Accommodating Direction Ambiguities in Direction of Arrival based GNSS Spoof Detection

Hridayangam Jain, Sherman Lo, Yu-Hsuan Chen, Fabian Rothmaier, J. David Powell
Stanford University

BIOGRAPHY (IES)

Hridayangam Jain is an undergraduate majoring in Aeronautics and Astronautics at Stanford University.

Sherman Lo is a senior research engineer at the Stanford GPS Laboratory. 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 automobiles. He also works on spoof and interference mitigation for navigation. He has published over 100 research papers and articles.

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 is a Ph.D. candidate in the Stanford GPS Laboratory in the Department of Aeronautics and Astronautics.

J. David Powell is a professor emeritus of Aeronautics and Astronautics at Stanford University. He received his Ph.D. in Aeronautics and Astronautics from Stanford University in 1970. He has conducted research on the use of GPS-based attitude determination augmented with inertial sensors, applications of the FAA's WAAS for enhanced pilot displays, flight inspection of aircraft landing systems, and the use of WAAS and new displays to enable closer spacing of parallel runways.

ABSTRACT

Using direction of arrival (DOA) for GNSS spoof detection has several desirable properties. First, DOA-based spoof detection makes any spoofing from a single antenna very detectable regardless of how sophisticated its generation. It is difficult for a GNSS spoofer to create different DOAs as it generally requires transmitting from different locations, simultaneously. Thus, it forces a spoofer to utilize a much more complicated transmission system than a single antenna to create signals that can deceive DOA-based spoof detection. Second, it is complementary to and independent of other commonly used GNSS spoof detection methods thus providing additional layer of protection and certitude to detection.

To utilize this method, we need means of getting DOA measurements of GNSS signals, preferably one that is both simple and has small form factor equipment. Controlled reception pattern antenna (CRPA) and dual polarization antenna (DPA) are two means of making such measurements [1][2]. While simple, low-profile methods such as the DPA and a two-element antenna are preferred, these methods result in ambiguity in measured direction of arrivals. DPA measurements have a 180-degree ambiguity while two-element antennas have a symmetric ambiguity in DOA along the axis between the two antennas. The ambiguity can affect detection performance and limit the utility of such a system.

This paper examines the ambiguity issue, focused on the DPA. It examines and develops a processing method to handle the ambiguity. First, we create two separate cases from the ambiguous measurements – a best genuine and a best spoof case. From these cases, we develop tests to examine each case and their likelihood. We use the processing results of both cases to manage the ambiguity. The processing method is tested and demonstrated using simulations and data from on-air tests in both nominal and spoofing conditions.

INTRODUCTION

The marked increased use of GNSS technology in many safety and economically important activities, such as portable devices for location-based applications, means that spoofing and interference attacks – even low power ones – can potentially affect many users and many critical applications. This is made even more likely as both the hardware and software required to simulate GNSS signals is becoming relatively inexpensive and easily accessible. Unfortunately, there does not seem to be a single reliable spoof detection and mitigation method to deter or defend against an intelligent attacker. Very powerful protection can be offered by combining methods and consistency checks to detect if signals are genuine or spoofed.

Direction of arrival (DOA) methods are useful because they are independent from other commonly proposed methods and can therefore provide a complimentary layer of protection [2][3][4][5]. Most documented GNSS spoofing are broadcast from a single transmitter such that the signals arrive from the same direction. Anticipated GNSS spoofing attacks in the near future will likely be from single antenna due to the technical challenges of multi-directional spoofing attacks that require coordinating transmissions from several antennas. On the other hand, genuine GNSS signals arrive from different directions and generally match the expected directions from the ephemeris data. Having the DOA-based spoof detection makes single antenna spoofers much more detectable or forces a spoofer to conduct a much more complicated attack using multiple synchronized locations. An attractive method to measure DOA of GNSS signals is with the Stanford dual polarization antenna (DPA) based on the polarization of the signal [8]. Unfortunately, the measured DOAs from the DPA have a 180-degree ambiguity. This introduces challenges with respect to the identification of spoofed signals.

This paper details a method for managing this 180-degree ambiguity for GNSS spoof detection using software. It focuses on analyzing two cases – best genuine and best spoofed – and considers statistical tests to distinguish spoofing. The performance of the method is then examined using experimental data from a government spoofing test. While additional hardware may be able to resolve the ambiguity, it can increase complexity or the size of the antenna. Hence, this paper examines the possibilities and limitations of managing the ambiguity in software when processing the ambiguous received DOAs.

BACKGROUND

Dual Polarization Antenna

Dual polarization antennas are attractive for GNSS interference mitigation for several reasons. In their earliest incarnations for GNSS protection, DPA were developed for mitigation of jamming [7][8]. More recently, it was found that they can also provide GNSS spoof detection through the Stanford DPA shown in Figure 1 [6] and other similar antennas [11]. While DOA measurements typically require measurements from multiple antennas or elements, the Stanford DPA can determine the relative azimuth direction of GNSS signals with a single antenna. Furthermore, the antenna can be created with commercial off-the-shelf hardware and is capable of receiving both right and left-hand circularly polarized (RHCP and LHCP, respectively) signals, which allows it to calculate azimuth and even rough elevation angles. Last, but certainly not least, this is a single antenna solution using two feeds which makes it complaint with International Traffic in Arms (ITAR) and export arm restrictions (EAR), which limits the use of multi-element GNSS antennas for civil applications, such as spoof mitigation, to three elements [9].

Figure 1. Stanford PCB Dual Polarization Antenna

The Stanford DPA takes advantage of the following insight: no matter how a signal is initially polarized, it becomes linearly polarized when it impacts the ground plane. The signal component traversing the ground plane thus has equal RHCP and LHCP components. This concept is illustrated in Figure 2. This allows for a rough calculation of the elevation angle of GNSS signals, which are transmitted RHCP. Signals arriving from higher elevations will have larger RHCP components and lower elevation signals will have more equal RHCP and LHCP components due to increased impact with the ground plane.

Furthermore, the phase relationship between the LHCP and RHCP signals can be used to back out a relative azimuth angle or DOA. The phase difference that aligns with destructive interference of the two components is related to the DOA by Equation 1 where φ is the true azimuth, φ_0 is the offset azimuth of the antenna relative to true north, ε_φ is the error on the azimuth measurement, and ψ is the applied DPA phase shift. The drawback of using the DPA direction finding is seen in the equation, which is derived from the physics and the effect of the hybrid coupler used to form the LHCP and RHCP signal [6]. A null will occur at the same applied phase shift for a given DOA or that DOA plus 180 degrees. This results in a 180-degree ambiguity in the DOA measurement.

The onboard electronics on the Stanford DPA process the incoming signals through a hybrid coupler to create a RHCP and LHCP signal [1]. The Stanford DPA electronics then applies a discrete sweep through all phase shifts to the RHCP signal and forms a combined signal composed of the sum of the phase-shifted RHCP and LHCP signals. This is used to determine DOA by finding the phase shift that results in the peak or the null in the carrier-to-noise ratio (C/No) of the combined signal. This process is visualized in Figure 3, where the C/No of the combined signal are shown for each of three satellites, each at different elevations. The local minimums on this plot correspond with the phase shift that results in nulls (i.e. destructive interference). Notice that the nulls are much weaker at the higher elevation angles (> 65 degrees) causing the DOA estimates to have larger errors and be less reliable. Hence, in the analysis in this paper, the DOAs from such satellites (elevation angles > 65 degrees) are excluded and not used.

$$\psi = 2\left(\varphi - \varphi_0 + \varepsilon_\varphi\right) + 90° \qquad (1)$$

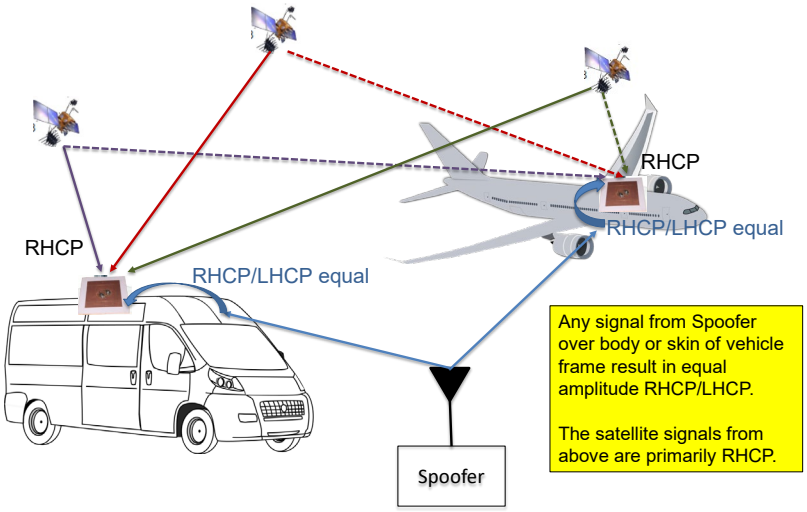

Figure 2. Dual Polarization Antenna Concept for Spoof Determination

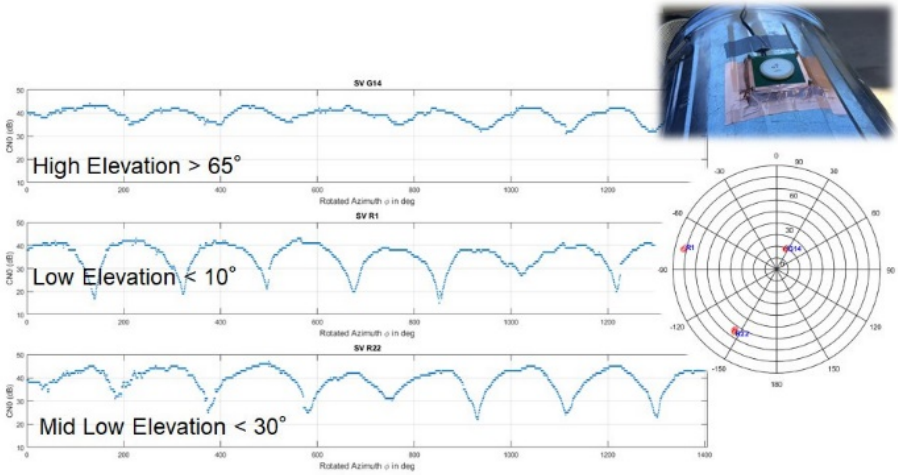

Figure 3. C/No vs. Rotation of LHCP component for different satellites: high (top), low (bottom), very low (middle). At low and very low elevation, rotation of LHCP component to the correct angle can significantly cancel out RHCP, ~ 180 second to complete all rotations.

AMBIGUITY PROCESSING

General Problem Statement

Direction of arrival-based spoof detection works by examining the consistency of the measured direction of arrival with what is expected. The measured DOAs are typically DOAs relative to a common but unknown direction rather than North, which is often not known. Absolute DOA, which is DOA relative to a known direction such as North, requires knowledge of antenna orientation, such as heading, that is derived from a trusted source. Trusted heading can be derived from many sources – even GNSS using previously authenticated measurements updated with inertial sensing. However, there are cases where that is not available, such as on start up or if the system does not have good attitude or direction sensors.

Regardless of whether we have absolute or relative DOA, we can use differences in DOA (ΔDOA) between the measured azimuth angle and the expected direction. Under nominal (i.e. genuine) conditions, if we compare the measured DOA results to the expected directions from ephemeris data, the ΔDOA distribution will be roughly the same for all satellites. In other

words, the difference between measured and true DOA should be a bias offset, roughly equal to our assumed zero azimuth direction and true North, that is largely common for all signals. If instead we assume a single direction as the expected direction, each DOA would have a different bias. Under spoofing conditions, assumed to be from a single source, the opposite would be true. With the ΔDOA results, various hypothesis tests can be used to determine if the measured angles belong to a nominal or spoofed distribution and the relative likelihood of each possibility [10]. For example, we can take as a test statistic the sum of the squared ΔDOAs, normalized by the estimated variance. For example, assuming we have all genuine measured signals, we can make two comparisons and test statistics. If we create the test statistic using the ephemeris for the expected DOAs and account for the bias offset, the statistic should be distributed centralized chi-squared (χ^2).

One method shown in reference [10] examines the DOA measurements in two ways. It finds two non-exclusive sets of DOAs: one that is consistent with coming from direction (spoofing) and one consistent with the broadcast ephemeris (genuine or nominal). Statistical estimates of likelihood are then calculated for these cases, seen in Equations 2-7, and used to evaluate which case (spoofed or genuine) is more likely.

Assuming we have a set *S* consisting of *M* satellite DOA measurements that is consistent with a single direction of arrival (spoofed), we derive Equations 2-4. Using the test statistic of the sum squared ΔDOAs relative to an estimated spoof direction (the mean DOA) normalized by the variance of spoof DOA, $\sigma_{AZ,spoof}$, Equation 2 provides the probability of that test statistic. This is calculated from a probability density function (pdf) of a centralized chi squared distribution with *M* degrees of freedom. This calculation quantifies how well this set of signals match up with the assumption that they come from the same direction. Equation 3 shows a similar probability. It shows the probability of S matching the genuine or nominal satellite DOAs. This is calculated with a similar test statistic except using the true azimuth from the ephemeris and the measurements variance of nominal DOA, σ_{AZ}. As we are testing to see how the signals are consistent with coming from genuine satellites, this pdf is also from a centralized chi-squared distribution. Essentially, this calculation quantifies how well this set of signals match up with measurements from the genuine constellation.

The relative probability of spoofing conditions given S is calculated using Equation 4. The variable *N* is used to relate the prior probabilities of spoof and genuine conditions with $P_{spoof,\ prior} = N{*}P_{norminal,\ prior}$. N=1 is generally used. Equation 4 is an important calculation for discriminating between spoofing and genuine. For example, while a set may be highly likely to come from one direction, it may also be true that those satellites are actually in a similar direction resulting in a high likelihood of matching the genuine DOAs. The result in this example is that Equation 4 would not indicate a strong probability of confidence that this set is from spoofing. Equation 5 shows the calculation for relative probability of having nominal signals.

Note, we derived values for σ_{AZ} and $\sigma_{AZ,spoof}$ from our various field tests. Values of roughly 14 degrees for σ_{AZ} and 5 degrees for $\sigma_{AZ,spoof}$ are used in the analysis [10].

$$P_{spoof,S} = \chi^2_M \text{ pdf}\left(\Sigma_{i,\epsilon S}\left(\frac{AZ_i - AZ_{spoof}}{\sigma_{AZ,spoof}}\right)^2\right) \qquad (2)$$

$$P_{nom,S} = \chi^2_M \text{ pdf}\left(\Sigma_{i,\epsilon S}\left(\frac{AZ_i - AZ_{i,true}}{\sigma_{AZ}}\right)^2\right) \qquad (3)$$

$$P_{spoof\ overall,S} = \frac{P_{spoof,S}}{P_{spoof,S} + N{*}P_{nom,S}} \qquad (4)$$

$$P_{nom\ overall,S} = \frac{P_{nom,S}}{P_{spoof,S} + N{*}P_{nom,S}} \qquad (5)$$

If the DOAs are more or less correct, then the above calculation would be sufficient. Unfortunately, there is a 180-degree ambiguity, and this introduces additional uncertainty that must be accounted for in spoof detection. The ambiguity affects each GNSS signal independently such that there are two possible azimuth angles per satellite – the calculated measured angle and that angle plus 180 degrees. As we have relative rather than absolute DOAs, the ephemeris cannot be used to set the

ambiguity directly. Even with a minimal GNSS constellation of four satellites this ambiguity results in eight possible relative configurations as we can assume a base direction for one satellite and can vary the ambiguity of the other three relative to that one. This would be computationally expensive to compare across all such cases.

Accounting for the Ambiguity in Software

To manage the 180-degree ambiguity issue, we developed a calculation process in order to answer our fundamental question: can my position be trusted? The process is shown in Figure 4. In our baseline process, we start by excluding high elevation angle satellites (> 65 degrees from ephemeris) as the DPA does not estimate these DOAs well. Next, we create two cases: a si called best-nominal case, where all the signals are consistent with the genuine or nominal satellites; and a best-spoofed case, where all the signals are consistent with being from a single direction (hence spoofed signal). These are created using the measured DOAs and choosing the ambiguity that best match the case. Then, we perform calculations to calculate the biases in relative DOA for each case and eliminate outliers. Outlier elimination is critical to the process to eliminate poor measurements and also perhaps account for the case where there is a mixture of genuine and spoofed signals. Now, we are left with a subset that is free from outliers to the assumption. For each case, we can then perform different hypothesis tests such as those described above and in a previous paper [10]. For example, for the nominal case, we calculate the relative likelihood of the having genuine signals versus all spoofed signals. Finally, we compare the results from both cases.

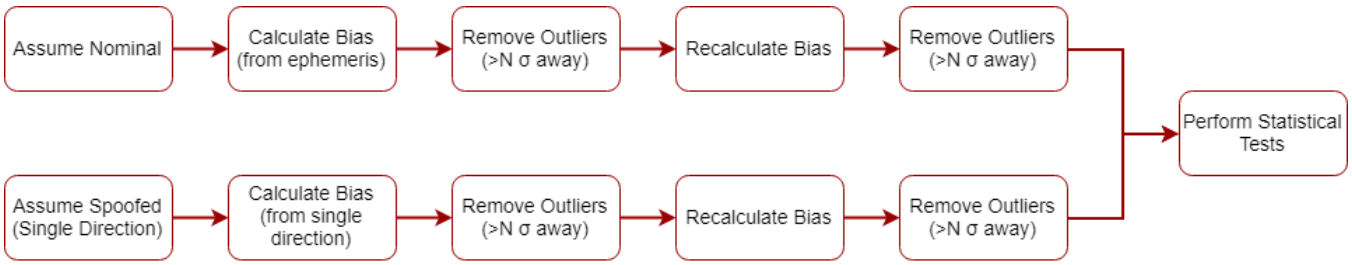

Figure 4. General process for processing ambiguity issue using two cases: best nominal and best spoofed.

Processing Statistical Results in an Executive Monitor

This comparison, typically done in an executive monitor (EM), considers many factors, such as the relative confidence and the number of outlier DOAs eliminated to make an overall decision. A simple EM is shown in Figure 5. For example, a simple EM may require that in any given case, 50% of non-high elevation satellites are retained. This means that, after rejecting high-elevation satellites, at least half of the remaining non-high elevation satellites must be retained for the statistical tests. Then, we must calculate a high relative confidence of the case's assumption (Equation 4 or 5). That is, the overall probability of the case must be greater than the detection threshold of each case – the probability of detecting nominal or spoofing (P_{nom_detect} or P_{spoof_detect}, respectively). The final decision may be made based on the table shown in Figure 5. If one case passes but the other does not, the decision is clear. If both cases fail, a conservative answer would be to not make a decision (i.e. inconclusive). If both the nominal and spoofed cases meet these criteria, then several possibilities exist depending on the application and safety level. One option is to choose the case in which the most satellites are retained, but this may result in missed detections, particularly when there are just a few spoofed satellites. Another option is to choose neither case and indicate that the results are suspicious. While this is not useful in itself, as a user may not know what to do with this information, it can be passed to other detection monitors to aid their results. A third option is to be conservative and to alert spoofing any time there is a reasonable (but not necessarily conclusive) chance that there is spoofing. This is not suitable for all applications, as it reduces nominal availability. But, for applications that can suffer some availability loss, this approach may be sensible.

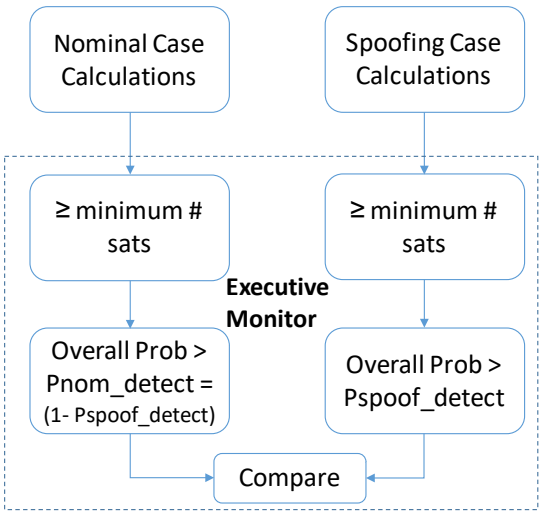

Executive Monitor Result Comparison

	Spoof Detect Passed	Spoof Detect Not Passed
Nominal Detect Passed	Compare each prob, # sats to decide	Nominal
Nominal Detect Not Passed	Spoof	No decision (leave for other tests)

Figure 5. Example Executive Monitor Processing and Decision Making

Nominal Scenario

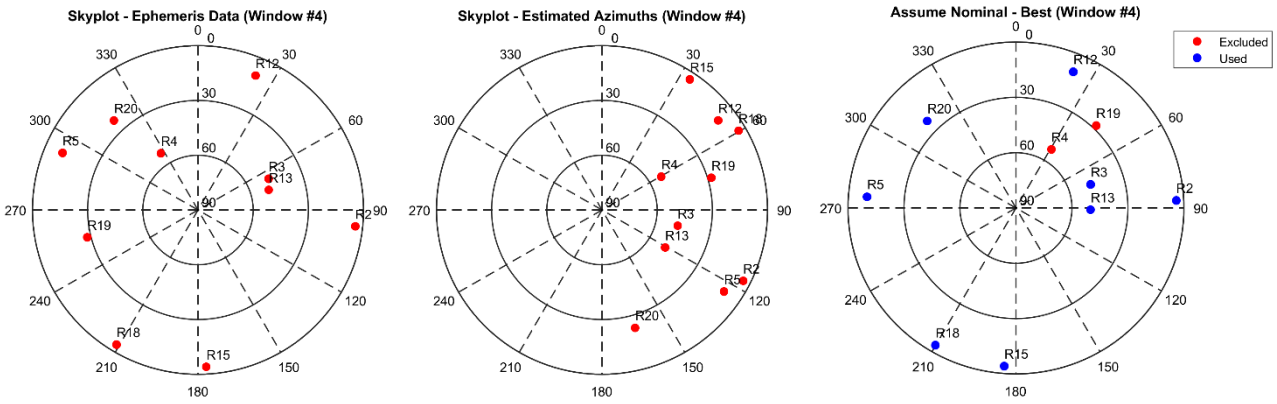

Figure 6. Skyplots of ephemeris data (left), DPA-measured azimuth and ephemeris elevation (middle), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity and heading bias assuming genuine signals (right) of GLONASS satellites under nominal conditions.

The concept is illustrated in the next few figures. Figure 6 shows a case where, with the exception of a few outliers, the corrected measured DOA match the expected ephemeris data well. The figure shows, from left to right, the skyplot based on ephemeris data (left), initially calculated DPA-measured azimuth (which is kept to between 0 to 180 degrees) and ephemeris elevation (middle), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity and heading bias assuming genuine signals – "best nominal" (right). GLONASS satellites are used in the example seen in the figure. The best nominal case is generated from the raw DPA measurements by choosing 180-degree phase shifts to best match the expected (relative) directions from the ephemeris data. Next, the difference between the expected directions and the phase-shifted DOAs is calculated and averaged resulting in an average ΔDOA, or bias. Recall that under nominal conditions, this bias will be roughly the same. This bias correction is then applied to all satellites to account for the unknown orientation of the antenna.

Spoofed Scenario

Under completely spoofed conditions, one expects all signals to arrive from a single direction, assuming the spoofer transmits from a single location and all the signals have been captured by that spoofer. The best spoofed case assumes completely spoofed conditions and chooses the 180-degree ambiguity value for each satellite to best match a single direction of arrival. Another way to put it is that we choose the ambiguity such that the relative DOA between measurements is as close to zero as possible. Hence, a heading bias is not required because the measured data is not compared to the ephemeris information. Instead, spoof detection is considered by averaging DOA measurements and analyzing the variance of the measurements. This concept is illustrated in Figure 7, which uses a similar set up as Figure 6, where despite some deviation due to imprecision in azimuth estimation, the signals appear to arrive from the same direction.

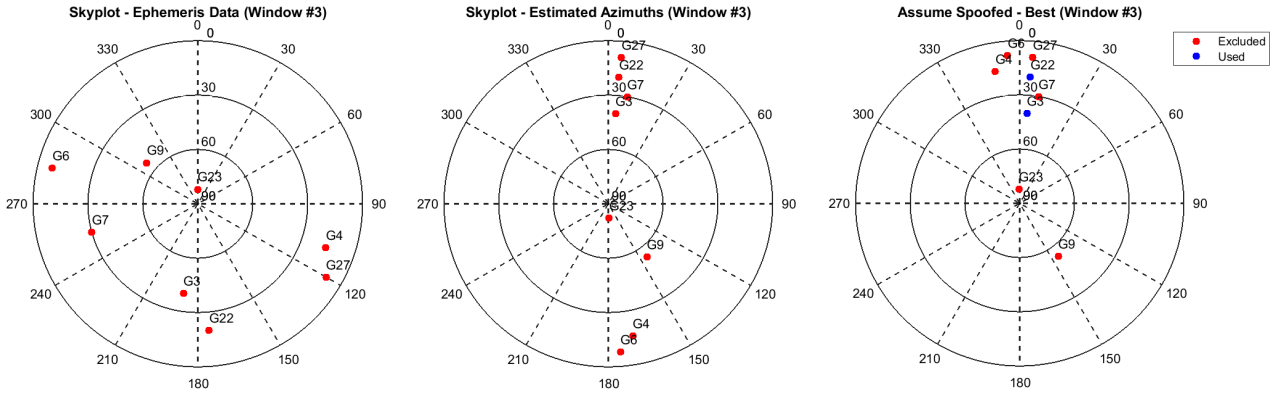

Figure 7. Skyplots of ephemeris data (left), DPA-measured azimuth and ephemeris elevation (middle), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming spoofed signals (right) of GPS satellites under spoofed conditions.

Outlier Rejection

With the ambiguities selected and initial bias estimates calculated for each case, outlier identification can be conducted. To determine what is or is not an outlier, we use the expected variation of the ΔDOA to see if a measurement deviates too far from expectations. As discussed earlier, the expected ΔDOA variance is larger under nominal than for spoofed measurements as the error in the spoofed measurements is correlated and should be mostly eliminated by estimating a bias. For the genuine signals, each DOA is derived from an independent signal and from different directions, the DOAs should not be correlated. Under our assumed spoofed conditions, signals arrive from a single direction and some signals would likely deviate strongly from the ephemeris data. Importantly, they derived from the same transmission and from the same direction; their measurement errors, such as from multipath, should be highly correlated. Hence, when examining the relative difference between the spoofed signal DOAs, the variance should be small as DOA measurement error is effectively the same, hence a bias, on all spoofed DOAs.

We mark as outliers measurements that deviate more than three standard deviations from expectation. In Figure 6 and Figure 7, these outliers are marked in red and the remaining satellites in blue. We note the number of outlier measurements and recalculate the bias excluding these measurements. The recalculated bias is then used to improve our ΔDOA for the remaining “good” measurements. Once outliers are removed, statistical tests can be performed to compare the two cases and to calculate their relative likelihood. It is important to note that the performance of the two cases are generally inversely related. That is, it is unlikely for any given set of GNSS signals to perform strongly under both best nominal and best spoofed assumptions as typically the satellites are well distributed through the sky. So, under nominal conditions, GNSS signals arrive from multiple directions.

Therefore, by comparing the performance of the two cases with a decision tree, one may be able to reasonably detect spoofing [10]. Of course, some scenarios are harder than others. There are situations and scenarios where genuine signals can naturally come from similar directions. However, genuine signals coming from approximately the same direction should have more variations than signals coming from exactly the same direction due to the error correlations mentioned previously. Additionally, if the likelihood of both cases is high in an absolute sense, one case should not be much more likely than the other in a relative sense. Hence the system would not provide a strong certitude of either state – which is the desired outcome.

Mixed Signal Scenario

Of course, the two assumed cases do not explicitly cover the scenario where there is both spoofed and genuine signals. There are many conditions that can result in having a mixture of both such signals. In order for an attacker to interfere with or spoof genuine GNSS signals, they must broadcast falsified signals with a power above that of the genuine signals. However, if the attacker's C/No is selected to be close to the genuine signal's C/No, the receiver may experience a mix of falsified and genuine signals. While this reduces the attacker's ability to reliably spoof location, it can be used to interfere with and reduce navigational ability on the receiver.

Spoof detection for mixed signal scenario is more difficult since signals arrive from multiple directions, but do not match the ephemeris data well. This may cause both the best genuine and best spoofed case tests to return either inconclusive results or results that test positive but for only after excluding many satellites. Similarly, under low C/No conditions the receiver may only pick up a few GNSS signals. Given only a few satellites, the likelihood that genuine signals from these satellites may arrive from the same direction increases. Therefore, the failure rate of the spoof detection algorithm increases as well. In particular, inconclusive results and test with a small number of satellites yield high rates of Type-I errors, or false alarms. One must be particularly careful to reduce the false alarm rate because it renders the navigational ability of the device useless and may incentivize users to mistrust the spoof detection. Given that DOA-based spoof detection will likely be implemented alongside other consistency checks, one strategy may be to indicate a potential anomaly but utilize other methods to corroborate the spoofing

One way to manage mixed case scenarios is through the outlier rejection process described previously. Under both the nominal and spoofed case assumptions, several satellites will deviate more than 3 standard deviations from their expected directions and will be rejected as outliers. Unfortunately, this may lead to a limited number of satellites for the statistical tests. Recall that the likelihood for scenarios where genuine signals naturally come from similar directions grows as the number of available signals shrinks. Hence, the best protection against mixed spoofing would be a combination of DOA spoof detection with other consistency checks.

EXPERIMENTAL SETUP

The analysis and validation of our developed techniques was conducted using the Stanford DPA with on-air test data collected from the roof of the Stanford University Durand Building and at the US government sponsored spoofing exercise in 2017. In both locations, the Stanford DPA was setup with a one-foot by one-foot ground plane and tested under multiple scenarios as shown in Figure 8.

In the tests, we use the Stanford DPA processed signals to determine DOA. The DPA generates a combined signal of the LHCP and a phase-shifted RHCP signal. The scan mode, controlled by the onboard microcontroller, uniformly steps through a full 360 degrees with a variable phase shifter. The combined signal is processed via a u-Blox receiver to determine its C/No. Later, the signal is post processed to find the phase-shift that results in a minimum C/No or null to get DOA.

In the field test with on air spoofing, the DPA was placed statically on top of a vehicle as shown in Figure 8. The DPA microcontroller was set so that a complete shift cycle took 256 seconds (180 degrees). This slower rate (rates of less than a minute and as low as 2.54 seconds were previously used) was used to increase fidelity of the measurements by allowing for a longer dwell time at each shift.

Figure 8. Two Dual Polarization Antennas (DPA) on SUV roof (Left) and DPA in lab (Right).

The spoofed GNSS signal was transmitted from a single spoofing location which was about 20-degree elevation relative to our vehicle. Due to our location, the received power of the transmitted signal was slightly but not greatly above that from the genuine GNSS signals. In most spoofing cases, the only GPS signals tracked by our receiver were spoofed. But, because of the power, there are some mixed cases where the receiver is tracking both genuine and spoof signals. This also exists sometimes because our 256 second phase shift period may contain both spoofing and non-spoofing times. Additionally, some signals were not spoofed such as those from the Wide Area Augmentation System (WAAS) and GLONASS satellites.

RESULTS

We analyze the performance of the ambiguity processing method described on over 300 different measurement sets. Four selected scenarios are described for illustration purposes. Scenarios 1 and 2 examine cases of accurate spoof detection for both a nominal and spoofed case. Scenarios 3 and 4 consider cases of mixed signals and inaccurate measurements to analyze failed spoof detection and highlight sources of error. These four scenarios are shown in Figure 9 to Figure 12. In each scenario, the top-left skyplot shows the ephemeris-derived azimuth and elevation. The top-right skyplot shows the initial DPA-measured azimuth (without accounting for the ambiguity) and ephemeris-derived elevation. The ephemeris-derived elevation angle is used simply to visualize the satellites on the skyplot and are not used in the calculation. The bottom-left skyplot shows the best-nominal case, where 180-degree phase shifts are applied to each satellite to best match the ephemeris data. The bottom-right skyplot shows the best-spoofed case, where phase shifts are applied to best match a single direction of arrival. The satellites excluded due to outliers are shown in red while the remaining satellites used for calculations are shown in blue. Note that high elevation satellites (above 65 degrees) are excluded as the DPAs measurements are not accurate above this elevation. Probabilities are then calculated using the remaining satellites as described previously and in detail in a previous paper [10].

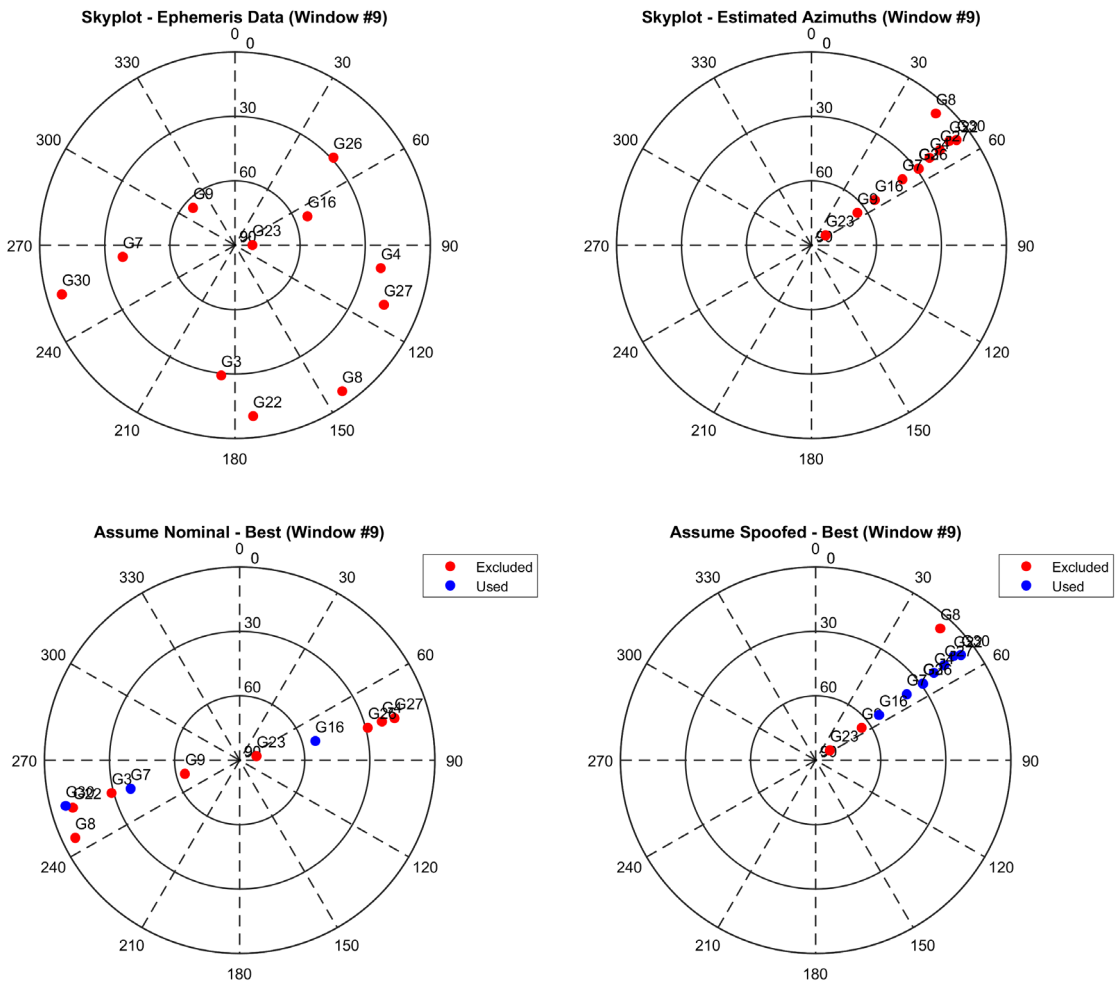

Figure 9. Skyplots of ephemeris data (top-left), DPA-measured azimuth and ephemeris elevation (top-right), DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming nominal signals (bottom-left), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming spoofed signals (bottom-right) of GPS satellites under spoofed conditions.

Scenario 1 examines a situation where all signals are spoofed and is shown in Figure 9. Due to high elevation angles, 2 satellites are eliminated from the calculations. Under nominal assumptions, 7 more of the remaining 10 satellites are rejected as outliers as they exceed 3 standard deviations from the mean ΔDOA. The remaining 3 satellites, shown in blue, lie within the acceptable tolerance and are used for calculations. Note from the geometry of the ephemeris data that the 3 that are retained are satellites whose lines of sight to our user are naturally aligned. So, while the probability of being genuine (Equation 4) is high, the probability of spoofing (Equation 3) will also be high. Hence, with this nominal set, and with both probabilities being high, the relative probabilities of each case relative to the other (Equation 4 or 5) will not be strongly conclusive (e.g. > 99%). This is not surprising as we have effectively removed satellites until only the ones that naturally align are used. So, they both match a genuine and a spoof hypothesis. However, we can tell this is the issue by considering the number of rejected satellites due to outlier detection versus the overall number of satellites. In the spoofed case, 8 of the 11 satellites are retained and provide a strong indication of spoofing. Hence, the outlier elimination indicates that the spoofed case is more likely, and we have strong probabilistic evidence from hypothesis testing of that case that it should be spoofing. Scenario 1 provides a strong example of accurate spoof detection with a calculated probability of spoofing of over 99.999999%.

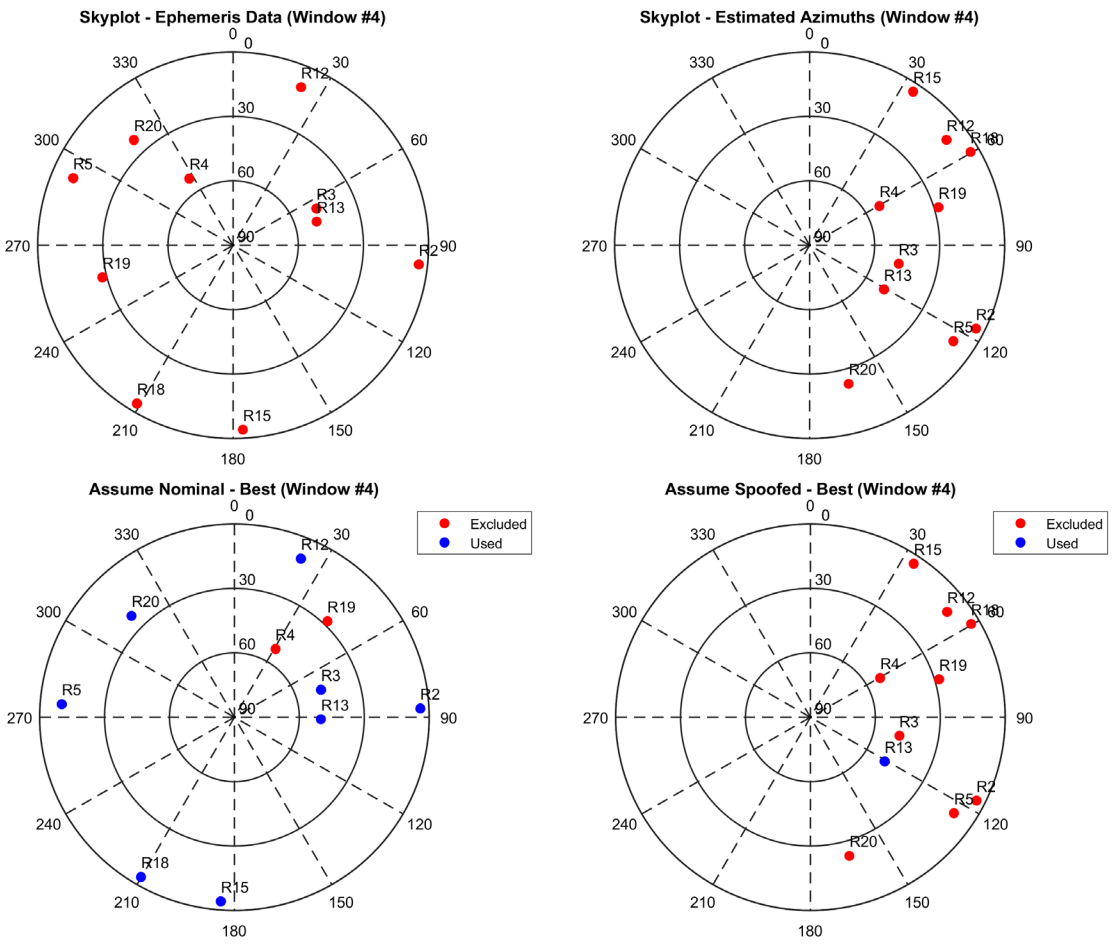

Figure 10. Skyplots of ephemeris data (top-left), DPA-measured azimuth and ephemeris elevation (top-right), DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming nominal signals (bottom-left), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming spoofed signals (bottom-right) of GLONASS satellites under nominal conditions.

Scenario 2 examines a situation where there are well measured nominal signals and is shown in Figure 10. In the nominal case, only 2 of the 10 satellites are discarded as outliers after applying the appropriate 180-degree phase shifts and bias corrections. On the other hand, in the spoofed case, nearly all of the satellites are discarded. Note that a single satellite is retained only as due to the trivial case where a single satellite appears spoofed as it arrives from a single direction. In this scenario, the results of the nominal case are given more significance than the results of the spoofed case due to the number of rejected outliers. The nominal case indicates a relative probability of nominal (i.e not spoofed) of 80.336% before accounting for rejected outliers. This results in accurate detection of a nominal case.

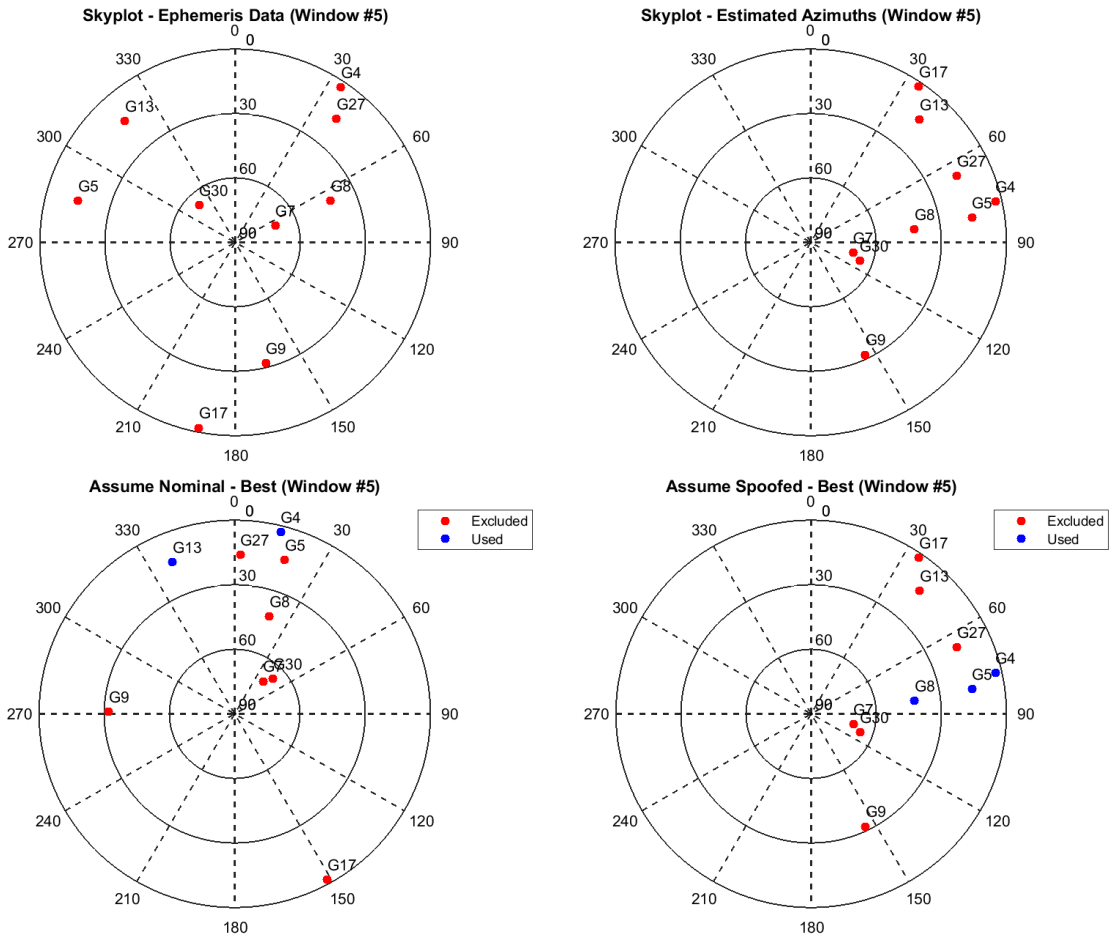

Figure 11. Skyplots of ephemeris data (top-left), DPA-measured azimuth and ephemeris elevation (top-right), DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming nominal signals (bottom-left), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming spoofed signals (bottom-right) of GPS satellites under nominal conditions.

Scenario 3 examines a situation where there are nominal signals with some inaccurate measurements and is shown in Figure 11. Comparing the ephemeris data with the measured azimuths (minus calculated bias offset), we notice that several satellites have large errors (> 30 degrees) in the measured DOAs. This error is likely explained by inaccurate null-detection or perhaps multipath errors. Unfortunately, these large errors propagate through the bias correction process. Due to the errors, the normalized ΔDOA distribution has a large variance. In the nominal case, this results in 5 of the 9 satellites being rejected as outliers. Note that another 2 satellites are rejected due to high elevation, leaving only 2 satellites for calculations. In the spoofed case, 3 of the 9 satellites are retained after outlier rejection, as their relative variance is much lower than the nominal case. If we examined the resulting calculated relative probabilities of each hypothesis, a larger number of satellites are retained in the spoofed case with a smaller variance, the spoof detection algorithm returns a false alarm, or Type-1 error. In this case, the executive monitoring can make one of several decisions along the lines discussed in “Accounting for the Ambiguity in Software.”

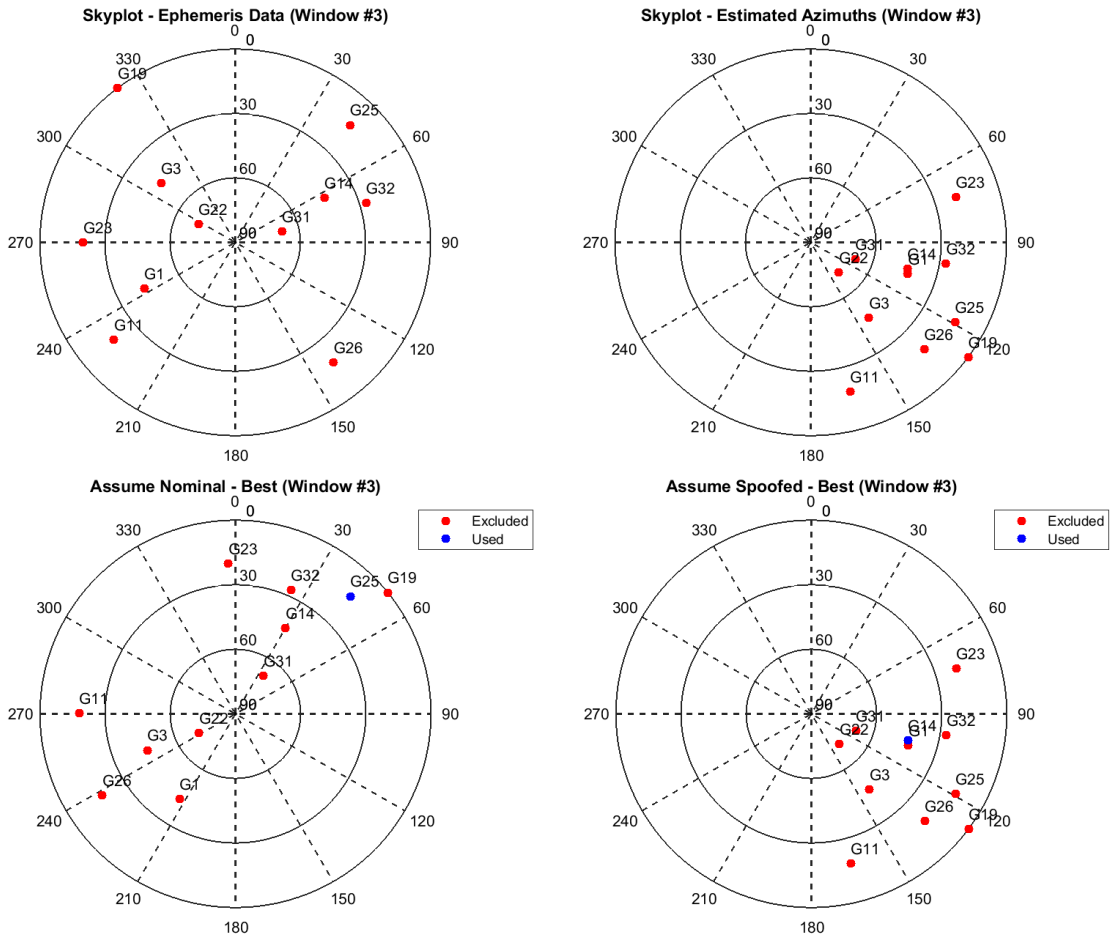

Figure 12. Skyplots of ephemeris data (top-left), DPA-measured azimuth and ephemeris elevation (top-right), DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming nominal signals (bottom-left), and DPA-measured azimuth and ephemeris elevation corrected for 180-degree phase ambiguity assuming spoofed signals (bottom-right) of GPS satellites under nominal conditions.

Scenario 4 examines a situation where we have a mix of genuine and spoofed signals and is shown in Figure 12. Once again, the normalized ΔDOA distribution has large variance under both nominal and spoofed assumptions. In an ideal situation, under nominal assumptions, all of the spoofed signals would be rejected as outliers. However, the spoofed signals cannot be identified a priori and must be included under the initial bias correction. This adds error to the average ΔDOA, or bias, that is applied to all satellites, pushing nominal signals outside tolerances. Similar to Scenario 4, this scenario rejects a large number of satellites in both the nominal and spoofed cases. Specifically, only 1 satellite is retained in both cases. Again, it will remain up to the EM to decide. For an aviation application, given the large number of rejected outliers, the spoof detection algorithm returns inconclusive. For this application, the result is likely preferred rather than alerting or not alerting without adequate confidence. The result would be to pursue further consistency checks.

Table 1. Performance of Spoof Detection Algorithm

Result	Percentage
% Nominal - Detected Correct	98.4% (314 of 319)
% Nominal - Detected Incorrect (false alert, Type-I error)	1.56% (5 of 319)
% Spoof - Detected Correct	47.9% (23 of 48)
% Spoof - Detected Incorrect (missed detection, Type-II error)	52.1% (25 of 48)
% All - Inconclusive	10.5% (43 of 410)

The above scenarios are a few illustrative examples of the various results of the test. Table 1 shows the results from hundreds of such scenarios under nominal, spoofed, and mixed conditions. Of the 410 total scenarios, there were 319 nominal scenarios, 48 spoofed scenarios, and 43 scenarios in which there were too few satellites retained by either case to determine if the receiver was spoofed.

An important result is that the false alert percentage is low. As discussed earlier, spoof detection that constantly renders the receiver inoperable is not a desired quality. Combining DOA-based spoof detection with other consistency checks, the missed detection statistic may also be improved significantly. The high nominal detection accuracy is also encouraging. The use of spoof detection to verify a nominal signal as genuine may also have important repercussions in safety of life applications. Recall that the tolerances selected in the EM can be used to trade off the selectivity of the detection in regards to false-alerts, missed detections, or inconclusive results. For example, one can select a less conservative EM to reduce the number of missed detections and inconclusive results, but at the cost of increasing the number of false alerts. As discussed earlier, the EM may be selected to best match the application of the receiver.

SUMMARY

This paper develops methods for mitigating the 180-degree ambiguity problem inherent in DPA direction of arrival measurements. Solving the ambiguity problem is important as it allows the DPA to provide high confidence spoof detection. This is attractive as the DPA can be a small form factor patch antenna and can be used in a variety of applications from aviation to automotive. In nominal conditions, this could provide heading, even when a vehicle is static. Under spoofing, the DPA, if it can robustly detect spoofing, helps make use of GNSS safe. The small size along with its static heading and spoof detection capability makes it attractive for autonomous vehicles.

The paper focuses on managing the ambiguity by creating two cases: best genuine and best spoofed. A process for deciding if spoofing exists is developed by first determining the number of satellites that may fit those cases, then calculating the relative probability of spoofing and finally having executive monitoring (EM) make a decision based on those results. Live spoofing measurements are made under different conditions and our algorithm is tested against these results. The results demonstrate that the technique manages the ambiguity well showing that the DPA approach has potential. However, it is not perfect, both with false alerts and missed detection higher than desirable. However, with better EM design or additional spoof detection mechanisms, the capability can be improved. The improvement is necessary to meet the stringent requirement of aviation and other safety of life applications.

ACKNOWLEDGEMENTS

The authors thank the Federal Aviation Administration (FAA) and the Stanford Center for Position Navigation and Time (SCPNT) for sponsoring this research. The authors gratefully acknowledge the support of the Stanford Research Experience for Undergraduate (REU) program. The authors also thank the US government providing us with an opportunity to test under

live GPS spoofing.

REFERENCES

[1] 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

[2] Paul Y. Montgomery, Todd E. Humphreys, Brent M. Ledvina, “Experimental Results of a Multi-antenna Receiver Defense Against a Portable Civil GPS Spoofer,” ION ITM, Anaheim, Jan 2009

[3] M. Psiaki, Brady W. O'Hanlon, Steven P. Powell, Jahshan A. Bhatti, Kyle D. Wesson, Todd E. Humphreys, Andrew Schofield, “GNSS spoofing detection using two-antenna differential carrier phase," Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2014), Tampa Bay, Florida

[4] Appel, Manuel and Konovaltsev, Andriy and Meurer, Michael (2016) *Joint Antenna Array Attitude Tracking and Spoofing Detection Based on Phase Difference Measurements.* In: Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016). ION GNSS+ 2016, Portland, Oregon

[5] Yu Hsuan Chen, Sherman Lo, Per Enge, " Using Multi-Antenna Iridium Measurements for Rapid Spoof Detection," Proceedings of the Institute of Navigation Joint Navigation Conference, Orlando, FL, June 2015 (Presentation Only)

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

[7] Trinkle, Matthew, Cheuk, W-C, "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

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

[9] Adrien Perkins, Yu-Hsuan Chen, Wei Lee, Sherman Lo, and Per Enge, “Development of a Three-Element Beam Steering Antenna for Bearing Determination Onboard a UAV Capable of GNSS RFI Localization,” Proceedings of the ION GNSS+ 2017, Portland, OR

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

[11] T. Kraus, F. Ribbehege and B. Eissfeller, "Use of the Signal Polarization for Anti-jamming and Anti-spoofing with a Single Antenna," Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation, Tampa, FL, September 2014, pp. 3495-3501.