraft Communications Addressing and Reporting System (ACARS)," *Proceedings on Privacy Enhancing Technologies*, Vol. 2018, 2018, pp. 105–122. https://doi.org/10.1515/popets-2018-0023.

[14] Federal Aviation Administration, "Limiting Aircraft Data Displayed (LADD)," https://ladd.faa.gov/, 2019.

[15] Sampigethaya, K., Poovendran, R., and Taylor, C., "Privacy of general aviation aircraft in the NextGen," *AIAA/IEEE Digital Avionics Systems Conference - Proceedings*, 2012, pp. 7B5–1. https://doi.org/10.1109/DASC.2012.6382413.

[16] NBAA, "ADS-B Privacy FAQ," https://nbaa.org/aircraft-operations/security/privacy/ads-b-privacy-faq/, 2020.

[17] Chaum, D., "Untraceable electronic mail, return addresses, and digital pseudonyms," *Communications of the ACM*, Vol. 24, No. 2, 1981. https://doi.org/10.1145/358549.358563, https://www.chaum.com/publications/chaum-mix.pdf.

[18] Dingledine, R., Mathewson, N., and Syverson, P., "Tor: The Second-Generation Onion Router," *Paul Syverson*, Vol. 13, 2004. https://doi.org/10.5555/1251375.1251396.

[19] Michel, G., "ADS-B Privacy," https://github.com/guillaumemichel/ADS-B_Privacy, 2020.

[20] Michel, G., "Aircraft Privacy Simulator," https://github.com/guillaumemichel/aircraft-privacy-simulator, 2020.