Improved RFI Localization Through Aircraft Position Estimation During Losses in ADS-B Reception

Michael Dacus, Zixi Liu, Sherman Lo and Todd Walter
Stanford University

BIOGRAPHY

Michael Dacus is a M.S. student in the GPS Laboratory at Stanford University. He received his B.S. degree from Clemson University in 2020 in Mechanical Engineering.

Zixi Liu is a Ph.D. candidate at the GPS Laboratory at Stanford University. She received her B.S. degree from Purdue University in 2018 and her M.S. degree from Stanford University in 2020.

Sherman Lo is a senior research engineer at the GPS Laboratory at Stanford University.

Todd Walter is the Professor of Research and Director of the GPS laboratory at Stanford University.

ABSTRACT

The ability to estimate aircraft position and velocity is vital for safe and effective air traffic operations and management. With the emergence of the Automatic Dependent Surveillance-Broadcast (ADS-B) network, aircraft are now capable of sending GNSS positioning and operational data to ground stations to be used by regional air traffic controllers. In the research community, ADS-B has also been used for quickly detecting and even more recently, localizing and characterizing GNSS interference threats [1]. However, ADS-B does present some limitations that inhibit GNSS Interference, Detection and Localization (IDL) capability, specifically the loss of regular ADS-B derived positions. Reducing this loss of real-time aircraft positioning information using various interpolation approaches could help immensely with localization.

This paper uses historical crowdsourced ADS-B data to examine how well missing flight path data can be interpolated, and how interpolated flight paths can be used to localize interference sources. Crowdsourced data is retrieved from OpenSky Network, a client-based application that receives air traffic data from a network of ADS-B receivers around the globe [2]. ADS-B flight paths are processed, reconstructed, and split into various datasets based on the change in the Navigational Integrity Category (NIC) and presence of gaps in the flight paths. The two primary interpolation methods tested are a simple cubic spline and a MATLAB Systems Object, geoTrajectory. Both were applied to randomized trajectories and aircraft on a nominal test set to understand how well each performed compared to its original ADS-B position. Once interpolation criteria are established for each flight profile, the chosen interpolation method is then applied to a dataset with flight gaps indicating a drop in the NIC value (potential GNSS/RFI interference), as well as another dataset [with flight gaps], but no dramatic change in NIC (possibly beyond LOS). An estimated centroid location of interference is calculated using only the flight paths indicating possible GNSS/RFI interference.

To assess the performance of each interpolation method and how it might improve IDL efforts, data was selected over a period of 3 days at a reported interference event near Denver International Airport [3]. Based on the results from the various interpolation methods, the nominal data showed that both cubic spline and geoTrajectory had similar performance, and that more complex flight path geometries are harder to interpolate. Once the interpolation method is applied to the gapped flight datasets, predicted flight paths from data with indicated drops in the NIC value can localize the interference source with good accuracy compared to the original jamming source, justifying a high correlation between NIC value change and localizing the jamming source location. Furthermore, the percentage increase in data added from interpolation suggests that predicting gaps (especially during potential interference) adds a significant amount of missing flight path information, which can be useful in other IDL algorithms.

I. INTRODUCTION

The Automatic Dependent Surveillance-Broadcast (ADS-B) is a system that incorporates regular aircraft broadcast of its position information to provide aircraft surveillance for air traffic control and airspace users. For aircraft, it is supported by a series of receiver stations (primarily within the United States and Europe (EU)) that serve as a network to collect local GPS positioning, velocity and operational messages on the 1090 Mhz frequency. The ADS-B network has had an increasing role in air traffic and management operations within the National Airspace System (NAS) over the past decade. With requirements from both the Federal Aviation Administration (FAA: United States) and the European Union Aviation Safety Administration (EASA: EU) for all aircraft traveling within controlled airspace to have installed ADS-B Out equipment by 2020, ADS-B has already transitioned into a fully operational and available status across CONUS and Western Europe.

While ADS-B provides several unique advantages over traditional radar-based methods, there are some limitations, more specifically the possibility of losing reception to local ADS-B ground stations, resulting in gaps in the observed flight path trajectory. ADS-B reception loss can be due to a variety of reasons, whether it is from physical signal blockage (Figure 3b), poor receiver coverage, or from various forms interference and jamming (Figure 3c). The former is particularly true with ground based ADS-B receivers when deriving ADS-B information from crowdsourced sources (i.e. OpenSky network), as they do not get to choose their reception site locations nor have their antenna placed high on a cellular tower (as is the case for the FAA ADS-B ground based transceivers or GBTs). Some research (as noted in Section II.1) use ADS-B data to determine whether aircraft are affected by regional jamming sources (GNSS or other RFI), whereas other fields (noted in Section II.2) use the data to predict flight path trajectories. Regardless of the end purpose, state estimation with ADS-B is vital for not only accurate and reasonable prediction of air traffic, but to improve the efforts in localizing interference sources.

This paper seeks to interpolate existing flight path trajectories and attempt to estimate the jamming source from reported ADS-B messages over a period of 3 days (January 21, 22 and 23, 2022) surrounding Denver International Airport. Interference could be immediately detected using ADS-B information around January 21, 2100 UTC, indicated by an increased presence of ADS-B messages containing low NIC values. Additional ADS-B data was collected for January 22 and 23 around Denver International Airport to capture the impact of interference on departing and arriving traffic. This particular set of the data was selected due to indicated reports of GNSS interference from multiple pilots and aviation authorities [3]. After data processing, reconstruction and splitting, data was split into three subsets: a nominal set (no flight gaps), data with indicated potential GNSS interference, and data with aircraft flying beyond LOS with ground station. Flight gaps were also classified into three main profiles: straight-and-level, climb/descent, or a combination of maneuvers. The two interpolation methods used in this paper—a basic cubic spline, and geoTrajectory (a MATLAB Systems Object)—are compared using the nominal test set. Once interpolation criteria are established for each flight profile, the chosen interpolation method (cubic spline) is then applied to both the dataset including aircraft flight paths indicating potential interference and flight paths indicating beyond LOS from ground station. The location of the predicted flight paths are then compared and used to infer the location of the jamming source using a latitude-longitude centroidal calculation.

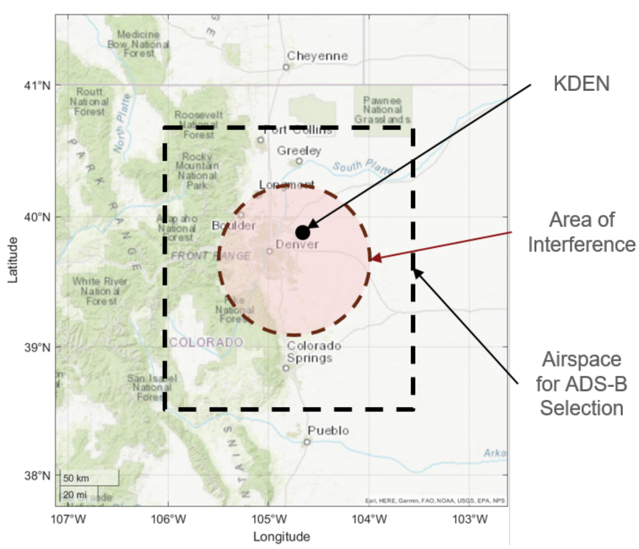

Figure 1: a general topographic map of approximate area of interference, and queried ADS-B data

Section II discusses various applications of ADS-B data in both state estimation and IDL fields. Section III discusses processing techniques for the incoming ADS-B data from the OpenSky Network. Section IV discusses the flight path interpolation methods and evaluation metrics used, and Section V shows validation results of each applied interpolation method, its application to selected set of interference data, and the estimated locations of the potential jamming source. Figure 2 is a flow chart of the methodologies used to filter, split, and interpolate the ADS-B data during reported interference.

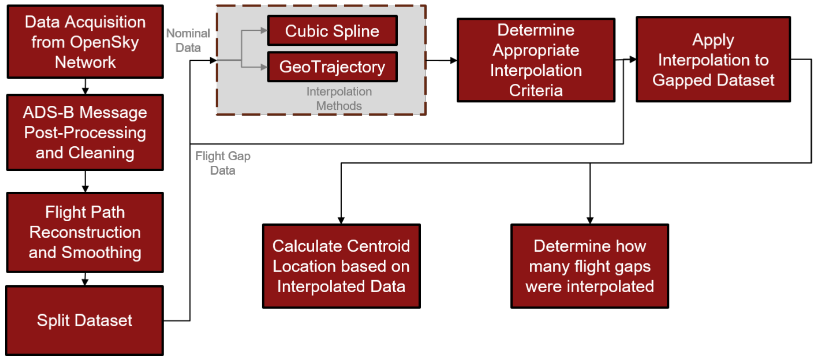

Figure 2: Flow Chart of ADS-B processing and interpolation

1. The Navigational Integrity Category (NIC)

While the 1090 MHz ADS-B used worldwide primarily provides position, velocity and other information, the 1090 MHz (Mode S ES) provides ADS-B position reports as well as a quality measure for the positioning solution, known as the Navigational Integrity Category (NIC). The NIC value is an integer contained within an ADS-B message that reflects the signal accuracy and quality of its associated positioning solution. A higher NIC (as shown in Figure 3a) corresponds with a smaller containment radius of the position estimate, while a low NIC value (or zero) corresponds to a large or even unknown error bound of the position [4]. To use the ADS-B reports in identifying and localizing GNSS interference, we need a quality measure of GNSS reception. We use the NIC as a proxy metric for the quality of GNSS reception as higher error levels generally means poorer (weaker) GNSS reception. Despite the fact that ADS-B messages do not contain signal strength information (such as a Carrier-to-noise ratio (C/N_0)), the NIC value is the only metric contained in the position reports that can be used to evaluate the integrity of a positioning solution.

Since the NIC value (and its relative proximity to missing gaps in ADS-B flight paths) is the only available performance metric transmitted within the ADS-B position message, it has been used to infer as to whether aircraft are beyond LOS or possibly impacted by various levels of interference. Referring back to Figures 3b and 3c, it has been noticed that aircraft that lose reception to ADS-B ground stations due to physical obstructions, poor receiver coverage, and/or are beyond LOS maintain a steady NIC value prior to losing signal. However, aircraft that enter known areas of interference have experienced a drop in NIC value prior losing ADS-B information. Aircraft can be also affected by communication interference, multipath, and other sources of noise. Other variables, such as aircraft dynamics (yaw, pitch, or roll, etc.) can also affect potential losses in signal reception (since it changes of the directionality of onboard antennas), but it is the aircraft's position that primarily affects reception loss.

II. RELATED WORK

In order to fully understand the nature of potential intermittent loss of ADS-B reception, it is important to provide a review on the state of both GNSS interference detection (using ADS-B) and state estimation research fields. Understanding how ADS-B data is used to detect GNSS interference sources may be helpful in learning how aircraft lose ADS-B reception during flight. Equally as important, a literature review of the state estimation field could provide an idea behind various interpolation methods, and how the field overall might be improved.

1. GNSS Interference Detection and Localization (IDL)

There has been new and extensive research into using ADS-B data for localizing and detecting interference sources. Darabseh et. al. takes a probabilistic approach by using ADS-B data to formulate distributions based on the Navigational Accuracy Category-Position (NACp) value, an accuracy metric versus the the NIC value which represents a containment radius [5]. The findings presented showed that analyzing changes in the NACp over time can be valuable in alerting the presence of interference, but the methodology does not address localizing the interference source, and must rely on accurate pilot reporting. Jonáš et. al. uses the gaps in aircraft paths and Power Difference of Arrival (PDOA) as a means of localizing a possible interference source [6]; While the method proposed included promising results, the real location of the jamming source was unconfirmed and could not validate the accuarcy of the PDOA method. Liu et. al formulates the center of the jamming area (i.e. predicted location of the jammer) as a convex optimization problem, and applies this approach to a reported interference event near Hayward Executive Airport in January 2019 [1]. The proposed method was able localize a region of airspace near the airport, but similar to [6], its true location of the jamming source was undetermined, and the interference itself was intermittent over a

NIC	Containment Radius (R_c)
0	Unknown
1	37.04 [km]
2	14.816 [km]
3	7.408 [km]
4	3.704 [m]
5	1852 [m]
6	1111.2 [m] 926 [m] 555.6 [m]
7	370.4 [m]
8	185.2 [m]
10	25 [m]
11	7.5 [m]

(a)

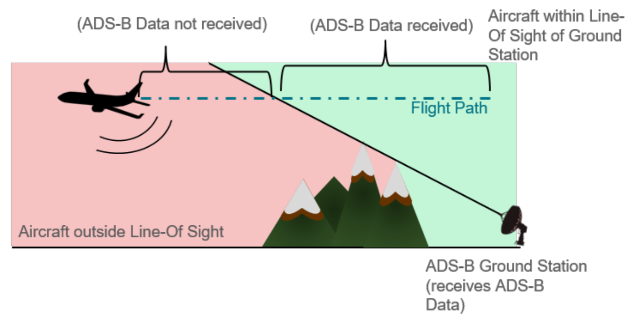

(b)

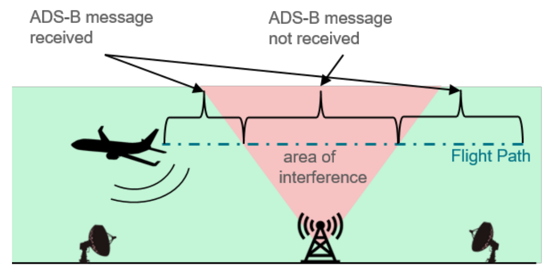

(c)

Figure 3: 3a Navigational Integrity Category (NIC) values and their corresponding containment radii [4], 3b Illustration of aircraft beyond Line-Of-Sight (LOS) with ground station, and 3c Illustration of aircraft affected by local jamming sources

period of a month, which made localization difficult to achieve.

2. Aircraft Flight Path Interpolation Techniques

In the field of state estimation and specifically aircraft trajectory prediction, for short-term intermittent gaps in flight data, many use a localized regression method for filling in missing data. Weinert et. al. implements a piecewise cubic interpolation on equally timed interval data with missing points [7]. Wu et. al also uses a cubic spline interpolation as a preprocessing technique prior to training on a neural network for 4D Trajectory prediction [8]. In both papers and others that use this localized approach, cubic interpolation of the missing flight path is reasonable for a short duration of ADS-B reception loss, but is unknown as how long of reception loss allows such a method to be reasonable.

Other methods (similar to Wu et. al [8]), introduce more data-driven approaches to flight path interpolation. The most popular modeling approach is incorporating some variation of a long short-term network (LSTM). Zeng et. al uses a sequence-to-sequence deep long short term memory network (SS-DLSTM) to prediction future trajectory based on historical aircraft trajectories, showing a higher robustness of flight path prediction for longer trajectories [9]. Xu et. al uses a social-LSTM model (S-LSTM) similar to [9], but instead incorporates a pooling layer in order to learn off of adjacent aircraft flight paths for various aircraft simultaneously [10]. Finally, Sahadevan et. al uses bi-directional LSTM (Bi-LSTM) to interpolate flight paths based on data trained in opposite temporal directions for similar flight paths [11]. Other papers use a combination of Deep Neural Networks [12] with similar results to LSTM models (but with longer computational time). Because of the longer computational time, using interpolation approaches that rely on the existing position information (cubic-spline) are desirable since the ultimate goal is to use these flight paths to localize a jamming location.

III. DATA PROCESSING

1. OpenSky Post-Processing

ADS-B data is gathered from the OpenSky Network, an open-source, ground-based ADS-B receiver network that gathers aircraft positioning from commercial and private aircraft around the world. It not only collects aircraft positioning information, but also reported velocity (grounds speed, vertical rate), course heading, and the operational status of the ADS-B receiver onboard the aircraft [2]. As part of the API, Opensky decodes the original raw message per the Minimum Operational Performance Standards (MOPS) [4]. For the purposes of this paper, only the aircraft's positioning information (latitude, longitude, altitude), ICAO number (Identifier), timestamp of message received at the ground station, and NIC value are needed. The interpolation methods in Section IV can also incorporate velocity information, but for simplicity, positioning is only used in the interpolation process.

After the ADS-B data is retrieved from OpenSky, it must be pre-processed prior to usage within the various interpolation methods. First, duplicate ADS-B messages [with the same timestamp and position fix] are removed from the dataset, in addition to messages that include inconsistent receiver timestamps (a particular ground station that indicates a timestamp not in sync with other ground stations). An inconsistent receiver clock is similar to an incorrect position fix, so the corresponding ADS-B messages should be removed. Finally, messages with incomplete information are also removed.

2. Flight Path Reconstruction

As part of data processing, the flight paths for each aircraft transitioning through the given airspace are reconstructed into equally spaced-intervals of one second. ADS-B messages are reported on average every half of second, but are rarely reported to this exact interval; flight path reconstruction to equally spaced intervals smooths the flight path trajectory, reduces the size of the dataset due to redundant position reports, and removes irregularities due to possible noisy ADS-B positioning solutions. Flight path reconstruction follows the methodology as Zeng et. al [9] by formulating it as a linear least squares with regularization optimization problem, as shown below:

$$\underset{Y}{\operatorname{argmin}}\{||AY - Y_0||_2^2 + \lambda||\Gamma Y||_2^2\} \tag{1}$$

where Y_0 is the original ADS-B positions (unequal time intervals), A is a pre-constructed design matrix, Y is equally spaced trajectory that is being solved, Γ is the Tikonov regularization matrix and λ is the weighting parameter. The first term describes the measurement of the fit, and the second term adjusts the amount of smoothing (via λ) involved.

3. Data Splitting and Classification

a) Data Splitting by Change in NIC

With ADS-B flight paths processed, the flight trajectories (across the 3 days) are divided into 3 separate datasets prior to interpolation. A nominal dataset is generated from January 21 flight paths around Denver, which contain flight paths indicated at or above NIC of 8 and do not contain any flight gaps. As the name suggests, this dataset is representative of aircraft not susceptible to interference, reception loss due to possible signal blockage, or other possible sources.

During reported interference (after 2100 UTC), two other datasets are created; the first subset contains aircraft flight paths that have at least one flight gap in its trajectory and has an indicated drop in NIC prior to (or after) the gap, which—as noted in Section I—could be an indication that an aircraft is interfered by local jamming sources. The final subset contains flight paths that do not indicate a drop in NIC prior to or after losing signal, which could represent an aircraft being affected by signal blockage, poor receiver coverage, or other sources, also noted in Section I. As such, the purpose for splitting the dataset based on NIC value change is study the relationship between the location of losses in ADS-B position, and how that relates to the change in the NIC value seen in reported messages. Due to the size and number of flight paths, the datasets during reported interference are viewed over two time frames, Friday evening when the jamming started (January 21, 2100 UTC-January 22, 0659 UTC) and Saturday, local time (January 22, 0700 UTC-January 23, 0659 UTC).

b) Flight Profile Classification

In addition to splitting the datasets by presence of gaps and whether there is a change in NIC value, all flight gaps are categorized into three flight profiles: straight-level, climb or descent, or a combination of maneuvers. The primary reason for categorizing by flight profile is to determine the quantity of various flight path trajectories in the datasets, and to determine if there is a relationship between the flight profile and Euclidean Error between the true and predicted position.

Each of the three flight profiles follow a different set of filtering criteria, from most (straight and level) to least constraining (multiple maneuvers). In order for a flight gap to be straight and level, the aircraft must not deviate more $\pm 10°$ in heading and $\pm$ 30.5 m in altitude (or 100 ft). For a flight gap to be climbing or descending, the aircraft must maintain the same heading criteria as the straight-and-level flight, but loses the altitude restriction. Finally, for a flight gap to a be combination of maneuvers (i.e. ascending/descending and turning), it losses both criteria and serves as a “general/otherwise” category.

IV. FLIGHT PATH INTERPOLATION METHODS AND EVALUATION METRICS

The interpolation approaches listed in this section are implemented using only the positioning information of a specific aircraft and its current flight path.

1. Cubic Spline Interpolation

As noted in Section II.2, the cubic spline interpolation is a popular data filling method for short-term gaps in ADS-B flight paths. As the name suggests, it fits a third-order piecewise polynomial function to a flight path using local positioning information before and after losing ADS-B reception, which represented by the following function:

$$x(t) = \begin{cases} S_1(t) = a_1t^3 + b_1t^2 + c_1t + d_1 & if\ t \in [t_1, t_2] \\ S_2(t) = a_2t^3 + b_2t^2 + c_2t + d_2 & if\ t \in [t_2, t_3] \\ \vdots & \\ S_{n-1}(t) = a_nt^3 + b_nt^2 + c_nt + d_n & if\ t \in [t_{n-1}, t_n] \end{cases} \tag{2}$$

where n is the number of reported positions within the flight path, $\{a, b, c, d\}$ are the the fitted parameters, t is the timestamp, and $x(t)$ is the reported position. Since the aircraft's position is 3D coordinates (latitude, longitude, and altitude), 3 unique cubic splines (with respect to time) are generated between reported timestamps. Despite the function being piecewise, it must also be continuous and differentiable along the entire length of the flight path, all of which is represented mathematically below:

$$\begin{gathered} S_i(t_i) = y_i,\ S_i(t_{i+1}) = y_{i+1}\ \textit{for i=1...n-1} \\ S'_{i-1}(t_i) = S'_i(t_i)\ \textit{for i=2...n-1} \\ S''_{i-1}(t_i) = S''_i(t_i)\ \textit{for i=2...n-1} \end{gathered} \tag{3}$$

Also, the beginning and the end of the cubic spline must be straight-lines (double derivatives at endpoints are zero), which can be represented below:

$$\begin{gathered} S''_1(t_1) = 0 \\ S''_{n-1}(t_n) = 0 \end{gathered} \tag{4}$$

Therefore, for each of the piecewise cubic functions, a sufficient amount of adjacent points needs to be provided in order to create a realistic flight path. For the nominal dataset, at least 20 points are used (prior to and after) for each artificial flight gap. However, when the interpolation method is applied to interfered datasets, the entire flight path is used.

2. GeoTrajectory Interpolation

While the cubic spline interpolation approach relies on nonlinear regression of third order polynomial, geoTrajectory is a proposed localized approach that interpolates data using a physics-based model [13]. GeoTrajectory is a MATLAB Systems Object that determines the trajectory of any object based on a given set of Geodetic waypoints. Based on the given set of coordinates and additional features ("AutoBank" and "AutoPitch" properties), it determines a suitable flight path, assuming a standard rate of turn and pitch flight dynamics, which the majority of commercial aircraft follow.

3. Evaluation Metrics

For the nominal test data, in order to evaluate how well each interpolation method predicts the flight path (compared to the original ADS-B positions), 100 aircraft were randomly chosen from this test set and used by creating random gaps in the data (make an "artifical gap" in the data) from 30 seconds to 4 minutes long. Using these gaps, we compare the predicted trajectory of flight gap to the ADS-B reported position, and caluclate the average Euclidean Error for each flight gap, as shown below:

$$EE_i = \frac{1}{L_i} \sum_{t=1}^{L_i} \sqrt{(x_1^t - \hat{x_1}^t)^2 + (x_2^t - \hat{x_2}^t)^2 + (x_3^t - \hat{x_3}^t)}\text{, for the ith aircraft} \tag{5}$$

where L_i is the duration of the flight gap (in seconds), $\{x_1^t, x_2^t, x_3^t\}$ is the ADS-B reported position at timestamp t, and $\{\hat{x_1}^t, \hat{x_2}^t, \hat{x_3}^t\}$ is predicted position using the specified interpolation. The Euclidean Error distributions (represented as a series of box and whisker plots) are then compared to the NIC value containment radius (Table 3a) to relate the containment radius to deviations from interpolated flight paths.

For ADS-B datasets timestamped during reported interference, specifically the subset of data that indicates potential GNSS interference (drop/increase in NIC around flight gap), the primary evaluation metric is a comparison between the number of ADS-B flight hours with a NIC value of 0 to the number of flight hours that were interpolated from the dataset. Since the flight path trajectories from this dataset would indicate NIC=0 prior to or after losing ADS-B position information, it can be assumed that the predicted flight paths of gaps in ADS-B data would also indicate a poor/weak signal (i.e. NIC=0). This metric (expressed as a percent increase) will give an idea as to how much data (indicating a poor/degraded signal, possibly due to interference) is being added; the idea is to predict as many gaps in data as possible without creating unrealistic trajectories, which would be bounded by the criteria established by the results from the nominal dataset.

V. RESULTS AND DISCUSSION

The following results are a combination of a) a comparison between interpolation methods on nominal dataset, b) the interpolation methods applied to flights that experienced a loss in ADS-B information near Denver International, and c) how much data was added (from interpolation) to the original reported ADS-B messages.

1. Flight Path Interpolation Comparison

Figures 4a, 4b, and 4c represent the Euclidean Error results between cubic spline/geoTrajectory interpolations and the original ADS-B reported position for a random sample of 100 aircraft at various flight gap durations. The box edges in each of the box and whisker plot distributions represent the 25th and 75 percentile respectively, and the center line represents the empirical median. Whiskers extend to a length beyond the interquartile range (IQR), and any other values outside the whiskers are indicated as an outlier ("+" symbol).

As shown in the figures, predicted flight paths that fly straight and level will have the smallest deviation relative to the reported position, followed by aircraft that are climbing/descending or are performing multiple maneuvers (largest amount of deviation). This result is expected since straight and level flight is the simplest (and least complex) flight path geometry; aircraft are harder to predict in flight if there is a change in alitude and/or heading. If we use the NIC=7 containment radius as the upper threshold for determining how much we can interpolate, it can be noticed that all flight gaps that are straight and level can be safely predicted up to the tested 4 minutes. For flight gaps that are classified as climbing/descending, it was decided that the gaps up to 180 seconds can be predicted since the distribution is below the threshold. Finally, for flight gaps that are characterized as a combination of maneuvers, it was decided that gaps up to 90 seconds can be predicted. While these requirements are fairly loose and easily adjustable, they will be helpful in filtering out longer trajectories that are not as easy to predict.

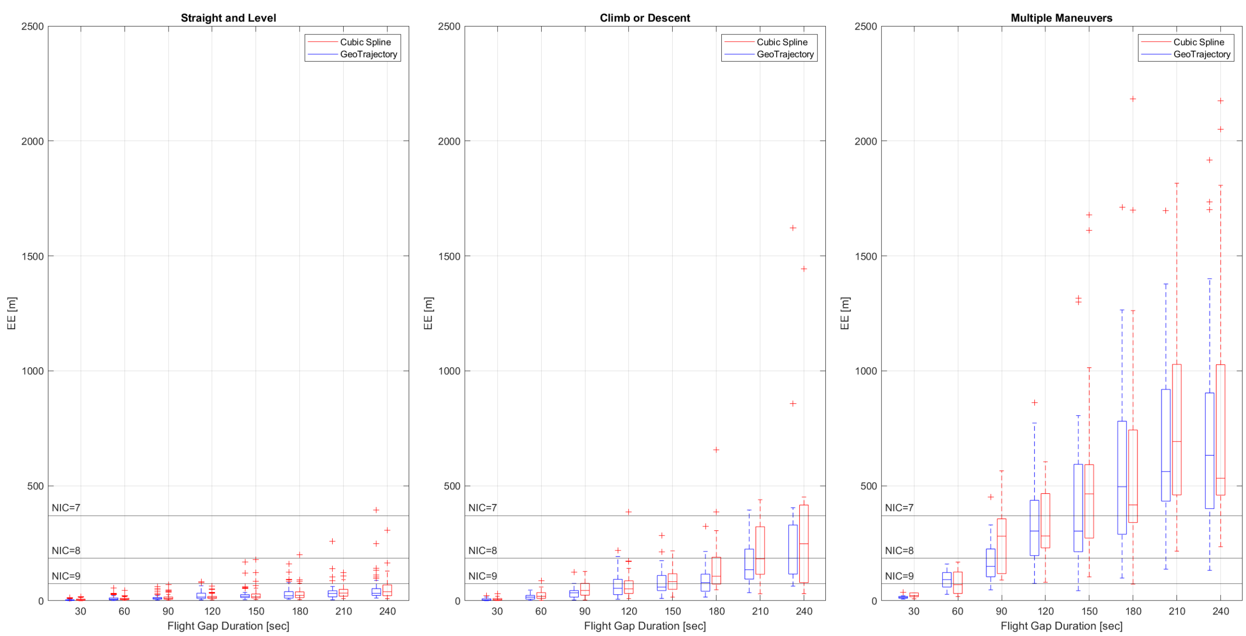

(a) **(b)** **(c)**

Figure 4: Euclidean Error distribution results between cubic spline and geoTrajectory results for various flight gaps durations, categorized by Straight and Level 4a, Climb/Descent 4b, or a combination of maneuvers 4c

It is also important to note that geoTrajectory and cubic spline Interpolations performed similarly across a variety of flight gap durations and profiles. Since they have the same relative variance throughout the results, cubic spline was chosen to predict aircraft flight paths during reported interference at Denver International.

2. Interpolation and RFI Localization Results

With the interpolation methods and data categorization methods outlined, the ADS-B flight paths flying near Denver International containing "gaps" can be processed and interpolated. To evaluate the data more efficiently, the interference subsets were further divided based on time of day: The first set applies interpolation to the evening of first reported interpolation (Friday evening local time), whereas the second set applies interpolation over a full day (Saturday, local time).

a) January 21, 0-2100 UTC

To gather a comparison of flight gaps in both nominal and potential interference scenarios, Figure 5 represents trajectories that had some loss in ADS-B reception during the aircraft flights prior to interference. Almost all flight paths (as represented in this figure) are flight paths that did not indicate a drop in NIC prior to or after losing ADS-B reception. Based on the location

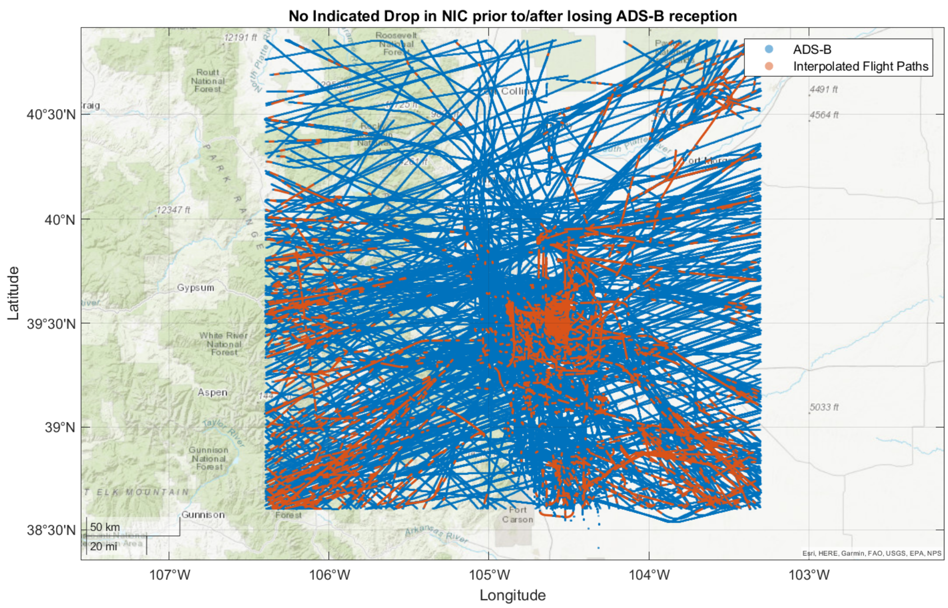

Figure 5: ADS-B reported position around Denver International (KDEN) prior to reported interference (2100 UTC)

of the interpolated flight paths, the majority of the flight gaps are located in high terrain and rural regions where aircraft could be susceptible to physical signal blockage (due to mountain features, etc.), as well as poor receiver coverage. As discussed in Section I and specifically about NIC values, the lack of change in the NIC and the location of these flight gaps suggests these aircraft may be beyond Line-Of-Sight from ground stations, but could also be affected by multipath and communication interference sources.

b) January 21, 2100-January 22, 0659 UTC

Figures 6a and 6b represent applying the cubic spline interpolation to flight gaps during the first 10 hours of reported interference, which would have been Friday evening, local time. Similar to flight gaps predicted in Figure 5, Figure 6a shows the majority of the flight gaps predicted in high terrain and rural areas, again suggesting that aircraft are susceptible to physical signal blockage and/or poor receiver coverage. However, as noted in Figure 6b, interpolated aircraft trajectories from data that indicate a change in NIC value prior to or after losing ADS-B position were located within the same area of interference visualized in Figure 1, suggesting that a change in NIC among those ADS-B messages is loosely related to some potential local jamming. Furthermore, the interpolated trajectories (as indicated by an orange marker) are used to calculate a centroid location of the jammer, which is indicated by the black star. From further analysis, it was determined that this calculated location was roughly 5-6 nm from the true jamming location.

c) January 22, 0700-January 23, 0659 UTC

Similar to the interpolation results determined from the first 10 hours of interference, ADS-B flight data over a full day saw similar results; Figure 7a also saw the majority of interpolated flight paths in high terrain and rural areas, as well routes most notable for arrival procedures in KDEN (northeast of the airport). As discussed previously, aircraft may suddenly lose position due to moderate banking turns, which may be a possible cause of losing positioning information along arrival procedure routes. Figure 7b shows the interpolated flight paths predicted of arriving and departing traffic over the same 24 hour period. Similar

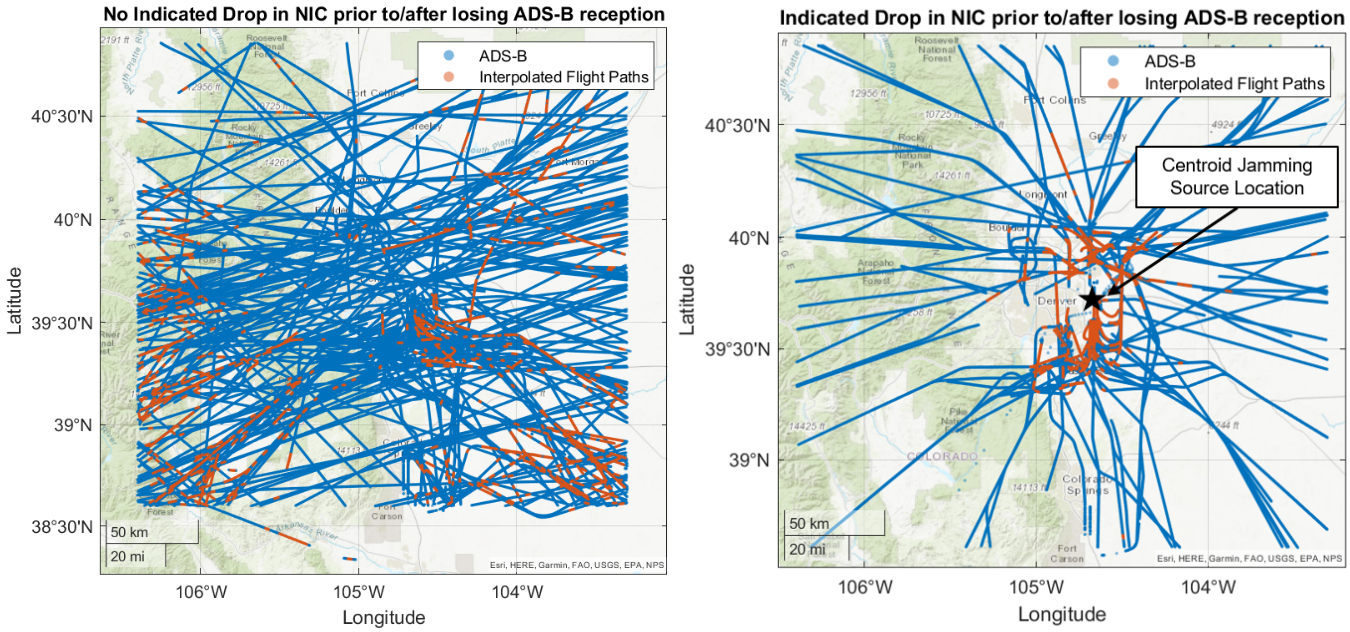

(a) (b)

Figure 6: ADS-B reported position and predicted interpolated trajectories during flight gaps split between no indicated 6a and indicated drop in NIC value 6b in airspace surrounding Denver International

to the results shown in Figure 6b, this data also has the interpolated flight paths congregated within the same approximate area of interference, reaffirming that data indicating a change in NIC is affected by the aircraft's location and presence of a possible jammer. A centroid calculation of the interpolated flight paths also shows a similar location as the centroid calculated using the first 10 hours of ADS-B flight paths, which also approximated 5 nm from the true jamming location.

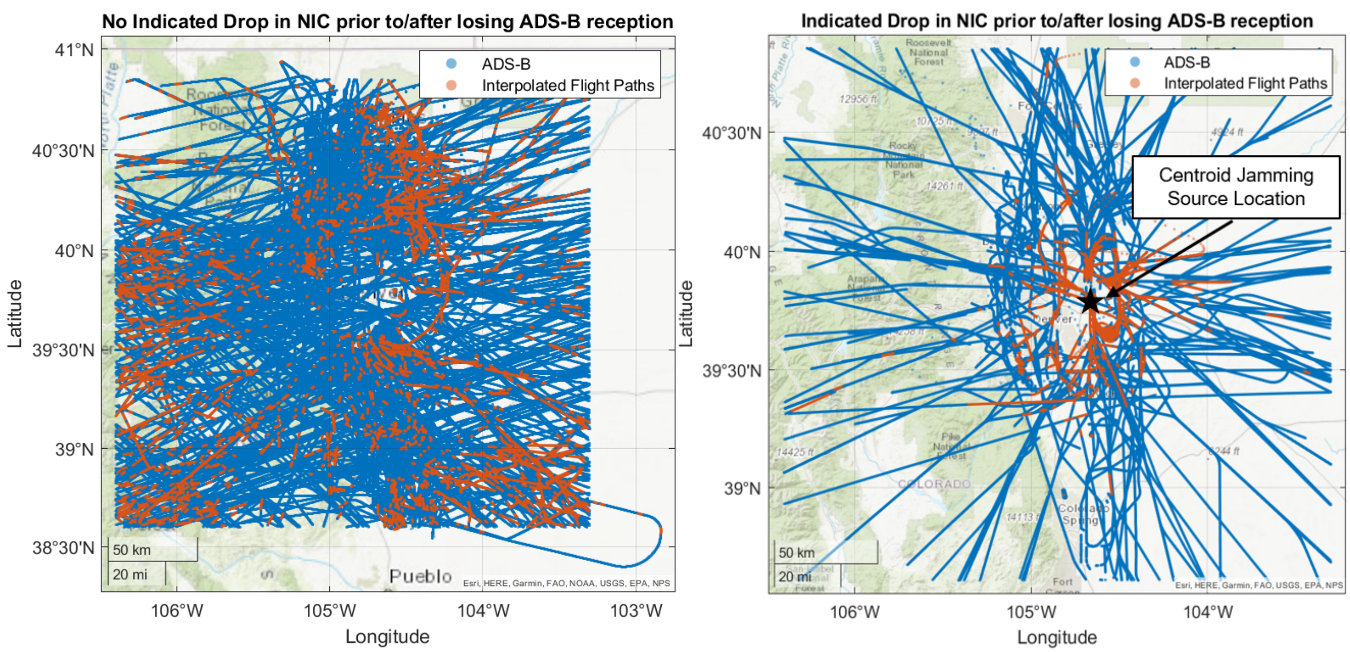

(a) (b)

Figure 7: ADS-B reported position and predicted interpolated trajectories during flight gaps split between no indicated 7a and indicated drop in NIC value 7b in airspace surrounding Denver International

3. Interpolation Improvement

Table 1 shows the results across each analyzed window of reported interference around Denver International. Despite the

	Jan 21, 2100-Jan 22, 0659 UTC	Jan 22, 0700-Jan 23, 0659 UTC
Duration of ADS-B Data indicating NIC=0 [hr]	4.95	11.65
Data points added from Interpolation [hr]	2.22	4.45
Percent Increase of NIC=0 from Interpolated Data	44.69%	38.21%

Table 1: Amount of flight path information added to ADS-B data, during interference, relative to the original amount of indicated interfered data

differences in the time windows, both add roughly 35-45 percent of interpolated data to the orignal size of ADS-B indicating a poor/degraded signal. In total over the 34 hours of reported interference, adding almost 7 hours of interpolated data was added to the original reported ADS-B messages, and the interpolated paths by themselves have been able to localize the source with good accuracy. Based on both the size of data added from interpolation and how this data specifically was able to localize the jamming source with a simple averaging solution, filling in gaps in flight data with indicated changes in NIC value will be beneficial for our IDL algorithms [14].

VI. CONCLUSION

Based on the predicted flight paths produced from two days of reported interference, interpolation of short-term flight gaps is beneficial for not only adding location information for missing aircraft, but more importantly be used to localize potential jamming sources. Based on noticeable trends that arise from reviewing ADS-B flight paths, the change in NIC value and its relationship with gaps in flight data can suggest the possibility of aircraft being affected by GNSS interference. By testing this trend on the reported interference, we were able to use the predicted flight paths of missing data (that also included indicated drops in NIC) to localize the jamming source down to approximately 5 nm from the original source. Based on localization results, it further substantiates the claim that indicated drops in NIC around missing ADS-B data within a flight path is correlated to the potential presence of RFI/GNSS ground-based interference. Furthermore, predicted flight paths generated from data not indicating a drop in NIC value highlights regions of airspace possibly susceptible to aircraft beyond Line-Of-Sight due to physical obstructions, poor receiver coverage, or other sources such as multipath and comms interference. From these conclusions, filling in missing gaps of ADS-B flight paths with realistic criteria can add a significant amount of data, and also improve the ability to localize jamming sources.

1. Future Work/Directions

Beyond testing a variety of interpolation process, future work will be concentrated on formulating a filtering process that detects erroneous ADS-B positioning data, which can be caused by a variety of reasons as discussed in Section I. The majority of the erroneous reported positions are caused by faulty or obsolete ADS-B equipment, but it can also arise from potential GNSS spoofing attacks. Erroneous ADS-B positioning data was noticed by several aircraft over the three day interference event at Denver International, resulting in creating unrealistic trajectory predictions. This filtering process could rely on a variety of approaches, either via as an outlier detection method or as a data-driven process. This process will rely on a number of metrics, not just reported NIC value, but the aircraft's dynamic variables (velocity and acceleration) and other aircraft information (tail number, etc.). Regardless of the method or metrics chosen, this will help further remove erroneous data prior to interpolation and localization.

VII. ACKNOWLEDGEMENTS

The authors would like to thank the FAA for sponsoring this research. The authors would also like to thank the OpenSky Network for using their ADS-B data.

REFERENCES

[1] Z. Liu, S. Lo, and T. Walter, "Gnss interference characterization and localization using opensky ads-b data," *Proceedings*, vol. 59, no. 1, 2020. [Online]. Available: https://www.mdpi.com/2504-3900/59/1/10

[2] M. Schäfer, 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*, 2014, pp. 83–94.

[3] *Aviation Cyber Initial PNT Meeting Briefing*, Office of the Secretary, Office of the Assistant Secretary for Research and Technology (OST-R), U.S. Department of Transportation, Washington, DC, USA, 2022.

[4] D.-B. R. (Firm), "Minimum operational performance standards (mops) for 1090 mhz extended squitter automatic dependent surveillance-broadcast (ads-b) and traffic information services-broadcast (tis-b)," vol. 59, no. Washington, DC, USA: RTCA, 2011.

[5] A. Darabseh, E. Bitsikas, and B. Tedongmo, "Detecting gps jamming incidents in opensky data," *Proceedings of the 7th OpenSky Workshop 2019*, vol. 67, pp. 97–108, April 2014.

[6] P. Jonáš and V. Vitan, "Detection and localization of gnss radio interference using ads-b data," in *2019 International Conference on Military Technologies (ICMT)*, 2019, pp. 1–5.

[7] A. Weinert, N. Underhill, B. Gill, and A. Wicks, "Processing of crowdsourced observations of aircraft in a high performance computing environment," in *2020 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, Sep. 2020, pp. 1–6.

[8] Z.-J. Wu, S. Tian, and L. Ma, "A 4d trajectory prediction model based on the bp neural network," *Journal of Intelligent Systems*, vol. 29, no. 1, pp. 1545–1557, 2020. [Online]. Available: https://doi.org/10.1515/jisys-2019-0077

[9] W. Zeng, Z. Quan, Z. Zhao, C. Xie, and X. Lu, "A deep learning approach for aircraft trajectory prediction in terminal airspace," *IEEE Access*, vol. 8, pp. 151 250–151 266, 2020.

[10] Z. Xu, W. Zeng, X. Chu, and P. Cao, "Multi-aircraft trajectory collaborative prediction based on social long short-term memory network," *Aerospace*, vol. 8, no. 4, 2021. [Online]. Available: https://www.mdpi.com/2226-4310/8/4/115

[11] D. Sahadevan, H. P.M., P. Ponnusamy, V. Gopi, and M. K. Nelli, "Ground-based 4d trajectory prediction using bi-directional lstm networks," *Applied Intelligence*, 2022.

[12] D. Adesina, O. Adagunodo, X. Dong, and L. Qian, "Aircraft location prediction using deep learning," in *MILCOM 2019 - 2019 IEEE Military Communications Conference (MILCOM)*, 2019, pp. 127–132.

[13] MathWorks, "geotrajectory: waypoint trajectory in geodetic coordinates," *MATLAB: Sensor Fusion and Tracking Toolbox: User's Guide (R2022a)*, 2012.

[14] Z. Liu, S. Lo, T. Walter, and J. Blanch, "Real-time detection and localization of gnss interference source," in *Proceedings of the 2022 International Technical Meeting of The Institute of Navigation*, 2022.