Demonstrating and Improving the Performance of a GNSS Dual Polarization Antenna (DPA) for Spoof Detection in Flight

Sherman Lo, *Stanford University*, Yu Hsuan Chen, *Stanford University*, Fabian Rothmaier, *Airbus,* Todd Walter, *Stanford University*

BIOGRAPHY (IES)

Sherman Lo is a senior research engineer at the Stanford GPS Laboratory and the executive director of the Stanford Center for Position Navigation and Time. He received his Ph.D. in Aeronautics and Astronautics from Stanford University in 2002. He has and continues to work on navigation robustness and safety, often supporting the FAA. He has conducted research on Loran, alternative navigation, SBAS, ARAIM, GNSS for railways and automobile. He also works on spoof and interference mitigation for navigation. He has published over 100 research papers and articles. He was awarded the ION Early Achievement Award.

Yu-Hsuan Chen is a research engineer at the Stanford GPS Laboratory. He received his Ph.D. in Electrical Engineering from National Cheng Kung University, Taiwan in 2011.

Fabian Rothmaier graduated from the Stanford GPS Laboratory in 2021 with his Ph.D. in Aeronautics and Astronautics at Stanford University. He currently works for Airbus.

Todd Walter is professor of research in the Department of Aeronautics & Astronautics. He directs at the Stanford GPS Laboratory. He received his Ph.D. in Applied Physics from Stanford University.

INTRODUCTION

Dual polarization antennas (DPAs) are being developed to provide robust GNSS spoof detection and interference mitigation. These antennas show great potential in determining the azimuthal direction of arrival (DOA) of incoming GNSS signals with just a single GNSS patch antenna. With satellite signal DOA, we can test for and detect non-genuine or spoofed signal. Under nominal conditions, it can provide static heading with as few as one GNSS satellite. The Stanford DPA prototypes have demonstrated an ability to find DOA from both genuine and spoofed signals [1][2][3]. The DPA provides right hand circularly polarized (RHCP) and left hand circularly polarized (LHCP) signal components that can be used to determine signal DOA, detect spoofing, or mitigate interference through directional nulls. Improved versions of the Stanford DPA have the ability to support high rate measurements, operate in dynamic situations and improved spoof detection algorithm that can utilize other detection metrics.

In this paper, we examine several means to improve performance and spoof detection with a DPA. It first discusses the design of the improved Stanford printed circuit board (PCB) based DPA [1] [4] to make is suitable flight trials. The design was modified to support high dynamics which means supporting a faster DOA determination capability than our first generation PCB DPA. Additionally, this DPA PCB supports a processing chains using a commercial off the shelf (COTS) GNSS receiver. Next, the paper examines the dynamic performance of the COTS GNSS DPA in flight test. It quantifies the factors that affect DOA performance such as satellite used, elevation angle and other factors. Accurate knowledge and incorporation of these factors can improve our DOA based spoof detection algorithm. It also examines performance in different flights conditions such as nominal and interfered. We will examine the resulting performance of the signal for spoof detection using algorithms developed to utilize DOA along with other detection metrics, specifically pseudo range residuals (prr). Finally, we examine the improvement in spoof detection if we are able to resolve the 180-degree ambiguity of DPA DOA measurements. Our analysis shows that resolving the ambiguity can greatly strengthen our confidence of spoof detection.

BACKGROUND

Dual Polarization Antenna and Direction of Arrival

GNSS dual polarization antennas are designed to provide the reception of two different polarization, typically right and left hand circular polarized (RHCP, LHCP, respectively), simultaneously. DPAs have been of interest to GNSS for some time and have been proposed for use to mitigate interference, particularly as part of a controlled reception pattern array (CRPA) where they can provide additional degrees of freedom (dof) without increasing number of physical elements or an increase in size [5][6][7]. These antennas are also attractive because they are reasonably simple being essentially a standard patch. However use of DPAs for GNSS DOA determination is relatively new [8][9].

Under appropriate conditions, the DOA of a signal can be derived from the reception of its RHCP and LHCP components [9]. While GNSS signals are typically RHCP whose direct reception would not provide any DOA information, any signal impacting a ground plane prior to entering the antenna will become linearly polarized due to boundary conditions. The important condition is that the ground plane is large enough. With the Stanford DPA, our prior experiments showed that a 1-foot square ground plane for good polarized components. The signal, linearized on the ground plane, can be separated into RHCP and LHCP components. The phase difference between the components is related to the azimuth or DOA. Field trials have demonstrated the capability of the DPA in finding incoming DOAs of both genuine and spoofed GNSS signals. With those DOAs, we have developed statistical tests to quantify our confidence of spoof detection [4][11].

We can mathematically derive this relationship by modeling a standard DPA antenna. A DPA can be built from a standard patch with two appropriately placed feeds. This configuration is common in COTS patch antennas and the Stanford DPAs are built using such antennas. The feeds are typically chosen so that they output the response of the antenna to two orthogonal linearly polarized signals. Figure 1 shows such an antenna with a port on the x- and y-axis with an incoming signal with a DOA of θ. The signal measured by each feed, denoted as being X- and Y- axis, denoted as Xin and Yin. These signals, perhaps amplified by a low noise amplifier (LNA), are then sent to a 90-degree hybrid coupler to output RHCP and LHCP signals, denoted as RHCPout and LHCPout. A 90-degree hybrid coupler is a type of directional coupler and has two input and two output ports. It takes the power from each input port and splits it evenly between the output ports along with applying a phase shift. The transform from X- and Y-axis components to RHCP and LHCP is shown in Figure 2. Mathematically this can be modeled by first expressing the inputs and outputs as a vector in Equation 1. Then we derive the transform function of the ideal hybrid coupler in matrix form, given by A as shown in Equation 2.
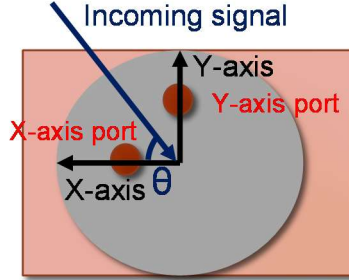

Figure 1. Configuration of DPA antenna and key elements

$$y = [RHCP_{out} \quad Xin \quad Yin \quad LHCP_{out}] \tag{1}$$

$$A = \frac{1}{\sqrt{2}} \begin{bmatrix} 0 & j & 1 & 0 \\ j & 0 & 0 & 1 \\ 1 & 0 & 0 & j \\ 0 & 1 & j & 0 \end{bmatrix} \tag{2}$$

Now suppose we have a linearly polarized signals with amplitude C coming in at an incoming angle θ as seen in Figure 1. Mathematically, we express it in Equation (3). Processing it through the hybrid coupler results in resolving the LHCP and RHCP components as seen in Equation (4). The resulting LHCP and RHCP signals, shown in Equation (5), do not have the same phase. In fact, the phase relationship between these two components is related the azimuthal direction of arrival or DOA

(θ). The phase difference between the two signals, ψ, is related to θ by Equation (6). Since we can track the RHCP and LHCP components, we can measure its phase difference, ψ, and then calculate θ. Equation (6) shows that the calculated θ has a 180-degree ambiguity. That is, an incoming angle θ and θ plus 180 degree (or π radians) will yield the same phase difference measurement ψ. While this is unfortunate, it does not prevent us from using the measurement for spoof detection, heading determination or even interference nulling. It does affect spoof detection performance, as shown later, as it creates greater uncertainty.

$$y = [RHCP_{out} \quad Xin \quad Yin \quad LHCP_{out}] = C * [0 \quad \cos\theta \quad \sin\theta \quad 0] \qquad (3)$$

$$Ay = \frac{C}{\sqrt{2}}[\sin\theta + j\cos\theta \quad 0 \quad 0 \quad \cos\theta + j\sin\theta] = \frac{C}{\sqrt{2}}[e^{j(\pi/2-\theta)} \quad 0 \quad 0 \quad e^{j(\theta)}] \qquad (4)$$

$$RHCP_{out} = \frac{C}{\sqrt{2}}e^{j(\pi/2-\theta)}; \quad LHCP_{out} = \frac{C}{\sqrt{2}}e^{j(\theta)} \qquad (5)$$

$$\varphi_{RHCP} - \varphi_{LHCP} = \psi = \frac{\pi}{2} - 2\theta; \quad \theta = \frac{\pi}{4} - \frac{\psi}{2} \qquad (6)$$
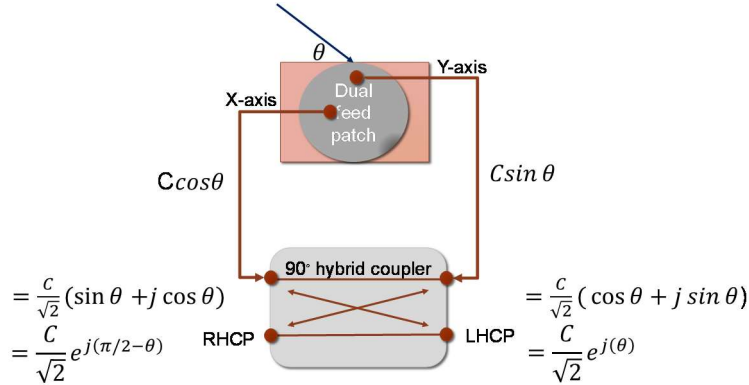

Figure 2. Derivation of RHCP & LHCP signal components using a Patch Antenna and a 90° hybrid coupler

Stanford DPA Platform

Stanford demonstrated many potential benefits of DPA for GNSS by building designs capable of supporting DOA determination [1][2]. This was then used for spoof detection, directional nulling and to provide static heading [3][8][9][10]. Several iterations of the design have been produced. Figure 3 shows a prototype that was flight tested in late 2019 and is the platform that provides the data used in this paper [3]. This iteration was designed to determine DOA faster than prior versions. It also served as a flexible platform to test several GNSS DPA design concepts. As with our other designs, it uses a hybrid coupler to generate RHCP and LHCP. However, this design counter-intuitively has its first stage LNA after the hybrid coupler instead of immediately after the antenna input. Our prior designs had the LNA immediately following the antenna output. The change was made to reduce the effect of phase variations induced by the LNA. Each LNA may induce a small phase error or shift on its input signal. This may be a problem for our design if the phase shift is different between the X-axis and Y-axis signals. If an LNA is used immediately after each antenna input (X- and Y-axis), the differential phase distortion induced on the X-axis and Y-axis signals could result in, for example, RHCP output from the hybrid coupler being a mixed composite of LHCP and RHCP. Equation (7) shows this result with the LNA phase error on the X-, Y-axis output by φ_x and φ_y, respectively. This would then affect the calculated phase difference ψ in a complex way. Having the first stage LNA after the RHCP and LHCP output also affects the phase difference ψ by adding the differential LNA phase errors. This is more straight-forward as it is just an offset to RHCP and LHCP. This makes it easier to average out.

$$Ay = \frac{C}{\sqrt{2}}\left[\left[\frac{1}{2j}(e^{j(\theta)} - e^{-j(\theta)})e^{j(\varphi_y)} + \frac{j}{2}(e^{j(\theta)} + e^{-j(\theta)})e^{j(\varphi_x)}\right] \quad 0 \quad 0 \quad \left[\frac{1}{2}(e^{j(\theta)} + e^{-j(\theta)})e^{j(\varphi_x)} + \frac{1}{2}(e^{j(\theta)} - e^{-j(\theta)})e^{j(\varphi_y)}\right]\right] \qquad (7)$$
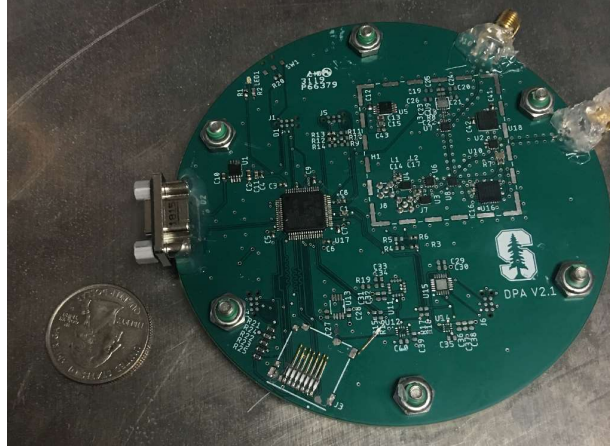
Figure 3. 2019 Stanford Printed Circuit Board DPA with COTS GNSS and RHCP/LHCP out mounted on aluminum ground plate for flight test
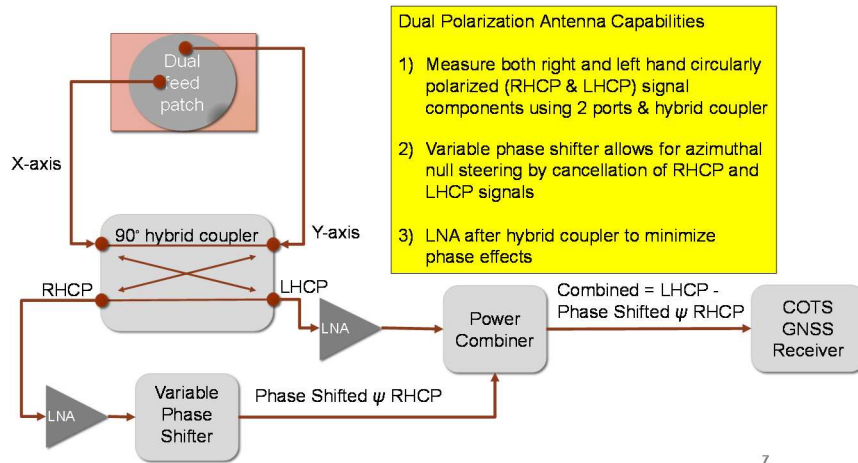

Figure 4. Revised dual polarization antenna for flight tests. Along with other changes, LNA placed after hybrid coupler

After generating the RHCP and LHCP components, the platform supports two possible processing options. One path is to process using a COTS GNSS chipset (i.e. ublox). This tests a low cost, COTS implementation and is shown in Figure 4. The platform can also output the RHCP and LHCP signals for processing by a SDR. This provides the opportunity for more sophisticated processing of the signal. However, these two processing modes are not supported simultaneously. This paper will focus on results from the COTS implementation. In both cases, the key is to be able to determine the phase difference between the LHCP and RHCP signals.

For the COTS based processing, the phase difference between the two circularly polarized components is determined serially. First, a composite signal composed of the LHCP signal and a phase shifted RHCP signal for the COTS receiver to track and calculate the carrier to noise ratio (C/No). At each processing interval, a new composite signal is generated with the RHCP signal systematically shifted in phase using a variable phase shifter controlled by a microcontroller. The variable phase shifter can electronically shift an input signal by a selected amount. The phase shifter is commanded by a microcontroller to apply phase shifts that span the entirety 360 degrees after a fixed number of steps (33 in the discussed implementation for flight tests). After spanning the whole interval, we can estimate the phase shift that results in alignment (maximum C/No) or anti-alignment (null or minimum C/No). The phase shift for the former provides - ψ while the latter provides (-ψ + 180 degrees). While both measurements can yield DOA, we use the nulls as they are sharper. Hence, a good null search algorithm is an integral part of this DOA determination process. The architecture allows for the use of common surface mount components and a small, low power COTS GNSS integrated circuit (IC). It simplifies signal processing into analog combinations for traditional processing by a COTS receiver. The main drawback of this approach is that the serial processing of the two signal components limits the speed with which we can determine DOA. The SDR implementation addresses this issue.

4

EXPERIMENTAL SET UP & FIELD EXPERIMENTS

The 2019 DPA design was tested in the field and flight trials to validate its operational capabilities and to demonstrate the utility of DPA for aviation. We were fortunate enough to partner with the US Air Force (USAF) Test Pilot School (TPS) at Edwards Air Force Base (AFB) to fly our prototypes of the DPA. We built antenna installations (e.g. radome, mounting plate) suitable for use on a USAF test aircraft. A custom 1-foot diameter aluminum plate was built to mount the antenna onto the top of the aircraft. A similar plate, but in a square form, was used to mount another identical DPA antenna on the bottom of the aircraft. For each DPA, the plate acts both as a mounting mechanism and a ground plane. A 3-D printed radome (housing) also built to enclose the antenna in a streamlined shape suitable for flight. The antenna, custom plate and radome for the top mounted antenna are shown in Figure 5. This unit was completely sealed using adhesives to survive the airborne environment. Air worthiness tests were performed on complete antenna units at Holloman AFB, New Mexico (NM) to ensure that it could safely operate within the C-12J flight envelope. For the flight tests, the PCB DPA was mounted on the top and bottom of a C-12J (Military version of the Beechcraft 1900) operated by the 586[th] Flight Test Squadron. This aircraft is shown in Figure 6. Both the mounted top and bottom antenna are shown in Figure 7.
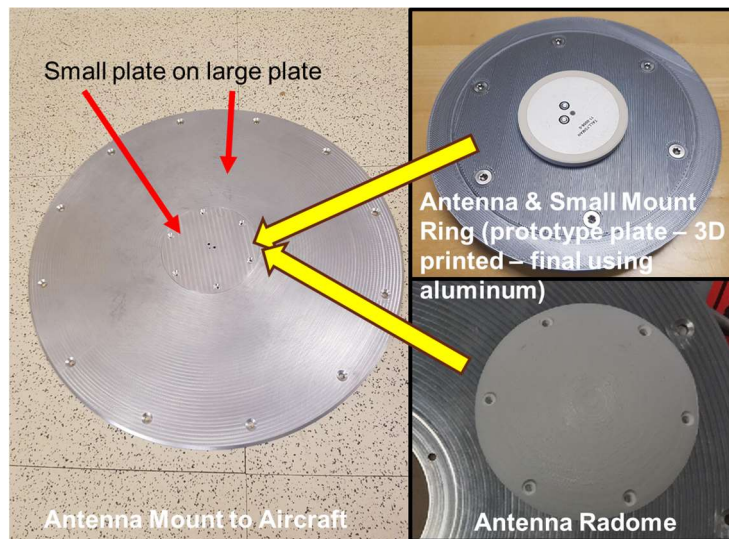

Figure 5. Components of dual polarization antenna for flight test of USAF C-12J with ground plane and antenna radome

Figure 6. USAF 586[th] Test Squadron C-12J used for flight test

Figure 7. Version 2 Stanford DPA mounted on the top (left) and bottom (right) of C-12J.

Ground and flight tests were conducted between 2019 September 6 to 19 at Edwards AFB. For the flight test, we operated the DPAs and a survey grade receiver connected to a separate antenna. Additionally, the USAF operated their ultra-high accuracy reference system (UHARS). The UHARS uses an enhanced embedded GPS with high grade inertial sensors (EGI) to provide position even under GNSS interference. While not used in the Edwards testing, UHARS can also take in non-GPS radio navigation sources such as Locata. For the analysis in this paper, we examined the processed data from the COTS GNSS. The top DPA operated in this mode about half the time whereas the bottom DPA was always operated in this mode. In this mode, the DPA was configured for the flight test to complete a full 360 degree scan every 1.9636 seconds. This was done using phase offset steps of 11 degrees which was adequate to accurately find the phase difference of the null. Each phase offset was used for a 60 millisecond (ms) interval to give the COTS GNSS receiver time to measure its effects.

In the flight trials, we have several possible sources for truth. We could use: 1) COTS GNSS (ublox) position velocity time (PVT) output, 2) the survey receiver and 3) UHARS. Each of these can be used to derive a reference for position and heading. We do not have heading with GNSS measurements from COTS or survey receiver. However, we can estimate true heading by differencing GNSS position estimates which yields the direction of travel, commonly known as course over ground (COG). While COG does differ from true heading due to cross wind effects, this difference is not significant in our evaluation given the magnitude of our other errors. In our prior paper [3], we used GNSS COG as our heading reference. Since then, we received UHARS data which provides both truth position and attitude. Figure 8 shows the comparison of UHARS yaw (blue circles) with the ublox PVT derived COG (red dots) for the top antenna on September 19, 2019. UHARS is used for truth in this paper, particularly for flights operated in the presence of interference. UHARS was designed to operate in GNSS degraded environments and hence its solution was more readily available under interference conditions. Figure 9 shows this clearly where there is a long period of zero COG (no COG output). The plot is for the bottom antenna on September 19, 2019 – a night with interference testing. The bottom antenna experiences more interference as it does not have the benefit of the aircraft body to block these ground transmissions. One drawback with UHARS measurements were provided at a lower rate than GNSS measurements from the ublox which is also seen in the plots.
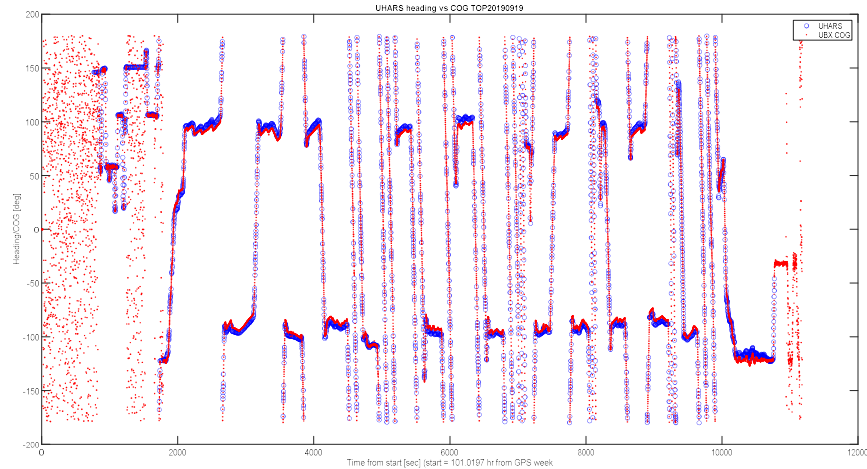
Figure 8. Comparison of course over ground calculated derived from GNSS velocity from the top antenna (red dots) to UHARS yaw truth (blue circles) on 2019-09-19. This flight occurred during GNSS interference testing

The testing was conducted in non-spoofing and predominately nominal conditions. The flight test consisted of 8 flights with two experiencing planned GNSS interference (September 17 and 19). These two flights were conducted as part of the overall Developmental Test (DT) NAVFEST GPS jamming exercise hosted by Edwards AFB testing the effects of denial of GPS L1 and L2 test on equipment and operations. Some of flight scenarios contains many different dynamic maneuverings including straight operations (level flights, climbs, descents) and spiral or turning operations. Figure 10 shows the flight paths of the tests on these various dates based on the UHARS data. These different flights allow us to examine performance under different flight modes and scenarios such as turning versus level flight or interference versus nominal. September 19 represents a flight with lots of turns and the bottom antenna experienced significant interference. September 18 represents a flight that is mostly level and straight. It is under nominal conditions.
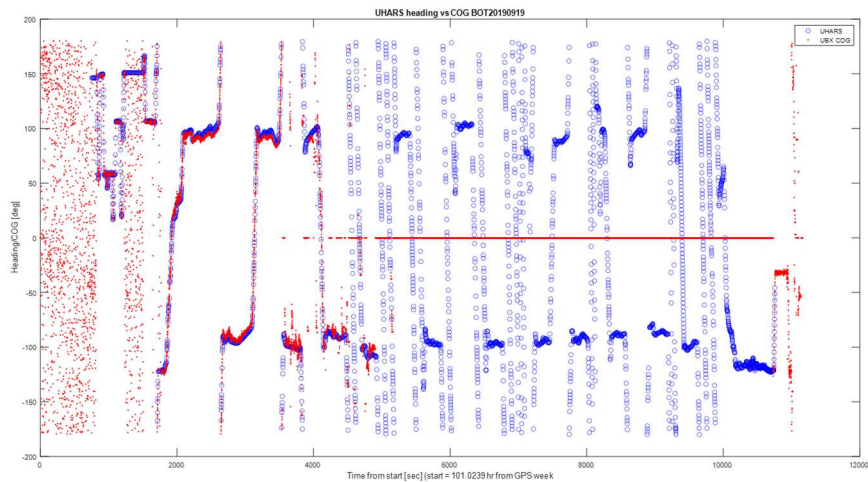

Figure 9. Comparison of course over ground calculated derived from GNSS velocity from the bottom antenna (red dots) to UHARS yaw truth (blue circles) on 2019-09-19. This flight occurred during GNSS interference testing
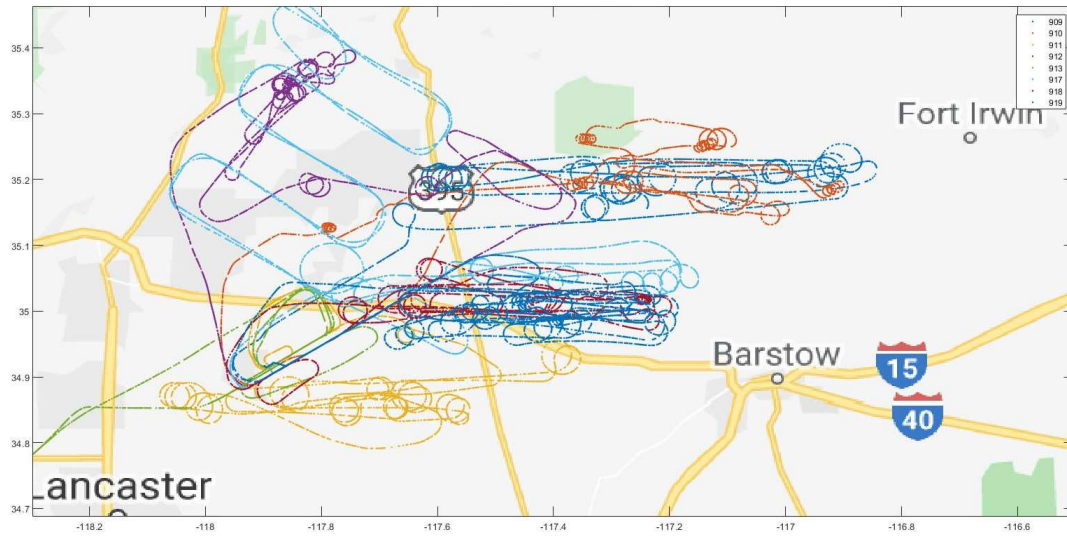
Figure 10. Flight Paths for tests from September 9 – 19, 2019 (dates are in the legend)

FLIGHT TEST RESULTS

Several analyses were conducted with the flight test data. We examined DOA error statistics and the correlation of DOA errors with factors such as elevation and azimuth angle. We examine the calculation of heading with DPA DOA under nominal and interference condition. We examine the spoof detection performance under nominal conditions with this data.

Direction of Arrival Performance

The developed DPA can determine DOA from any tracked satellites such as GPS, Galileo, GLONASS, Quazi-Zenith Satellite System (QZSS) and the Wide Area Augmentation System (WAAS). From the calculated DOA, we can derive error statistics provided we know the true DOA. We can calculate the true DOA since we have heading (or COG) as well as satellite azimuth from true location and satellite ephemeris. This allows us to generate DOA error histograms and error statistics. Figure 11, Figure 12 and Figure 13 shows the error statistics on DPA estimated DOA for various GPS, Galileo and WAAS satellites from the flight of September 18, 2019.

Despite less dwell time and flight dynamics, the revised DPA performed similar to past versions (in static and ground tests) with standard deviation of errors between roughly 20-30 degrees. The dwell times used to measure each phase shift about 10 times shorter than in prior tests [2][4]. The DOAs from the WAAS satellites seem to more accurate than those from GNSS perhaps due to these satellites being stationary or at a good elevation (middle elevation). We next used the flight data to examine DOA performance as a function of satellite, azimuth and elevation as well as location of antenna top (TOP) or bottom (BOT) of aircraft. Figure 14 and Figure 15 show the dependence of the DOA errors with elevation showing the DOA error statistics as a function of elevation angle as well as the percentage of data points at each elevation angle for the top and bottom antenna, respectively. The accuracy of the calculated DOA is expected to be worse if signal has less interaction with the ground plane prior to entering the antenna. This typically occurs with a high elevation satellites (> 60 degrees). Results from the top antenna shown in Figure 14 supports this theory. What is somewhat surprising is that this is also seen for September 19 results for the bottom (BOT) antenna, shown in Figure 15. This is unexpected as all signals entering the bottom antenna is presumed to have impacted the aircraft and ground plane prior to be received by the antenna. However, given that this was during a GNSS interference test and there are fewer data points for the bottom antenna in this case, it may not be conclusive. Results from September 18, 2019 does not show such a dependency and supports the hypothesis. Though not shown in the paper, we also examined the correlation of DOA errors with azimuth and satellite. There should be no correlation to azimuth and minimal correlation to satellite which is really a correlation to signal tracked, signal power and satellite location. The data generally supports this.
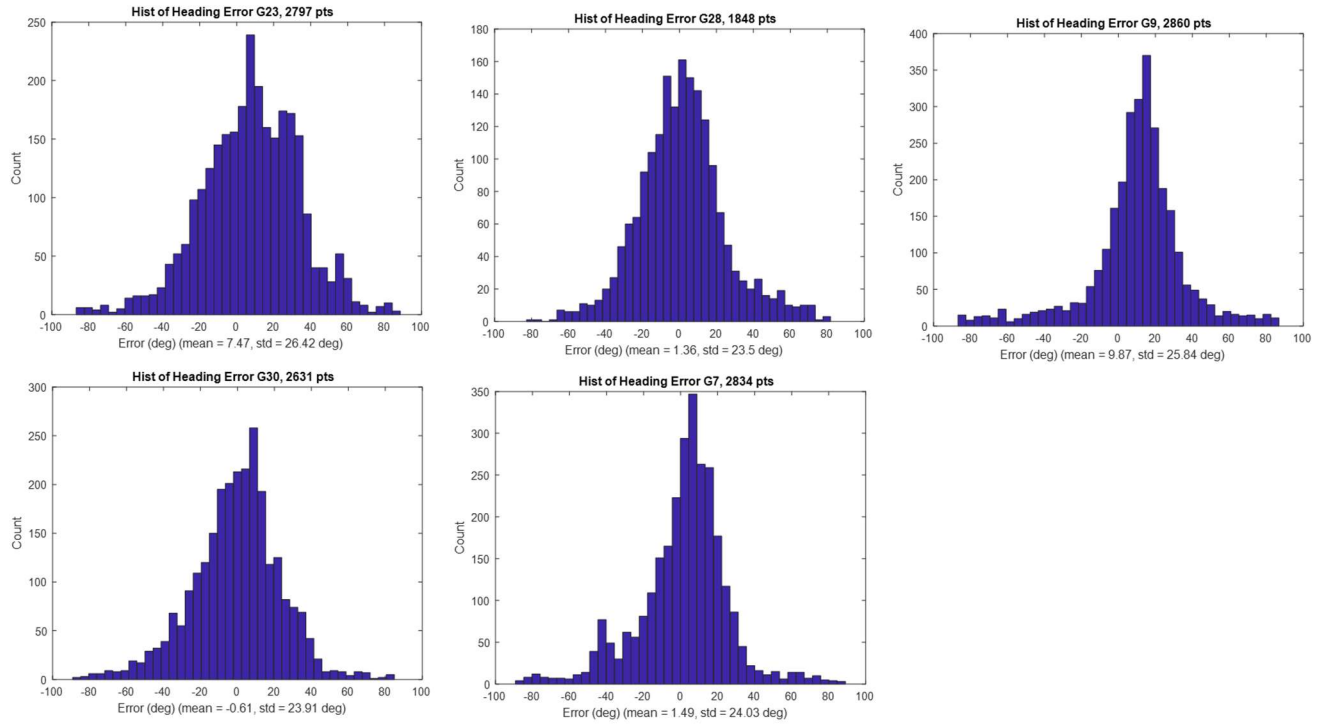
8

Figure 11. Direction of Arrival Error Statistics for Various L1 Signals from GPS Satellites ~ 2 hour flight (2019-09-18)
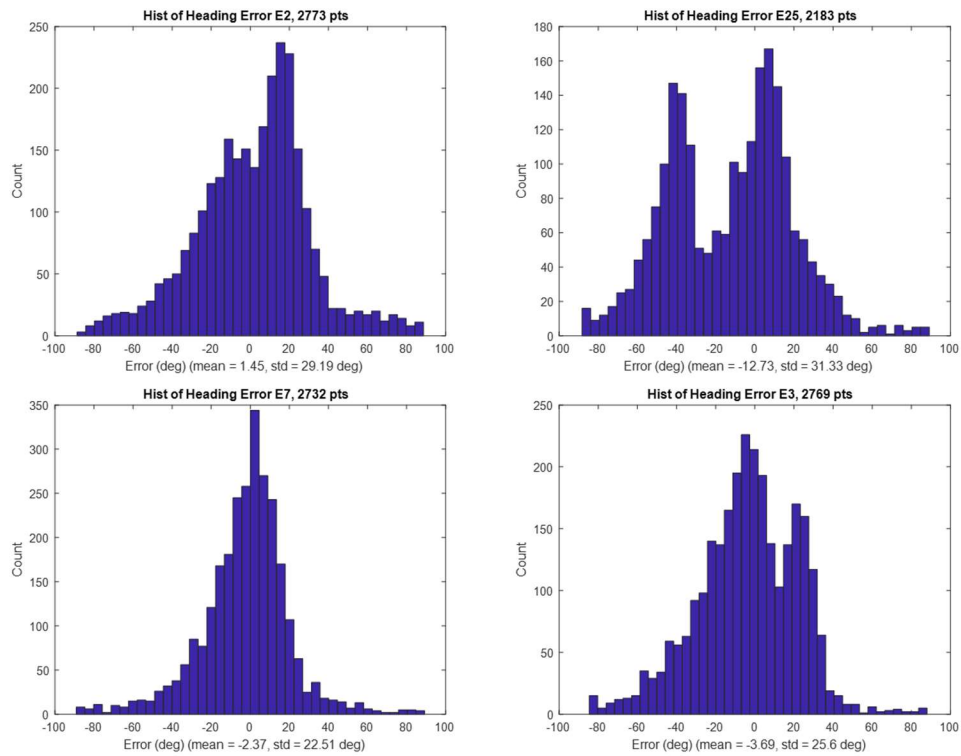

Figure 12. Direction of Arrival Error Statistics for E1 Signals from Various Galileo Satellite ~ 2 hour flight (2019-09-18)
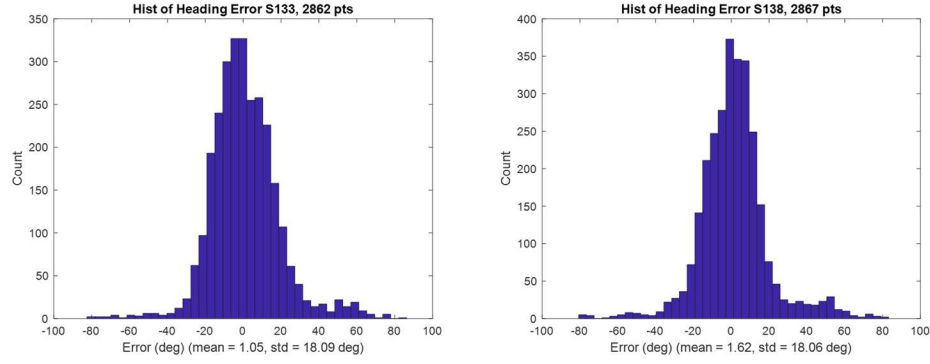
Figure 13. Direction of Arrival Error Statistics for L1 Signals from Various Wide Area Augmentation System (WAAS) Satellites ~ 2 hour flight (2019-09-18)
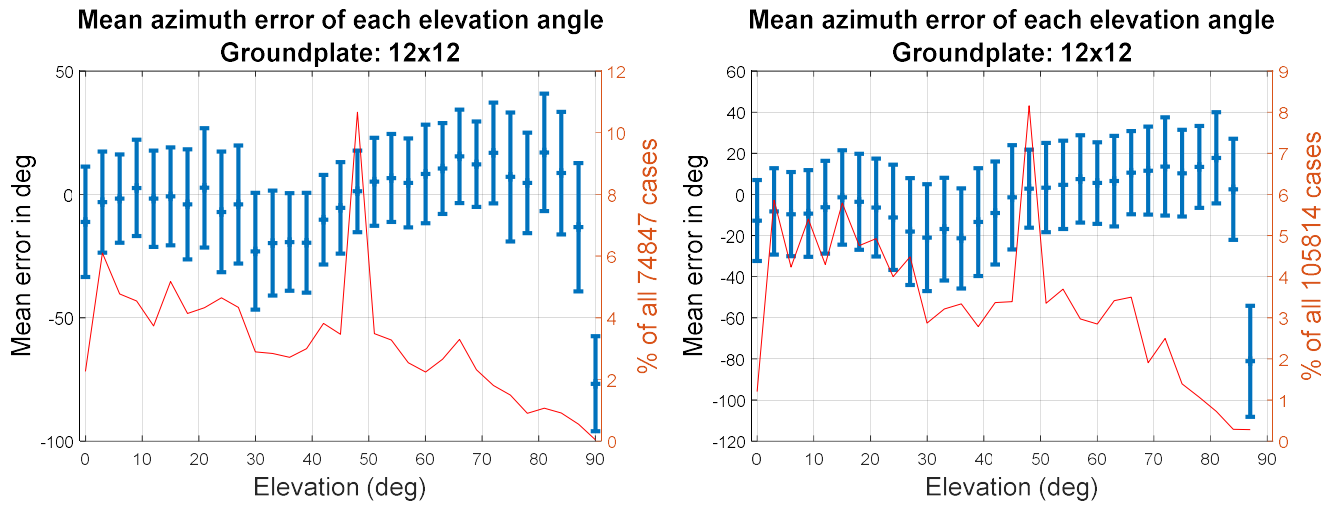

Figure 14. DOA (Azimuth) Error as a function of Elevation Angle for TOP antenna 09-18 (left), 09-19 (right)
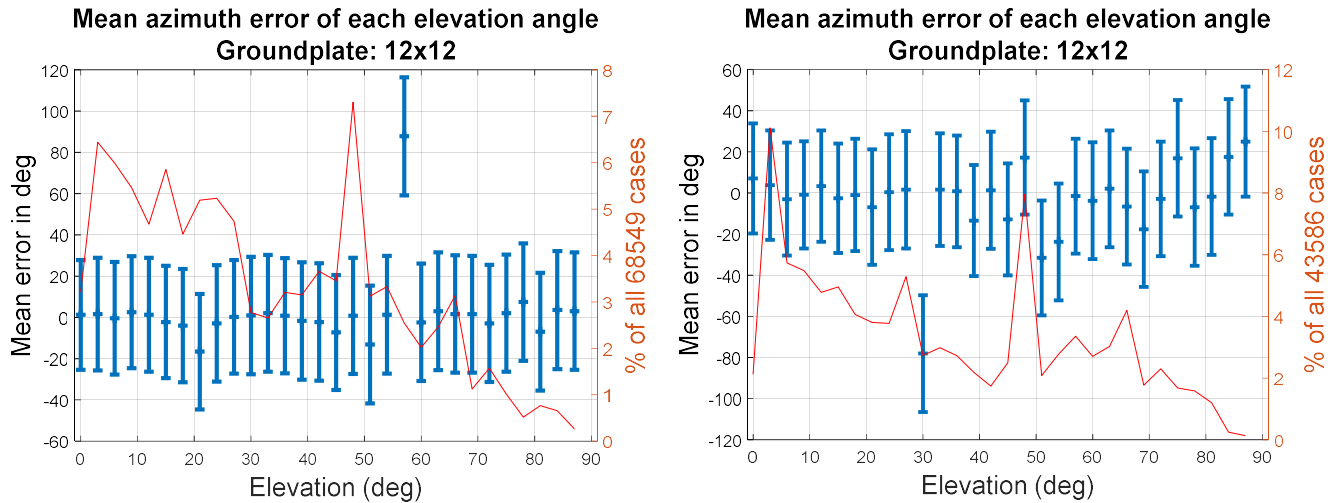

Figure 15. DOA (Azimuth) Error as a function of Elevation Angle for BOTTOM antenna 09-18 (left), 09-19 (right)

These statistics and their characterization as a function of satellite or elevation angle is used in our DOA detection algorithm. The simple way used in our prior implementations is to use one bounding variance level for the algorithm. However, with

10

these results, we can have different bounding variance for different types of satellites (GPS, SBAS, etc.) that are clearly seen in these results. Another is to also tune the spoof detection algorithm by elevation angle and other factors. Using these more precise characterizations will have make our DOA based spoof detection more sensitive.

DPA Heading Estimates

DOA measurements can be used to calculate heading under dynamic and static conditions. The DPA brings the ability to calculate heading in two situations where it is not possible to get heading using GNSS velocity measurements: 1) static situation and 2) less than 4 visible satellites. In prior papers [3][10], we discussed how to use the satellite DOA to estimate heading. This is possible with only 1 satellite but having more satellite improves the estimate. To get the best results, we use a Kalman filter as well as past information used to aid heading estimate. For the plots shown, we only used knowledge of COG to handle 180-degree ambiguity of the DPA. However, this is not necessary and can be solved with proper filtering. Figure 16 shows the DPA derived heading versus the UHARS yaw for September 18, 2019. These match very well, even in static situations such as the beginning of the data set. Figure 17 and Figure 18 show the heading comparison results for the bottom and top antenna for September 19, 2019 which was during an interference test. While a little worse than the prior night, both antennas still provide some reasonable heading estimate with the top antenna, not surprisingly, performing better. The remarkable result is that bottom antenna still provides heading, albeit noisy. This is because for most of the flight, the bottom antenna receiver, as seen in Figure 9, does not have enough satellites to provide position or velocity. Figure 19 shows the number of satellites used by that antenna and it is generally well below four for the middle flight segment. The results shows how well we can use DPA-derived DOA for heading, even under degraded GNSS conditions, as well as the attenuation benefit of the aircraft body.
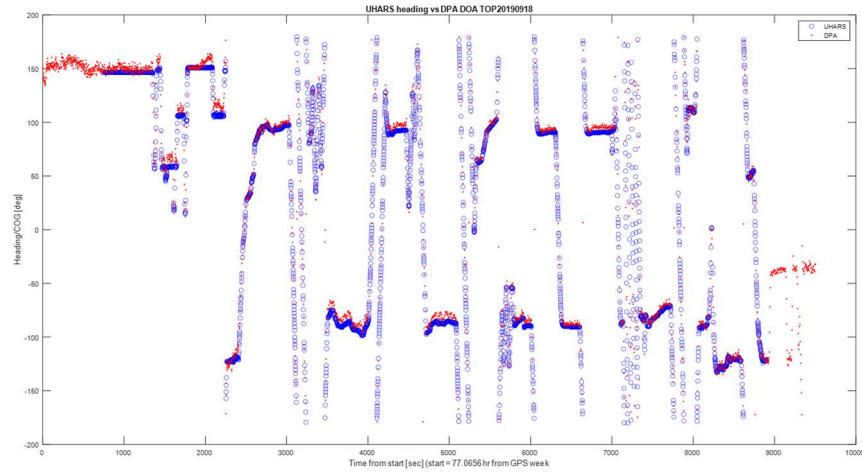

Figure 16. Comparison of heading derived using the direction of arrival from the DPA mounted on top of the aircraft (red dots) to UHARS yaw truth (blue circles) on 2019-09-18.
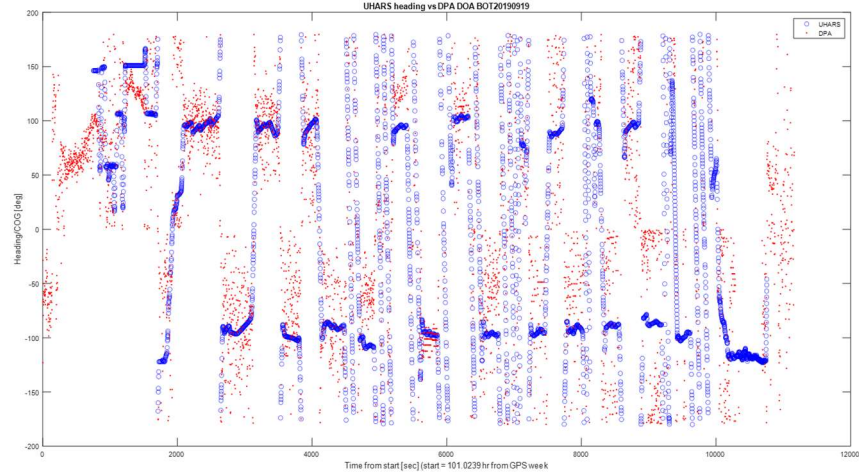
11

Figure 17. Comparison of heading derived using the direction of arrival from the DPA mounted on bottom of the aircraft (red dots) to UHARS yaw truth (blue circles) on 2019-09-19. This flight occurred during GNSS interference testing
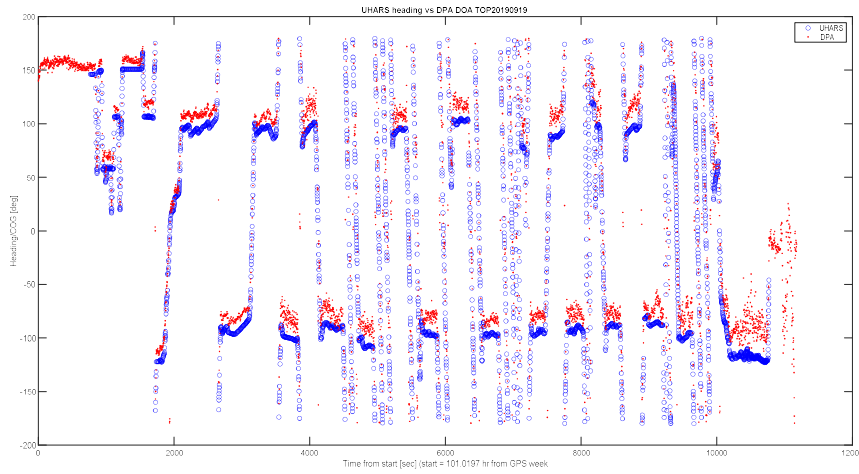

Figure 18. Comparison of heading derived using the direction of arrival from the DPA mounted on top of the aircraft (red dots) to UHARS yaw truth (blue circles) on 2019-09-19. This flight occurred during GNSS interference testing
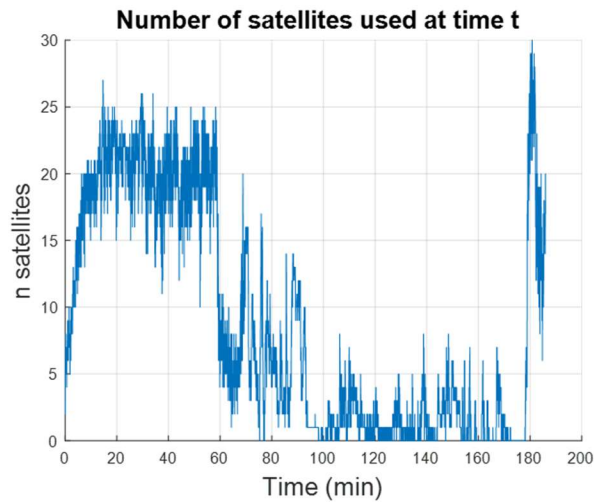

Figure 19. Number of Satellites Used (Available) for BOTTOM antenna on September 19, 2019 flight (GNSS interference testing)

SPOOF DETECTION METRIC UNDER NOMINAL CONDITIONS

The flight data is used to examine the performance of DOA based spoof detection using DPA measurements along with other measurement metrics. In this example, we use the joint Generalized Likelihood Ratio Test (GLRT) framework and process the combination of pseudorange residuals (prr) with three sequential DOA measurements from a DPA. We follow the framework introduced in [16] for a test deciding between a nominal (H0) and spoofed (H1) hypothesis. DOAs and pseudorange residuals are an attractive combination [17]. In the single antenna spoofer case, DOA based detection is most effective when many signals are spoofed whereas prr works best when only a few signals are spoofed. Hence these two detection methods complement each other.

Sequential DOA measurements are several snapshot DOA measurements from consecutive epochs, combined through a histogram filter [13]. By assuming measurement epochs to be independent, this allows for a better decision about the presence of a spoofing attack by considering more evidence at once than the snapshot approach. As detailed in [16], following the standard nomenclature used for example in [18], different detection metrics result in different types of hypothesis tests. These derivations are expanded from the original discussions in [19].

$$\min_{\gamma} P(log\Lambda(y) \geq \gamma|H_1) \; s.t. \, P(log\Lambda(y) < \gamma|H_0) \leq P_{FA_{max}} \quad (7)$$

It helps to have a brief overview of the methodology from [16]. The basic statistical detection test framework for all test – minimize the probability our decision variable $\log \Lambda\,(y)$, under spoofing conditions (H1), is above the threshold γ such that the probability of the decision variable being below the threshold in nominal condition (H0), is below the maximum probability of false alert (P_{FAmax}). This is provided in Eq. (7). The resulting choice of decision variable depends on the form of H0 and H1. The hypothesis test for DOA measurements is a simple H0 vs. simple H1 scenario (SS), meaning that a model of the expected measurements exists for both hypotheses. Pseudorange residuals represent a simple H0 vs. composite H1 scenario (SC), meaning that while their distribution under nominal conditions (H0) is known, the spoofed distribution (H1) is not. For example, a composite scenario may have uncertainty about the parameters of its distribution (e.g. mean is in a range). Each hypothesis test compares a decision variable, $\log\Lambda\,(y)$ - a function of the measurements y, against a detection threshold γ. We follow the Neyman-Pearson paradigm of setting γ as a function of the maximum false alert probability P_{FAmax} and the modeled distribution of $\log\Lambda\,(y)$ under nominal conditions. In the simple vs. simple case, $\log\Lambda\,(y)$ follows a Normal distribution. In the simple vs. composite case, $\log\Lambda\,(y)$ follows a χ^2 distribution. For this test, [16] derives Eq. (8) and (9) as the decision variable for the simple vs composite (SC) and simple vs simple (SS) scenarios, respectively. In these equations, μ_i is the mean of y under hypothesis Hi (i = 0 or 1) and Σ is the standard deviation of y which is assumed to be the same under H0 and H1. The distribution of the combination of the two under nominal conditions (more generically a combination of SC and SS metrics) is derived in [16]. The resulting probability density function (pdf) and the cumulative density function (cdf) is repeated here in Eq. (10) and Eq. (11), respectively, where Γ is the gamma function. For the equations, M is the number of metrics being combined. k, μ and σ are given by Eq. (12) – (14) with k being the degrees of freedom of the χ^2 distribution for each metric m from the set of simple vs composite metrics (M_{SC}) as the decision statistic for SC combinations is χ^2. μ and σ are mean and standard deviation from each metric m from the set of simple vs simple metrics (M_{SS}) as the decision variable in this case is normally distributed.

$$log\Lambda_{SC} = -\frac{1}{2}\|y - \mu_0\|_{\Sigma^{-1}}^2 = -\frac{1}{2}(y - \mu_0)^T\Sigma^{-1}(y - \mu_0) \; (8)$$

$$log\Lambda_{SS} = -\frac{1}{2}(\mu_0 - \mu_1)^T\Sigma^{-1}\left(y - \frac{1}{2}(\mu_0 + \mu_1)\right) (9)$$

$$f_{1:M}(z|H_0) = 2^{\frac{k}{4}-1}\sigma^{\frac{k}{2}-1}\frac{1}{\sqrt{2\pi}}\frac{1}{\Gamma(k/2)}exp\frac{-(z-\mu)^2}{2\sigma^2}\sum_{l=0}^{\infty}\frac{1}{l!}\left(-\sqrt{2}\right)^l\left(\frac{z-\mu}{\sigma} + \sigma\right)^l\Gamma(k/4 + l/2) \; (10)$$

$$F_{1:M}(x|H_0) = \int_{-\infty}^{x} f_{1:M}(z|H_0)dz \; (10)$$

$$\mu = \mu_{0,SS} = \sum_{m\in M_{SS}}\mu_m|H_0 \; (12)$$

$$\sigma^2 = \sigma_{SS}^2 = \sum_{m\in M_{SS}}\sigma_m^2 \; (13)$$

$$k = k_{SC} = \sum_{m\in M_{SC}}k_m \; (14)$$

From our flight, we obtained pseudorange residual measurements at 10 Hz, which is at a significantly higher frequency than what we can currently obtain with our ublox DOA measurements. The prr measurement model is characterized in [14] for

this antenna/aircraft configuration, and the DOA measurement model in [13]. The joint GLRT is computed at every DOA measurement epoch, using the sequential DOA measurements as well as the latest pseudorange residuals.

Under nominal conditions we analyze this metric combination with respect to the constraint on false alerts. We calculated the γ of the joint test which formally defined in Eq. (15) and obtained by numerically solving Equation (16). We plot the difference between threshold and measurements $-2log\Lambda - (-2\gamma)$ in Figure 20 for GLRTs using only DoA and prr measurements as well as the joint GLRT of DOA and prr measurements. The figure shows results using data from the DPA for the September 18 flight test. The interested reader is referred to [16] for an explanation of the seemingly arbitrary factor of -2 in the difference. Values below 0 indicate no alarm (the metric is below the threshold). $P_{FAmax} = 10^{-3}$ is set unrealistically high to provoke false alarms within the given flight time. We can see that the pseudorange residual test by itself causes two. This is a small number of false alerts among >57,000 measurement epochs with pseudorange residuals, the measurement model appears to be conservative. The 2886 epochs of sequential DOA and joint tests each result in no false alarm, over-satisfying the false alert probability constraint.

$$\gamma = F_{1:M}^{-1}(P_{FAmax}) \quad (15)$$

$$\gamma = \arg min_\gamma \|F_{1:M}(\gamma|H_0) - P_{FAmax}\|_2 \quad (16)$$
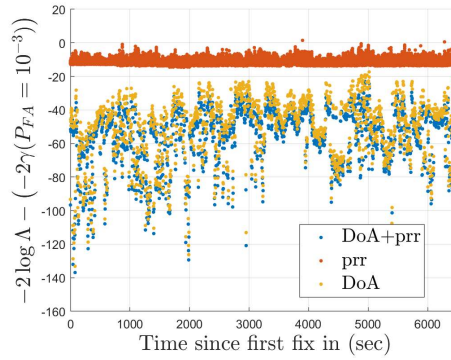

Figure 20. $-2 log\Lambda - (-2\gamma)$ when using both DoA and prr measurements, and when using either measurement alone. The threshold γ was computed with $P_{Famax} = 10^{-3}$. Negative values indicate no alarm (using data from September 18, 2019 flight). Results are updated from [19]

These real-world test results help validate the assumptions inherent to the joint GLRT framework in practice. An extensive data collection campaign would be necessary to cast tighter measurement models and detection threshold.

The results shown further only show that we have "done no harm" with respect to our constraint on false alerts by combining the two metrics. While the benefits of the combinations under spoofed conditions has been shown in simulation [16][17], this has yet to proven in flight.

BENEFITS OF RESOLVING 180 DEGREE AMBIGUITY FOR SPOOF DETECTION

One limitation with using a DPA is that there is a 180-degree ambiguity in the determined DOA. This can be seen in the derived equation for the DPA derived earlier. This is not a unique problem as other DOA determination based on two elements (i.e. two separated antennas or interferometry) also have ambiguities. This ambiguity affects spoof detection performance by adding uncertainty and practical resolution of the ambiguity can improve detection. In this section, we examine the benefits of resolving the ambiguity and of reducing the DOA measurement error.

To examine the benefits of resolving the 180-degree ambiguity, we examine our spoof detection test using DOA error statistics measured from the flight test which has a nominal standard deviation of 17-29 degrees. Based on this capability and different geometries, we can calculate the relationship between the probability of missed detection (PMD) of spoofing versus the probability of false alerts (PFA). This is the receiver operating characteristic (ROC) curve. The ROC result for our baseline case with 180-degree ambiguity is seen in Figure 21 which shows this curve, for detection with 6 satellites, as a function of the number of epochs over which we are able to perform the sequential spoof detection test. One can see from the ROCs that increased observation epochs results some improvements but for the most part, if we want to have $P_{FA} \sim 10^{-9}$, then we will have very high P_{MD}. To get P_{MD} below 0.1, our P_{FA} will also be very high ~ nearly 1.
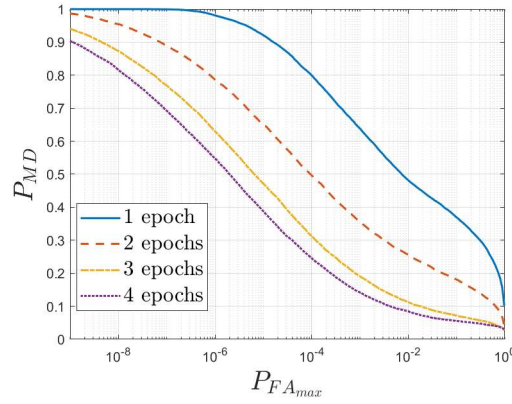
14

Figure 21. Receiver Operating Characteristic (ROC) Curve for DPA Spoof Detection based on DOA with 180-degree ambiguity

However, if we are able to resolve the 180-degree ambiguity, leaving everything else the same, the ROCs improve greatly. This is shown in Figure 22 which has the ROCs given knowledge or resolution of the 180-degree ambiguity. As can be seen, P_{MD} is now better than 0.2 for all cases at a P_{FA} of 10^{-9}. P_{MD} below 0.02 is achievable at P_{FA} of 10^{-4} for the 1 epoch case. It turns out this performance is similar to the performance achieved if we could not resolve 180-degree ambiguity but could reduce the standard deviation of DOA error by half. This is shown in Figure 23. So for our spoof detection algorithm, eliminating the ambiguity is effectively like reducing the root sum squared DOA measurement errors by half. Figure 23 also highlights the importance of having accurate characterization of the DOA error. Having better and more precise characterization of DOA errors improves spoof detection performance.
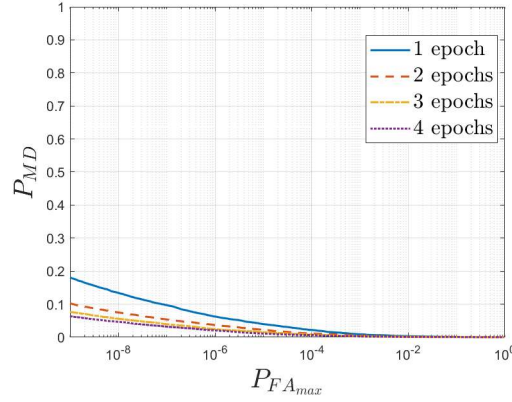

Figure 22. Receiver Operating Characteristic (ROC) Curve for DPA Spoof Detection based on DOA without 180-degree ambiguity
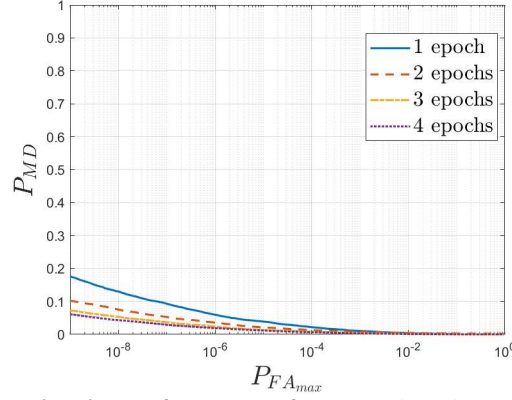
15

Figure 23. Receiver Operating Characteristic (ROC) Curve for DPA Spoof Detection based on DOA with 180-degree ambiguity but half of prior standard deviation of DOA error

SUMMARY

The dual polarization antenna offers a simple but powerful tool for mitigating GNSS radio frequency interference - both spoofing and jamming. This paper shows the practical implementation of DPA technology using COTS parts that was tested flight and how this and other developments can help DPA derived DOA based spoof detection. The DPA DOA performance is reasonable and similar to prior ground-based tests. With this data, we also better characterized the DOA error statistics as a function of different factors such as satellite. It shows limited dependency to elevation but no dependency on azimuth. Additionally, the performance of DPA DOA derived heading corresponds well to truth and can be derived with very few satellites in view. Results from flight tests show that our DPA can calculate reasonable DOAs and that we can use those DOAs to conduct spoof detection. It shows spoof detection performance under nominal conditions with the combination of DOAs and pseudo range residuals from the flight results. The results highlight how these two techniques can complement each other and also reduce probability of false alerts and missed detections. Finally, this paper examines another means of improving DPA derived DOA for spoof detection by examining the benefits of resolving 180-degree ambiguity inherent in the DPA DOA calculations.

ACKNOWLEDGMENTS

REFERENCES

[1] Yu Hsuan Chen, Fabian Rothmaier, Dennis Akos, Sherman Lo, Per Enge, "Towards a Practical Single Element Null Steering Antenna," Proceedings of the Institute of Navigation International Technical Meeting, Monterrey, CA, January 2017

[2] Yu Hsuan Chen, Fabian Rothmaier, Dennis Akos, Sherman Lo, Per Enge, "Demonstrating Single Element Null Steering Antenna Direction Finding for Interference Detection," Proceedings of the Institute of Navigation International Technical Meeting, Reston, VA, January 2018

[3] Sherman Lo, Yu Hsuan Chen, Godwin Zhang, Fabian Rothmaier, Chiawei Lee, "Developing a Dual Polarization Antenna (DPA) for High Dynamic Applications," Proceedings of the Institute of Navigation International Technical Meeting, San Diego, CA, January 2020

16

[4] Fabian Rothmaier, Yu-Hsuan Chen, Sherman Lo, "Improvements to Steady State Spoof Detection with Experimental Validation using a Dual Polarization Antenna," Proceedings of ION GNSS+ 2019, Miami, FL September 2019

[5] Matthew Trinkle, W-C Cheuk, "Null-steering GPS dual-polarised antenna arrays" Presented at SatNav 2003 The 6th International Symposium on Satellite Navigation Technology Including Mobile Positioning & Location Services, Melbourne, Australia 22-25 July 2003

[6] M. Rosen, M. Braasch, "Low-Cost GPS Interference Mitigation Using Single Aperture Cancellation Techniques," Proceedings of the Institute of Navigation National Technical Meeting, 1998 pp. 47-58.

[7] T. Kraus, F. Ribbehege, B. Eissfeller, "Use of the Signal Polarization for Anti-jamming and Anti-spoofing with a Single Antenna," Proceedings of Institute of Navigation GNSS+, Tampa, FL, September 2014, pp. 3495-3501.

[8] Emily McMilin, David De Lorenzo, "Field Test Validation of Single-Element Antenna with Anti-Jam and Spoof Detection," Proceedings of Institute of Navigation GNSS+ 2015, Tampa, FL, September 2015

[9] Emily McMilin, "Single Antenna Null Steering for GPS & GNSS Aerial Applications," Ph.D. Dissertation, Stanford University, March 2016

[10] Fabian Rothmaier, Yu Hsuan Chen, Sherman Lo, J. David Powell, "Single GNSS Antenna Heading Estimation," Proceedings of Institute of Navigation GNSS+ 2019, Miami, Florida, September 2019.

[11] Sherman Lo, Hridu Jain, Yu Hsuan Chen, Per Enge, "Robust GNSS Spoof Detection using Direction of Arrival: Methods and Practice," Proceedings of the Institute of Navigation GNSS+ 2018, Miami, FL Sept 2018

[12] Wim De Wilde, Jean-Marie Sleewaegen, Bruno Bougard, Gert Cuypers, Alexander Popugaev, Markus Landmann, Christopher Schirmer, Daniel Egea Roca, José A. López-Salcedo, Gonzalo Seco Granados, "Authentication by Polarization: A Powerful Anti-Spoofing Method," Proceedings of the 31st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2018), Miami, FL, September 2018, pp. 3643-3658. https://doi.org/10.33012/2018.15917

[13] Fabian Rothmaier, "Optimal Sequential Spoof Detection based on Direction of Arrival Measurements," Proceedings of the Institute of Navigation GNSS+ 2020, Virtual, September 2020.

[14] Fabian Rothmaier, Leila Taleghani, Yu-Hsuan Chen, Sherman Lo, Eric Phelts, Todd Walter, "GNSS Spoofing Detection through Metric Combinations: Calibration and Application of a General Framework," Proceedings of the Institute of Navigation GNSS+ 2021, St. Louis, MO, September 2021.

[15] Cole Johnson, Chiawei Lee, Marcea Ascencio, "Developmental Test NAVFEST: A Large-Scale, Multi-Aircraft, GPS Jamming Test Event," *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Monterey, CA, April 2018, pp. 858-868.

[16] Fabian Rothmaier, Yu-Hsuan Chen, Sherman Lo, Todd Walter, "A Framework for GNSS Spoofing Detection Through Combinations of Metrics," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 57, no. 6, pp. 3633–3647, Dec. 2021, doi: 10.1109/TAES.2021.3082673.

[17] M. C. Esswein and M. L. Psiaki, "GNSS Anti-Spoofing for a Multi-Element Antenna Array," in Proceedings of the Institute of Navigation GNSS+ 2019, Miami, FL, September 2019, pp. 3197–3214, doi: 10.33012/2019.17062

[18] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley & Sons, Inc., 2001.

[19] Fabian Rothmaier, "Statistical Inference for Safe and Continuous Navigation in the Presence of GNSS Spoofing," Ph.D. Dissertation, Stanford University, December 2021