Demonstrating Single Element Null Steering Antenna Direction Finding for Interference Detection

Yu-Hsuan Chen, Sherman Lo, Adrien Perkins, Fabian Rothmaier, *Stanford University*
Dennis Akos, *University of Colorado at Boulder*
Per Enge, *Stanford University*

BIOGRAPHY (IES)

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.

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 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.

Fabian Rothmaier is a Ph.D. candidate in the Stanford GPS Laboratory. He is currently working for the European Space Agency

Dennis Akos is professor in Aerospace Engineering Sciences at the University of Colorado, Boulder. He received his Ph.D. in Electrical and Computer Engineering from Ohio University in 1997.

Per Enge is a professor of Aeronautics and Astronautics at Stanford University, where he is the Vance D. and Arlene C. Coffman Professor in the School of Engineering. He directs the Stanford GPS Laboratory, which develops satellite navigation systems. He has been involved in the development of the Federal Aviation Administration's GPS Wide Area Augmentation System (WAAS) and Local Area Augmentation System (LAAS).

ABSTRACT

Use of Global Navigation Satellite Systems (GNSS) in safety of life applications such as aircraft navigation, railway control and autonomous vehicles is increasing as these technologies become more necessary or mainstream. To serve these applications, GNSS must provide high integrity, even in the face of deliberate attacks such as spoofing. The Stanford dual polarization antenna (DPA) is a technology that uses a single patch antenna but with two feeds to examine both left and right hand circularly polarized (LHCP and RHCP, respectively) signals. Proper design and installation allows the DPA to use polarization of the incoming signal to determine direction of arrival (DOA) and elevation. These measures can be used to discriminate a genuine from a bogus broadcast. The technology can also null out some interference signals.

This paper examines the continued development and field tests of our DPA. Test with genuine on-air signals have shown the ability to determine DOA and elevation. Field test in on-air spoofing and jamming conditions were also conducted. These scenarios allow us to demonstrate the performance of the DPA processing, DOA and elevation estimates.

INTRODUCTION

The openness of Global Navigation Satellite Systems (GNSS) satellite signals has enabled its rapid adoption worldwide. It allows a variety of manufacturers to develop and improve receiver designs enabling a variety of applications. However, this openness also makes GNSS vulnerable to attacks such as spoofing. This threat will only increase in the future. The means of conducting such an attack are becoming more obtainable. For example, security experts with basic GNSS knowhow were able

to develop a low cost, flexible spoofer [1]. Furthermore, the incentives for such attacks increase due to increased economic and safety uses of GNSS. Indeed, deliberate GNSS spoofing attacks are no longer theoretical or purely in the military domain. In the past year, spoofing attacks have been seen in the Kremlin and Black Sea [2] [3].

Given the severity of spoofing on safety of life and economic activities, some anti-spoofing (A/S) mechanism is desirable in any critical GNSS receiver. One important A/S function is to detect the presence of spoofing. Such detection is usually based on finding telltale signatures left by a spoofing signal that differ from the genuine. Different categories of detection techniques have been devised [4][5][6]. In this paper, we develop and analyze detection using a dual polarization antenna (DPA) to examine the signal in space properties. Specifically the DPA provides right and left-hand circularly polarized (RHCP and LHCP, respectively) components of the received signals. For a good DPA design and installation and a ground-based (i.e. low elevation) spoofer, these properties will often differ between the genuine and spoofed signal. These differences allows for determining the direction of arrival (DOA) and the rough elevation of the incoming signal. A single antenna spoofer can thus be detected as all its spoofed satellite signals will come from the same direction.

Miniaturized versions of our DPA suitable for field-testing were developed [7]. These are built on a printed circuit board (PCB) and utilize surface mount components to combine the signals from the two feeds to a hybrid coupler to create both a RHCP and LHCP output. The design provides a small form factor antenna, suitable for aircraft installation that has elevation dependent sensitivity to an incoming signal.

To demonstrate and quantify performance, these antennas were tested in several on-air scenarios. Tests with nominal, genuine GNSS signals only as well as with on-air spoofing and jamming were conducted. The paper details the signal processing and algorithms used by our DPA to determine DOA for spoof detection system. It also demonstrates the performance in various scenarios including on-air jamming and spoofing. Specifically, it shows the ability of the DPA in determining DOA.

BACKGROUND

One means to detect spoofing is to examine the physical properties of the incoming signal. Genuine GNSS signals have specified directions of arrival and specific polarizations (RHCP). Multi-antenna techniques have been suggested to examine DOA to detect spoofing [7][8][9]. Single spoofing antennas can only generate one DOA whereas the genuine satellite signals come from many DOAs. DOA-based spoof detection can be powerful but such techniques require either spatially separated antennas or multi-element antenna arrays. Large spatial separations require additional space and are more costly to install. For aviation installations, each antenna would need a separate hole and cable run through the aircraft body. Multi-element arrays have similar drawbacks as well as being restricted by International Traffic in Arms Regulations (ITAR) should there be four or more elements. An antenna sensitive to polarization to detect non-RHCP signals as being not genuine or multipath. Mayflower communications proposed such a concept for aviation spoof detection in the 1990s. However a spoofer can, with a little more work, replicate the polarization. The Stanford DPA concept addresses these limitations. It is small, utilizes a patch antenna and may be installed like a standard GNSS antenna. While it measures polarization, it does not rely on the spoofer using non-RHCP signals. Instead, it uses the measurements to determine DOA for spoof detection.

The Stanford DPA design generates and uses the RHCP and LHCP components of a signal to determine the presence of interference and spoofing. GNSS signals are RHCP and generally come from above the antenna. The spoofing signal needs to be also seen as RHCP to be consistent. However, just generating a RHCP signal is not enough to fool the Stanford DPA if installed properly. This is because signals that impact the antenna ground plane before entering the antenna becomes linearly polarized regardless of their initial polarization. A linearly polarized signal has equal RHCP and LHCP components; the antenna can use this to determine and cancel spoofing. So any spoofing signal that comes in from the level of the vehicle (automobile, aircraft) or below will impact on the ground plane first. Hence, their orientation from below the horizon will be detectable by the DPA and this information can be used to detect spoofing. The concept is detailed in [10] and shown in Figure 1. This forces the attacker to take a position above the vehicle, perhaps using an unmanned aerial vehicle (UAV), which is much more challenging. Static spoofing placements such as on rooftops may be used but are limited in range.

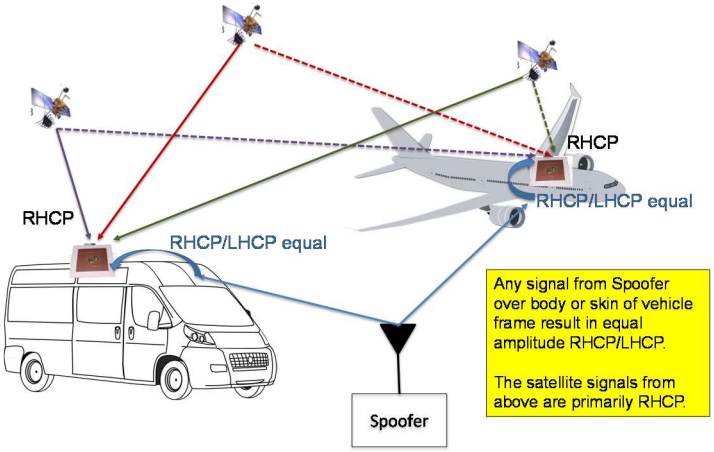

Figure 1. Dual Polarization Antenna Concept for Spoof Determination

Having demonstrated the capabilities using discrete components [10], Stanford developed a DPA design on printed circuit board (PCB) using commercial off the shelf (COTS) components [7]. The PCB DPA is shown in Figure 2 and processes the incoming signal and generates RHCP and LHCP components. It also contains all the necessary circuitry to monitor the LHCP signal, find DOA and mitigate interference. The RHCP is provided to the nominal receiver/radio while a combined LHCP/RHCP signal is fed to a monitor receiver that is controlled by executive monitoring (EM). The monitor receiver in our PCB is a u-blox receiver while the EM is implemented on an onboard microprocessor. In spoof detection mode, the EM examines the components to determine if there is interference or spoofing. One way this can be done through comparison of satellite carrier to noise ratio (C/No), DOA or both. Should an unwanted signal be present, the EM can change to mitigation mode where it can attempt to cancel the unwanted signal. The DPA has less than four elements and hence is suitable for export under current ITAR.

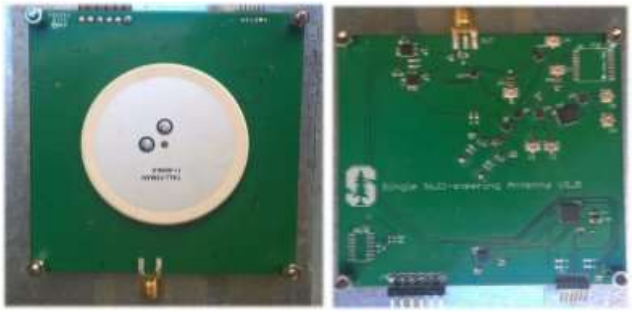

Figure 2. Stanford PCB Dual Polarization Antenna

Stanford PCB DPA

The Stanford DPA is built using a standard COTS patch antenna with 2 feeds (x and y axis) along with COTS surface mount components. The feeds outputs go to a 90° hybrid coupler, which combines them to create both a RHCP and LHCP signal.

With our DPA concept, the vehicle body performs an essential role as extended ground plane. A high elevation signal is not greatly affected, as it does not impinge on the body before entering the antenna. Hence its signal should be mostly RHCP. Incoming low elevation signals, such as from low elevation satellites and spoofing signals, impinge on the body before entering the antenna. The body causes the signal, regardless of transmitted polarization, to become linearly polarized (LP), thus having roughly equal RHCP and LHCP components. As our antenna can separate these components, it can use these components for spoof detection and interference mitigation.

We see this sensitivity to LHCP signal in conventional RHCP antennas. When illuminated by a purely RHCP signal, these antennas will still have LHCP energy due to antenna imperfections, ground plane and other conductive elements. This is quantified by its cross polarization discrimination (XPD) which is the ratio of co-polarized to cross polarized energy and is

dependent on the incoming direction of the signal (azimuth and elevation). For a RHCP antenna, co-polarized is RHCP and cross polarized is LHCP. XPD in decibels (dB) is given in Equation (1) where G_d is the gain of d polarization and θ, φ are elevation and azimuth of the incoming signal, respectively. Typically, the XPD for a GNSS antenna is high at high elevation and close to zero at low elevation showing how the ground plane and other components can produce in LHCP energy. It is this observation that help guide the development of the DPA.

$$XPD(\theta,\varphi) \triangleq [G_{RHCP}(\theta,\varphi) - G_{LHCP}(\theta,\varphi)]dB \quad (1)$$

The DPA with a large ground plane, by being able to measure the phenomena mentioned, can provide useful information about the incoming signal. First, its sensitivity to an incoming signal is elevation dependent because of the ground plane effect. An incoming RHCP signal from high elevation is mostly RHCP with little LHCP component. At lower elevation, the ratio of LHCP to RHCP energy increases such that they may be roughly equal around the horizon. Hence, a comparison of the energy of the components can provide rough elevation angle. One way to do this is by combining the signals in a destructive way.

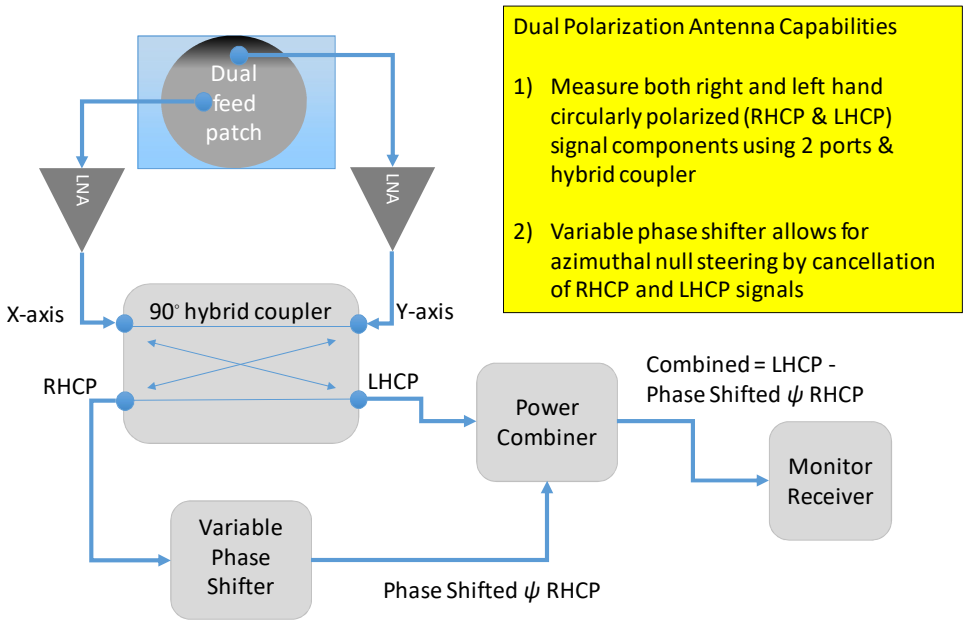

Figure 3. Basic design of the Stanford DPA with dual port patch antenna, low noise amplifier (LNA), 90° hybrid coupler, and variable phase shifter. The figure shows the DPA is set to spoof detection mode where the microprocessor steps through phase shifts to determine direction of signal arrival.

Second, the DPA can provides DOA as the RHCP and LHCP signals are offset by a phase shift that is related to the azimuthal direction of the incoming signal. This happens are part of the process of the ground plane forcing signals that traverse it to be LP. This results in the signal having "RHCP and LHCP energy that is phase coherent and similar in magnitude [10]." Given coherence and roughly equal energy, we can cancel the signal components simply by choosing the correct phase. A phase shift is applied to one of the signals, in our case, the RHCP. Our DPA microprocessor commands the rotation of the RHCP signal throughout all phase shifts, ψ, via the variable phase shifter and generates a combined signal as seen in Equation (2).The phase shift required for best cancellation also indicates the relative DOA of the signal providing another signature to identify interference. The microprocessor in the Stanford DPA sets the amount of time required to step through an entire cycle (360 degrees of azimuth scan).

$$combined\ signal(\psi) = LHCP - phase\ shifted(\psi)\ RHCP \quad (2)$$

The phase shift angle is related to the DOA or azimuth angle of the incoming signal as shown in Table 1 and Equation (3) with ψ being the variable phase shifter value, φ_0 being the azimuth angle of the x-axis feed and φ being the azimuth angle of the null [10]. Because φ_0 is unknown, our measurements are relative DOA unless a known reference is used to determine φ_0. The additional 90 degrees (°) is an artifact of creating the circularly polarized signals using x- and y-axis signals through the hybrid coupler. Rotating through 360° of phase shift tests 720° (two cycles) of incoming signal DOA. Another way to look at it is that

the combined DPA signal essentially has a dipole response to the direction of the incoming signal. Rotating through all phase shifts results in variations in C/No as it cycles through constructive and destructive combination of the LHCP with the RHCP. The phase shift, ψ_{null} that results in a null or minimum C/No as determined by the signal tracking is used to estimate the DOA via Equation (2). The DPA was initially tested using live GNSS signals on static and vehicle platforms to verify its performance on low and high elevation satellites. Figure 4 shows such a test where the DPA provides cancellation of low elevation signals, in this case, GNSS satellites. Hence the antenna can both detect, by finding LHCP components, and cancel spoofing from low elevation.

Table 1. Azimuth angle (relative to the top of the antenna) and corresponding phase shift

Azimuth angles (DOA) (φ)	*Variable phase shifter value (ψ)*
0° & 180°	90°
45° & 225°	180°
90° & 270°	270°
135° & 315°	0°

$$\psi = 2(\varphi - \varphi_0) + 90° = f(\varphi) \qquad (3)$$

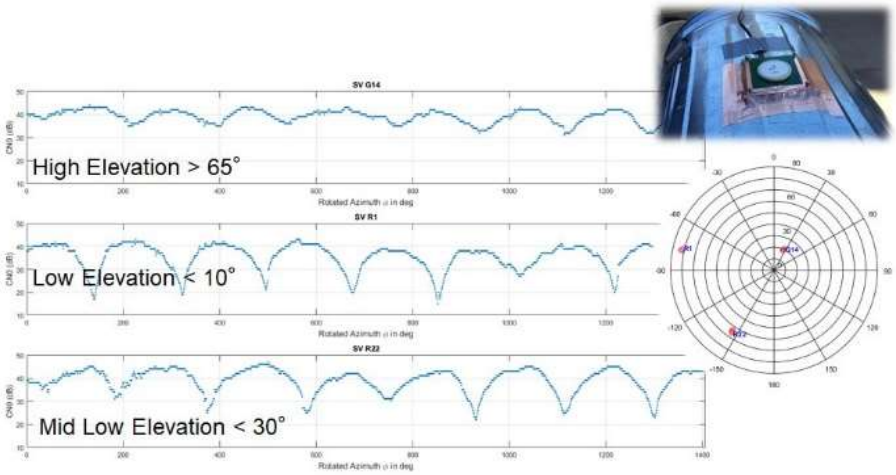

Figure 4. C/No vs. Rotation of LHCP component for different satellites: high (top), low (bottom), very low (middle). At low and very low elevation, rotation of LHCP component to the correct angle can significantly cancel out RHCP, ~ 180 second to complete all rotations

Determining DOA from Combined Signal

Determining DOA means devising an algorithm to find the null or minimum C/No in our data. The algorithm is illustrated in Figure 5 which shows the variation of C/No of the Wide Area Augmentation System (WAAS) geostationary satellite, PRN 138. The algorithm starts by finding a suitable period of time where reasonable C/No exist. The algorithm then selects a window of time where at least 360° of azimuth scan is made (the blue dotted line on the figure). This guarantees that at least one null exists within the window. For the window, a C/No threshold is determined and used to find a segment which contains the null and no peaks. The algorithm thus finds a segment that crosses over the threshold twice with samples between the crossings below the threshold. All the samples between two crossing points, indicated by the circles in the figure, are used to generate a curve fit. We use a third order polynomial to fit the curve for simplification and dealing with unbalanced or non-symmetrical curves. An unbalanced curve is one whose falling and rising edge are not similar (have different gradients), so a second-order polynomial does not adequately fit the curve. The curve should be balanced but because of receiver averaging, there is a lag in the results causing the asymmetry. This unbalanced curve can be seen in Figure 5. A second-order function may be used for some signals and has a lower computational load but it does not handle highly unbalanced curves well. The resulting fitted polynomial is then used find the minimum or null, indicated by a star (*) in the figure. The time associated with the null corresponds to a phase shift and a corresponding direction of arrival.

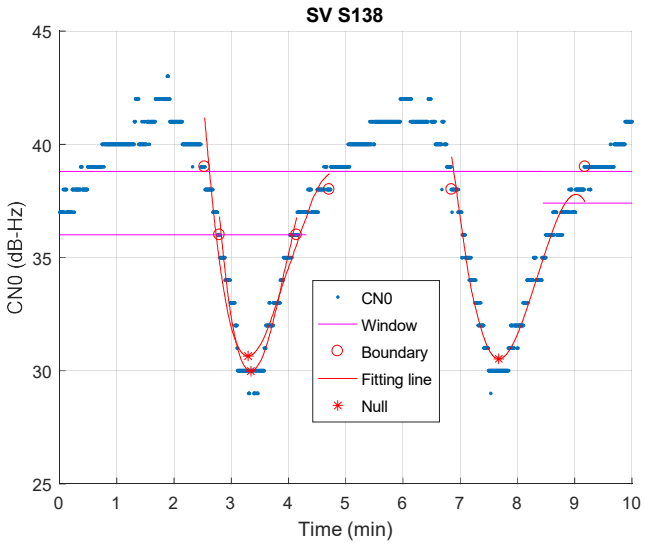

Figure 5. Illustration of algorithm for determining time of minimum C/No (null) applied to captured data. Example shows on-air C/No of the WAAS Geostationary satellite pseudo random number (PRN) 138 over a period of 10 minutes (roughly 2 cycles of phase shift). The elements of the null finding algorithm is shown in the plot as well as the determined null

EXPERIMENTAL EQUIPMENT & FIELD TESTING

Several field tests were conducted to further evaluate the Stanford DPA performance. The first tests, utilizing only broadcast satellite signals, were conducted at Stanford University. Static test were conducted with the antenna installed on a metallic trash canister or a large metallic plate to simulate the ground plane provided by the body of an aircraft or other vehicle. These were tested on the roof top of the Stanford Aeronautics & Astronautics building, about 50 feet above the ground. This allows for good line of sight to low elevation satellites. The setup is shown in Figure 6. Tests in nominal conditions provide a good indication of performance for high, medium and low elevation signals. An example of the static test result is seen in Figure 4.

Figure 6. Static testing at Stanford rooftop with flat ground plane

We also tested at a government sponsored exercise where different L1 spoofing test scenarios are conducted. Numerous scenarios were performed. Position and time spoofing, similar to those described or demonstrated in [6][11], were exercised both with ramps and jumps in deviation. Interference was often introduced prior to spoofing to knock off receivers under test. This also allowed us to test the interference detection and mitigation capabilities of the DPA. The tests were conducted statically with participants located such that the spoofing strength is slightly higher than the received GNSS signal strength. The C/No of the spoofed signals were kept the same so that participants can quickly confirm if their receivers are tracking the spoofed signal.

The field tested DPA are typically installed on the roof of our test vehicle such as the sports utility vehicle (SUV) as on a box similar to that shown previously in Figure 6. The antenna is 3 in x 3 in with a 12 in x 12 in ground plane. Placement on the top of the vehicle was expected to extend its ground plane. The typical DPA setup, seen at the front of the vehicle, has the combined signal processed by a monitor receiver (u-blox). The receiver is single-frequency on L1 band and set to track GPS, GLONASS, Galileo, and Beidou satellites. The u-center software provided by u-blox logged the signal strength, positioning results and other measurements. The microprocessor was programmed so it cycled through all ψ every 256 seconds (4 minutes 16 seconds). We also fielded a second DPA, seen at the back of the vehicle, which utilizes an Ettus universal software radio peripheral (USRP) to record raw intermediate frequency (IF) data from the LHCP and RHCP ports rather than a monitor receiver processing the combined signal. This data is stored on a solid state drive (SSD) of an Intel Next Unit of Computing (NUC) within the plastic housing on the roof. This data is not discussed in this paper.

For analysis, the monitor data is processed to determine DOA. The C/No result shows peaks and nulls with time due to the null-steering effect. We used the previously discussed curve-fitting technique on the C/No results to find the signal source DOA. Due to 256 sec/cycle scan rate, we have plenty of samples on each cycle for finding the optimal DOA. Different spoofing scenarios were investigated including time/position jump or walk. The jammer or spoofer was located at medium low elevation angle (15-20 degrees) respect to antenna. While these elevations result in the C/No null being not very deep (only around 10 dB peak-to-null), it is still adequate for DOA determination and the antenna can still detect spoofer and jammer in this situation. Should the antenna finds a number of satellites pseudo random numbers (PRNs) from the same DOA, this is a strong indication that there is a spoofing transmission. This information thus can form the foundation of a spoof monitor along with null depth information which is correlated to elevation angle of the incoming signal. Furthermore, we could potentially excluding these PRNs to allow us to develop an unspoofed position from pseudo ranges that are genuine. For the jammer, we find an angle of direction has highest C/No which corresponds to null steered to jammer direction.

FIELD RESULTS

On-air (Stanford) Test Results

Testing on GNSS signals allows us to study DOA performance as a function of elevation. The signal also can represent the effect of a sophisticated spoofing signal with the only difference being that each genuine signal comes from different directions while all spoofed satellite signals come from the same direction. Figure 4 shows results from static testing. For a low elevation satellite (< 10°, middle of the figure), the LHCP energy will be close in magnitude to RHCP energy resulting in deep fades in the combined signal. For a mid to low elevation satellite (< 30°, bottom of the figure), the LHCP energy is still reasonably close in magnitude to RHCP energy resulting in clear fades in the combined signal though one can see the fade getting weaker as the satellite goes up in elevation. The ground plane effect is less for higher elevation satellites and hence the signal becomes more purely RHCP. This is seen more clearly for a high elevation satellite (> 65°, top of the figure) where there are periodic dips but of only a few dB. So we can still get rough DOA and furthermore, from the three plots, the depth of the null provides a coarse indicator of elevation.

Figure 7 shows the estimated azimuth angle results of the part of a low elevation satellite pass. Note that only relative angle is determined and so a reference using the WAAS geostationary satellite is used. This can result in a bias. It shows the ability to estimate DOA from C/No variations. The error at higher elevation seems a little worse. Also, there are larger deviations at the horizon. This is not surprising as the gain pattern of the antenna is more disturbed. The results of Figure 4 suggest that the DOA estimates should get better with decreasing elevation, this is not clear in Figure 7. Aggregating data over several satellites, one can get the statistics of the variation as a function of elevation. Figure 8 shows the mean and standard deviation of estimated azimuth error as function of elevation. There is rising trend in standard deviation, especially after 50° of elevation. There is also greater variation at the horizon. These results depend on our phase step rate, which sets our dwell time at each phase, the

ground plane effects, and the antenna. The phase shift cycled through 360 degrees of azimuth in 51.2 seconds. This is 5 times shorter dwell time per phase than used in the on-air spoofing tests.

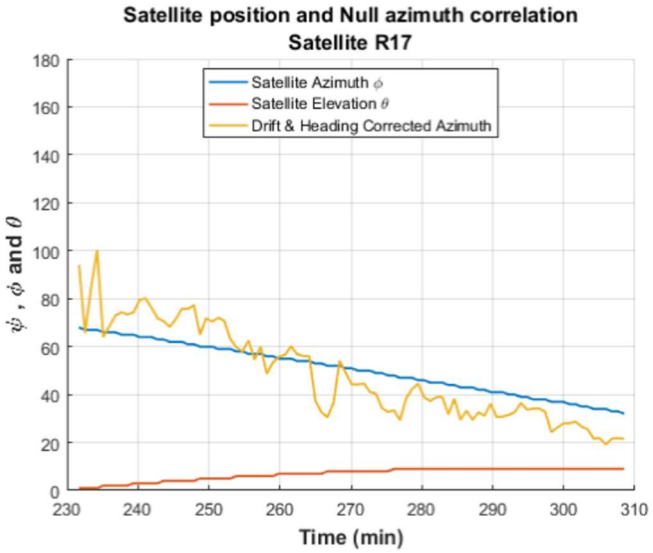

Figure 7.Estimated versus actual azimuth angle over part of a satellite pass. Elevation angle (from ephemeris) is shown for reference

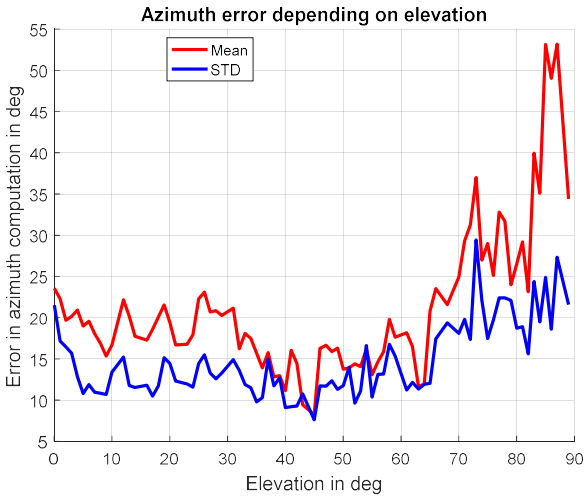

Figure 8.Mean and standard deviation of the error on estimated azimuth as a function of satellite elevation angle during static testing

Spoofing Test Results

The DPA operates on the physical properties of the signal to detect spoofing. Hence, it theoretically should not care what the spoofer is actually transmitting in terms of position and time. Indeed, as seen before, it is not detecting spoofing effects on ranges or position but 1) presence and coherence of LHCP and RHCP signals (as seen by the C/No variations) and 2) the azimuth of the incoming signal. Hence we get results with genuine GNSS signals as well as with live spoofing signals. These difference results between genuine and spoofed signals can be analyzed to help design robust detection monitors in the future. At the government test, many different spoofing scenarios were conducted. In general, the DPA behaved similarly in each scenario provided that the received spoof power was similar. An example of our results follows.

Figure 9 shows the positioning error starting from Scenario 1. The spoof scenario is a position push like that described in texbat scenario 4 [11] and starts at 6th minute of the figure. There is no jamming prior to spoofing. The positioning results show the

relative position to first fix. The position is spoofed 10 m to the east from true location. There are also larger variations in height after the spoofer is turned on. Figure 10 shows the C/No for all GPS satellites or pseudorandom numbers (PRNs). In the first 5 minutes when spoofer is off, the genuine signals have nulls at different times with different depth indicating a different azimuth directions, and hence locations. After the spoofer goes on during the 6th minute of the plot, all the C/No start to align together with similar nulls at the same time. While the similar signal strength is an artifact of the testing to help corroborate that the receiver is spoofed, having the null at the same location should occur regardless of the relative power of each spoofed satellite signal. Hence, even without doing DOA calculations, the system can see that something is amiss. Interestingly enough, going back to Figure 9, we see this also results in periodic positioning pattern in every 256 seconds when all PRN have similar signal strength. For example, the height error has peaks in the 6th, 11st, 15th, 19th, 23rd minutes.

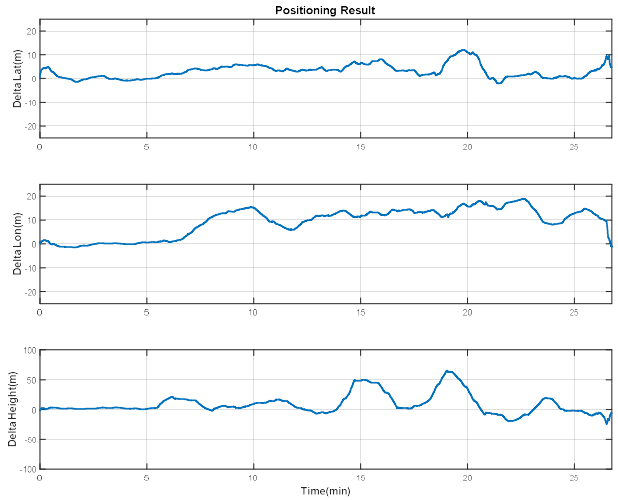

Figure 9. Positioning error with Dual Polarization Antenna during Spoofing for Scenario 1

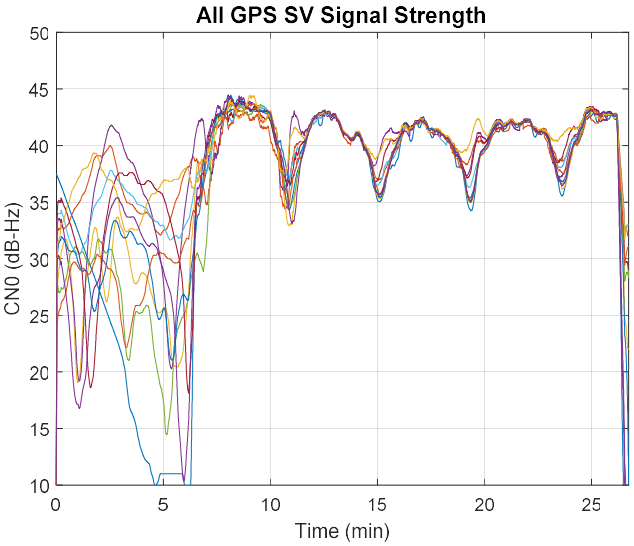

Figure 10. C/No for all GPS PRNs during Spoofing (starting at minute 6) for Scenario 1

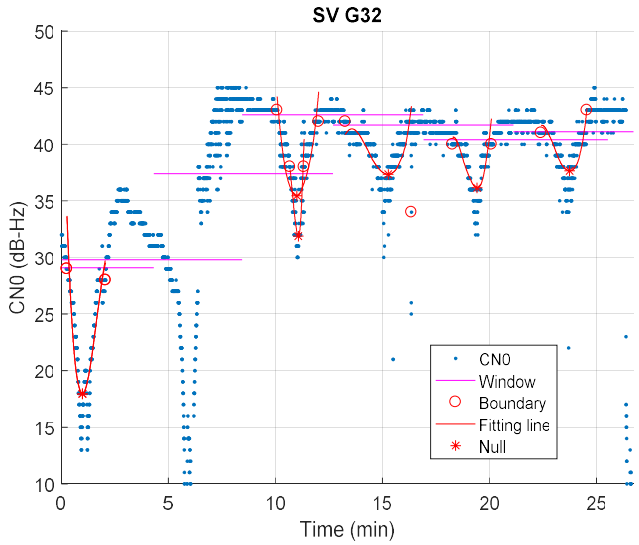

Figure 11. Signal strength (CN0) for GPS PRN 32 during Spoofing (starting at minute 6) for Scenario 1

We next examine the effect on individual satellite signals. Figure 11 shows the C/No of GPS satellite PRN 32 during the same time interval as in the previous figure. It also shows the curve fit that is used for estimating the azimuth of null. Recall that spoofing starts at minute 6. Due to the test design, the spoofing signal has a 8 dB higher C/No than genuine one (45 dB-Hz vs. 37 dB-Hz at peak). The null depth of spoofed signal, which is defined as highest to lowest signal strength, is 5 dB shallower than genuine one (15 dB vs. 20 dB). This is because that the spoofer on the tower is about 5 degree higher than satellite elevation (20 degree vs. 15 degree) and null depth should increase with decreasing elevation angle.

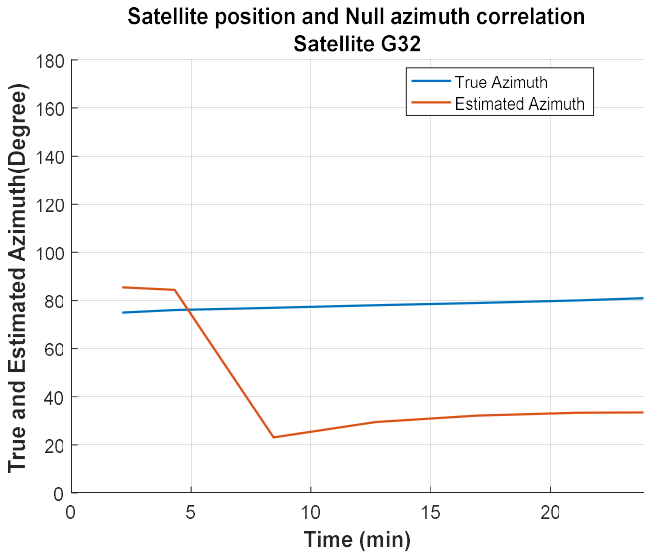

Figure 12. True (Almanac) and Estimated (from DPA) Azimuth or DoA of PRN 32 for Scenario 1

Figure 12 shows the signal direction of arrival estimation result over time for GPS PRN 32. Recall, that the DOA estimate is a relative estimate. So to get the absolute azimuth, the WAAS satellite, which is not spoofed, is used (see next paragraph). After the 6th minute when spoofer is turned on, the estimated azimuth changes from 90 to 30 degrees. As the azimuth is only estimated once every cycle, this change is a step transition. Even without an absolute reference, such a jump would tip us off to the start of spoofing.

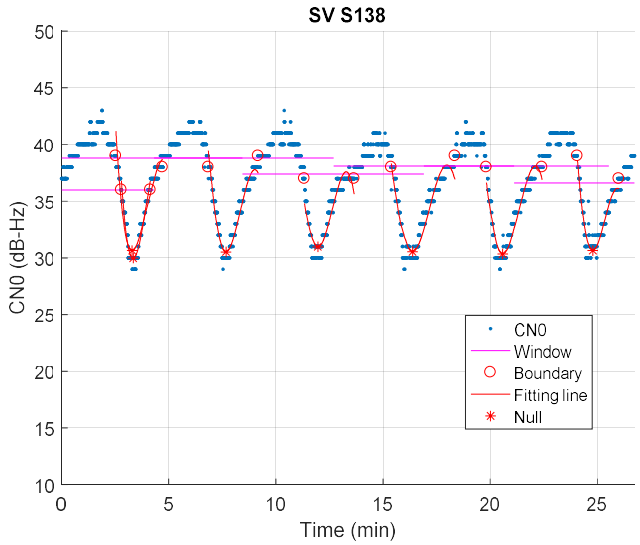

Figure 13. C/No of WAAS satellite PRN 138 for Scenario 1

Figure 13 shows the C/No of WAAS satellite PRN 138, which is not spoofed, over the test interval and curve fit used for finding the nulls. The C/No and null location are unchanged when spoofer is turned on at the 6th minute.

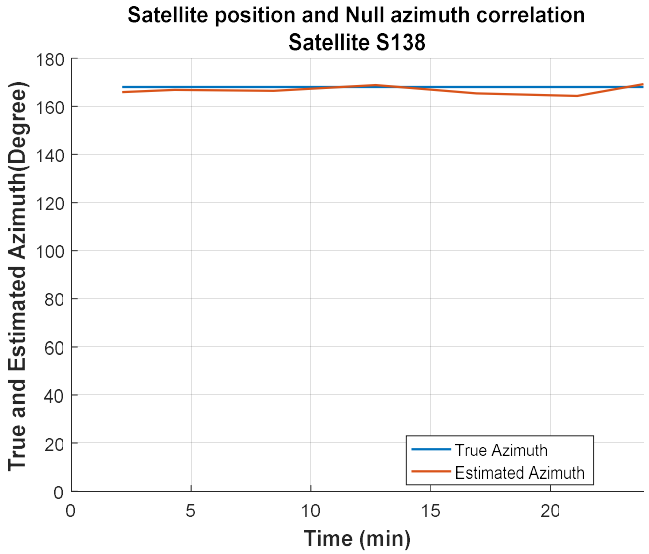

Figure 14. Direction of Arrival Estimate of WAAS satellite PRN 138 for Scenario 1

Figure 14 shows the signal direction of arrival estimation result over time for WAAS PRN 138. The estimated azimuth is unchanged even when spoofer is on. The three WAAS signals PRN 133, 135 and 138 are all not spoofed in the scenarios and these are used to determine the absolute heading of antenna. The heading can then be used to calculate absolute azimuth of other PRNs. This is only for illustration purposes. In reality, the spoofing detector can be done by relative azimuth.

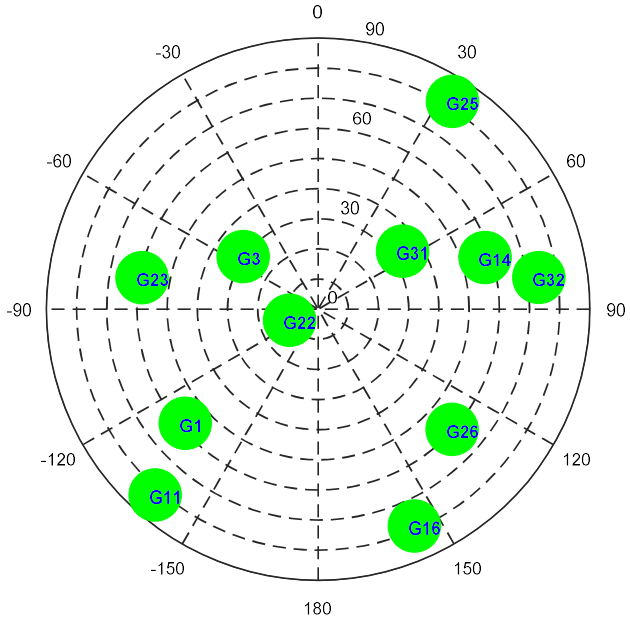

Figure 15. Sky plot based on broadcast ephemeris for GPS satellites during spoofing (Scenario 1)

Figure 15 shows the sky plot for GPS satellites with azimuths and elevations derived from the spoofed broadcast ephemeris. The spoofer broadcasts the same ephemeris from all available GPS satellite and only changes the pseudorange measurement in the receiver. Thus, we are using the true broadcast ephemeris but through the spoofing transmission.

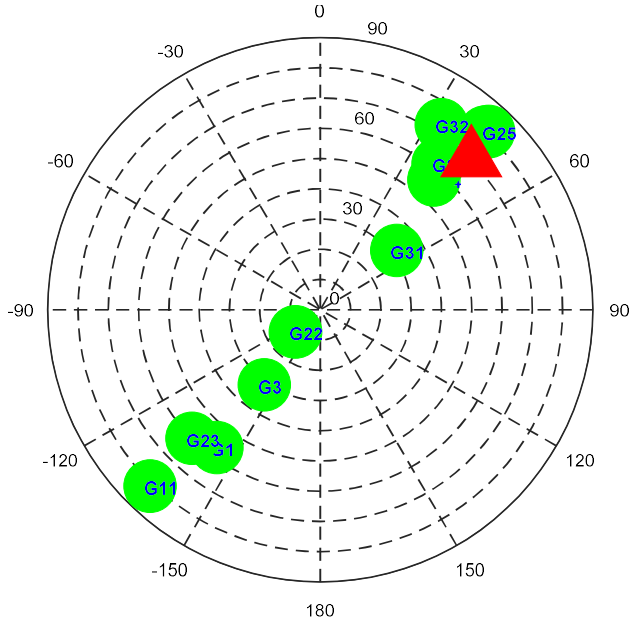

Figure 16. Sky plot for GPS satellites with estimated azimuth from the dual polarization antenna and elevation from ephemeris during spoofing transmission (Scenario 1)

Figure 16 shows the estimated sky plot for GPS satellites from the dual polarization antenna. The azimuth is estimated using null location in C/No (along with the absolute reference from WAAS) and elevation is derived by broadcast ephemeris. While DPA null steering only finds a relative DOA, it is still useful. Even without having an absolute azimuth reference (as derived from WAAS), the DPA still would find that all the satellites are coming from around the same direction which should indicate something is amiss. Our dual polarization antenna has DOA ambiguity of 180 degrees in its null azimuth estimation - see Equation (3) or Table 1. For illustration purpose, the azimuth from Figure 15 is used to determine whether the azimuth on the right or left hand side of sky plot as due to the 180° ambiguity. The red triangle marks the estimated direction of broadcasting

tower. All PRNs are aligned to a straight line from 45 degree and -135 degree which is quite different from the sky plot in Figure 15. This shows how DOA from multiple satellites can be used for spoofing detection.

Of course, the technique is agnostic to what type of spoofing is used. It will work regardless of whether the spoof is of time or position or both. Indeed, it will work even if the spoof signal provides the true time and location. An example of the performance under time spoofing is discussed next.

The effect of the spoofing on position is shown in Figure 17 which plots the positioning error over time (39 minutes) starting for assessed Scenario 2. The scenario is a "time push", like the one described in texbat scenario 2 and 3 [11], where time is gradually moved from the nominal time, starting from the 8th minute in the figure. The scenario ends at the 38th minute in the figure. The positioning results show the difference in position relative to first fix. The first fix was made prior to spoofing and is offset from the reference spoof location. The spoofing does cause a small position error as the reference location used for the spoofing is 10 m north and 10 m east of our true location. A user exactly at the reference location would only experience time spoofing. There are larger variations in height after the spoofer turned is on. This is due to the application of null-steering which results in all the spoofed PRNs having nulls at the same time. This results in 256 second periodicity in the error pattern. For example, the height error has peaks in the 6th, 11st, 15th, 19th, 23rd minutes.

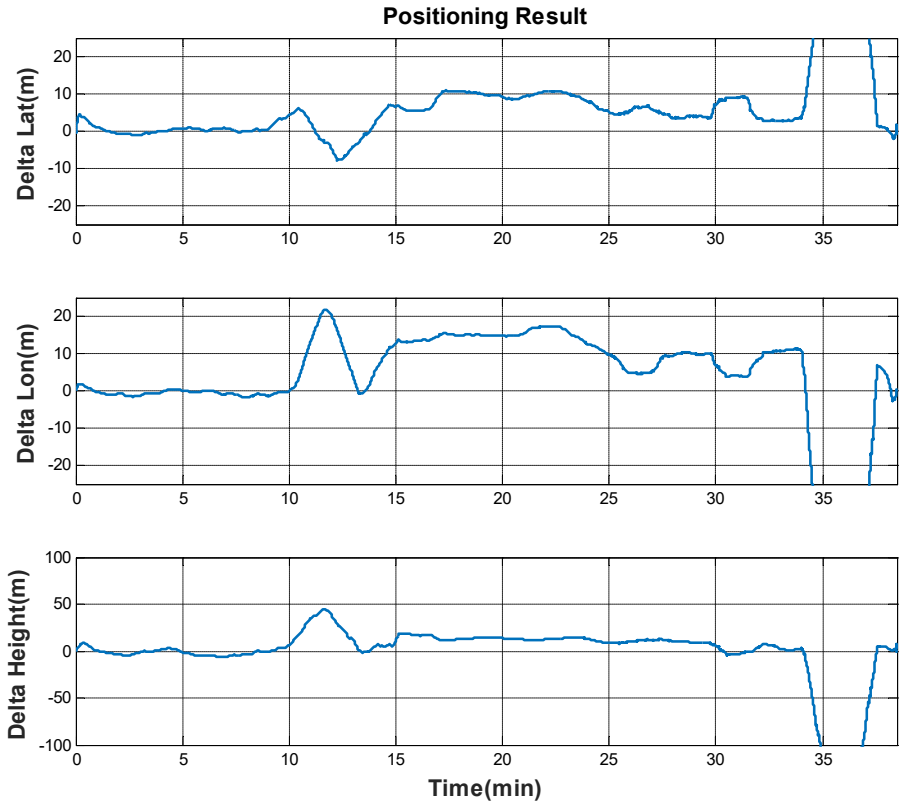

Figure 17. Positioning error with Dual Polarization Antenna during Spoofing for Scenario 2

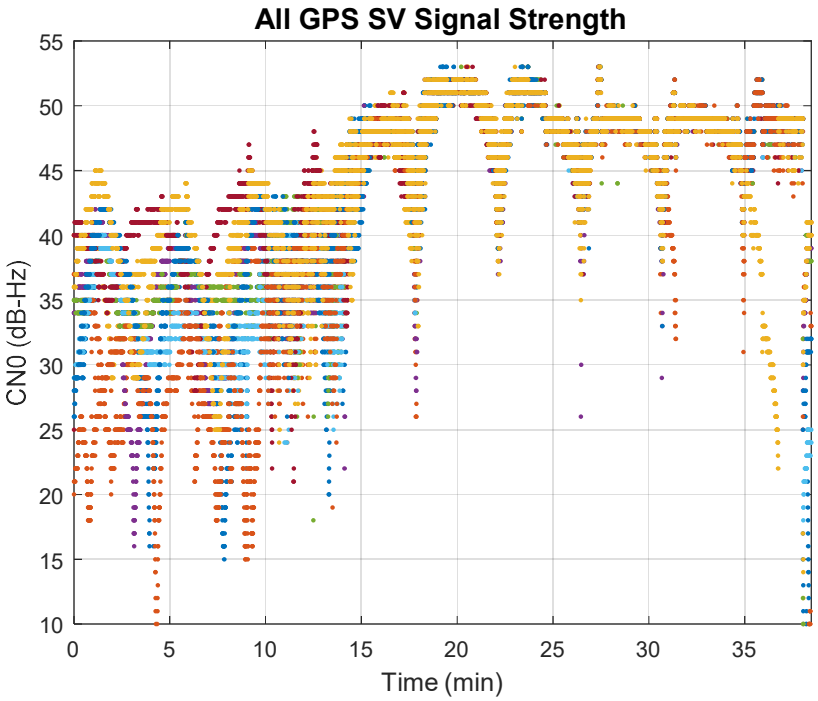

Figure 18. C/No for all GPS PRNs during Spoofing (starting at minute 6) for Scenario 2

Figure 18 shows the C/No for all GPS PRNs. In the first 5 minutes when spoofer is off, each genuine signal has a null with different depth and location in different time compared to the other satellites. After spoofer is turned on in the 8th minute, all the C/No align together by the 15th minute with similar nulls in term of depth and location.

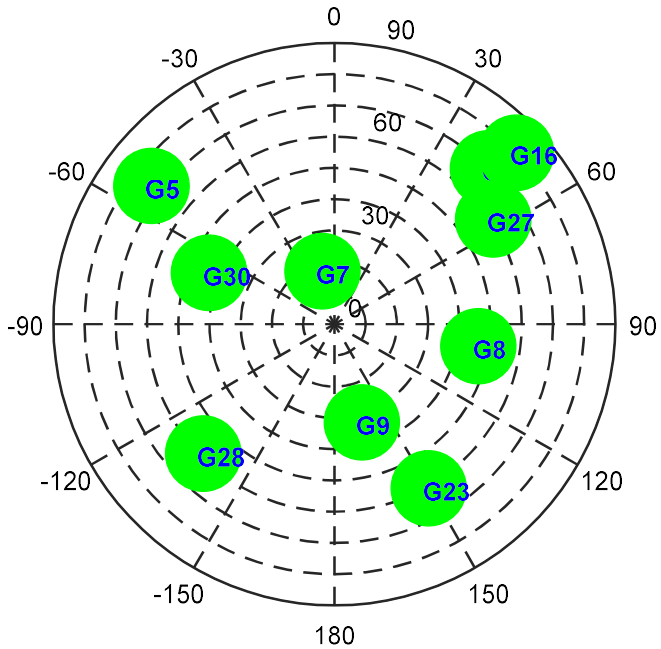

Figure 19. Sky plot based on broadcast ephemeris for GPS satellites during spoofing (Scenario 2)

Figure 19 shows the sky plot for GPS satellites with azimuth and elevation is derived by the broadcast ephemeris. This is derived from the spoofing rather than genuine signal. Again, the spoofer transmits the same ephemeris as that available from each of the genuine GPS satellite so either source provides the same information. The spoofer only affects the pseudorange measurement in the receiver.

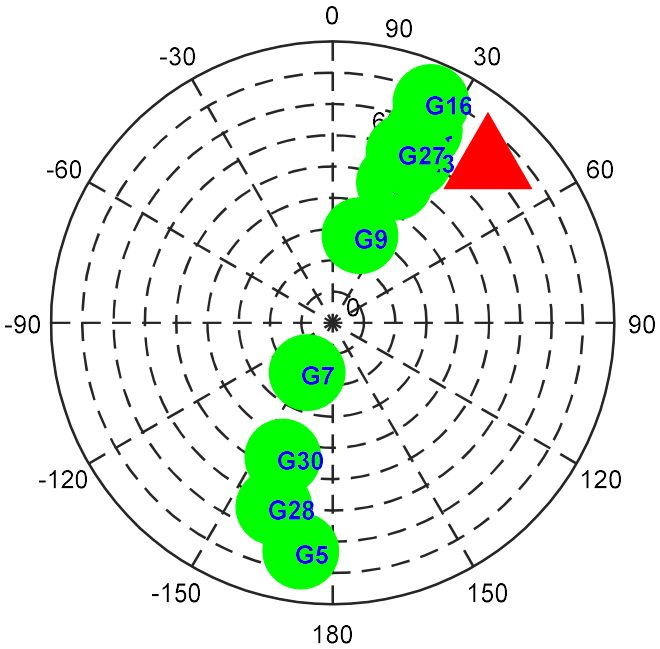

Figure 20. Sky plot for GPS satellites with estimated azimuth from the dual polarization antenna and elevation from ephemeris during spoofing transmission (Scenario 2)

Figure 20 shows the estimated sky plot for GPS satellites from the dual polarization antenna. The azimuth is estimated by null location in C/No and elevation is derived from the broadcast ephemeris (same as Figure 19). Due to our 180 degree azimuth ambiguity, for illustration purposes, we use the azimuth from Figure 19 determine the azimuth on the right or left hand side of sky plot. Again, the red triangle indicates the direction of broadcasting tower and its rough elevation. The alignment of the PRNs, though slightly offset from the true spoof direction, is a strong indication that we are receiving a spoofing signal.

Interference Test Results

As discussed in [10], the dual polarization antenna can provide mitigation against interference. However, given its DOA capabilities, can the DPA also help locate the interference source? The process of interference mitigation is similar to the phase offset search discussed for DOA. But for the DPA direction finding, we needed to track the signal. If the interference suppression allows for GNSS signals to poke through and be tracked, then we can find the direction of the interference source. As before, the microcontroller steps through each phase offset to create the combined signal. This will result in a null in the direction given by Equation (3). If that is the direction of the interference signal, then the interference signal will be suppressed, reducing noise. The suppression will perhaps allow the receiver to acquire and track GNSS signals that are also not greatly suppressed (i.e. not in the direction of the null) or even enhanced. Signals about 90° offset in azimuth from the direction of the null are in the gain area of the signal combination.

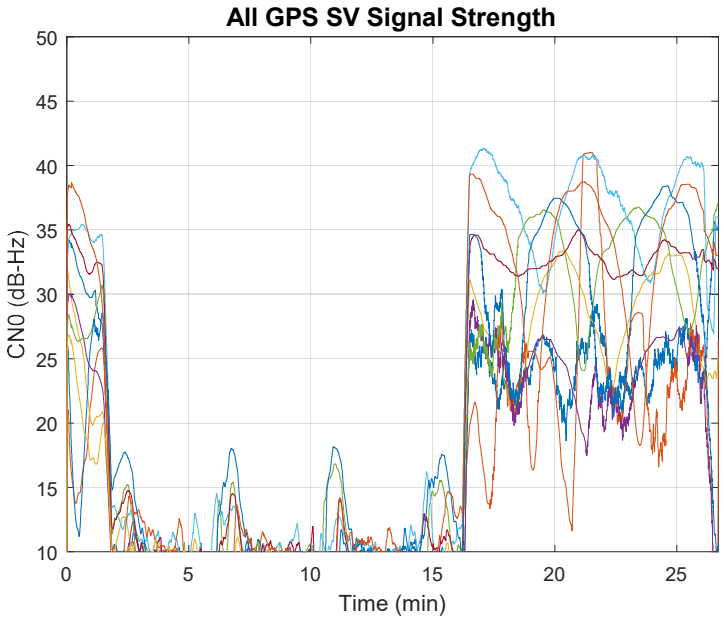

Figure 21. C/No for all GPS PRNs during jamming (starting at minute 1 and ending at minute 16) for Scenario 3

Figure 21 shows the C/No for all GPS satellites or pseudorandom numbers (PRNs) during Scenario 3. This scenario uses jamming and then spoofing as discussed in [6]. It is similar to what was believed to have happened in the Black Sea spoofing incident as the true GPS position was lost for a while before a new spoofed position is seen [12][13]. The jammer is on from about minute 1 to minute 16 and the resulting C/No of all GPS PRNs are lower by 20 dB-Hz or more. However, there are four peaks in this period when the null of gain pattern is steered towards the jammer. The effect of null suppresses the signal power of jammer, so the signal to noise ratio (SNR) of GPS satellites generally increases. Note that both the signal strength (S) and the noise (N) are a function of the steering of the null azimuth/phase offset. So, from Equation (4), if the null steering does not also significantly attenuate the GPS signal, then it may be possible to greatly increase SNR and possibly acquire the satellite. Four peaks correspond to four cycles in 15 minutes.

$$SNR = S(\psi) - N(\psi) = S\big(f(\varphi)\big) - N\big(f(\varphi)\big) \qquad (4)$$

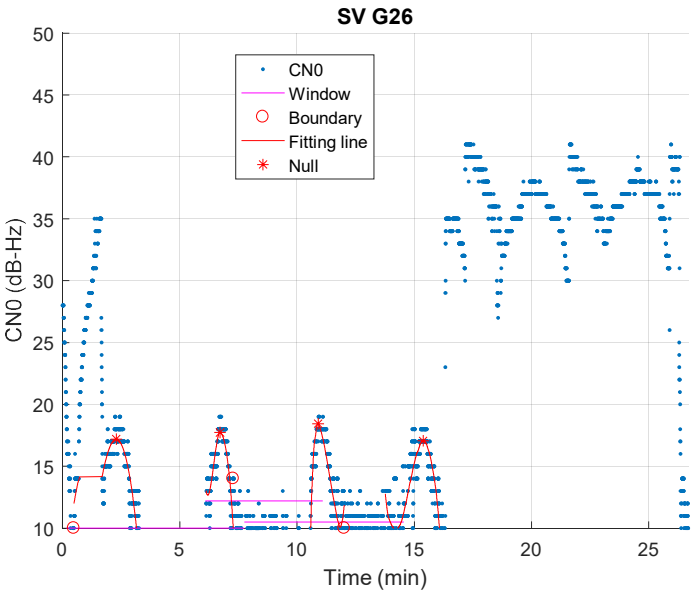

Figure 22. C/No for GPS PRN 26 during jamming (starting at minute 1 and ending at minute 16) for Scenario 3

Figure 22 shows the C/No of GPS satellite PRN 26, which is jammed from minute 1 to minute 16, over test time interval seen before. It also shows the curve fit used to form an estimate of the azimuth. Instead of using the null as it is very wide, the peak is used for azimuth determination. PRN 26 is used because it has highest signal-to-noise ratio during the jamming period.

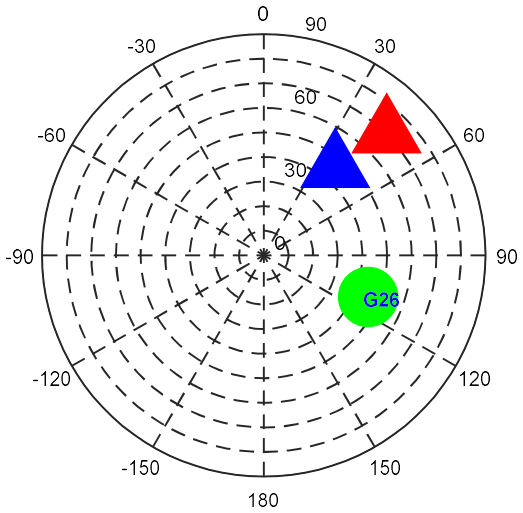

Figure 23. Estimated sky plot (DPA for azimuth, ephemeris for elevation) for GPS satellites from the dual polarization antenna during jamming test (Scenario 3). The red and blue triangles indicate the actual and estimated azimuth (DPA) of the jammer, respectively. The yellow G26 indicates the actual azimuth and elevation of PRN 26 (from ephemeris). Estimated jammer elevation is not determined and a value is chosen to visually separate it from the true jammer location

Figure 23 presents a sky plot showing the actual (red) and estimated (blue) azimuth of the jammer derived from *a priori* position measurements and the dual polarization antenna, respectively. Also shown is the location of GPS PRN 26 from its ephemeris. The estimated azimuth of the jammer is derived from the peak location of C/No as seen in Figure 22. The jammer elevation is not estimated and the plotted elevation is chosen to visually separate it from the true jammer location. This only provides a relative angle measurement so the direction of WAAS satellites prior to jamming was used as a reference. Notice that GPS PRN 26 is aligned approximately 90 degrees from the direction of jammer. Hence it is outside the null and likely in the high gain portion of the combined signal. These results show how DOA from the strongest satellite can be used for jammer direction finding or localization.

CHALLENGES & NEXT STEPS

There are still many challenges to overcome to make the DPA an effective and useful spoof detection method. The 256 second scan period is a limitation. We used 256 seconds in the field test to have high fidelity and to assure good direction determination. A fraction (half, a quarter or even less) of the dwell time should be adequate, especially for strong spoofing signals. However, a quarter of the dwell time still implies 64 second scans. We can also increase the steps between scan angle trading off some fidelity for speed. The scan time can also be reduced with architecture changes. For example, if we can create parallel processing chains that scan each quadrant. We could generate four different combined signals simultaneously if we split the RHCP signal through separate four phase shifters and process them through four different monitor receivers. This would further reduce the scan time by a factor of four. Overall, it seems reasonable to reduce the scan time by a factor of 16 to 64 using all of these techniques together. A 4-16 second scan time seems much more reasonable. The solution is reasonable for applications such as aviation as the components used are generally low cost, relative to avionics. This is shown in Figure 24. For automotive applications where more intense processing may be available, software processing can allow for parallel searches over many phases with the same digitized RHCP and LHCP signal. The tradeoff here is to have more powerful and power-consuming processing in place of dedicated hardware such as COTS receivers. We need to work on the optimal selection of these values to balance detection fidelity with speed.

Another area of work is to develop robust spoof detection algorithms. This paper show that we can generate good indicators of spoofing, there is the further step of using this information to robustly detect spoofing. That means developing algorithms that

use this information to find spoofing with low probability of false alerts and missed detection. The algorithm can combine both the elevation information derived from the null depth with the direction of arrival estimate. Another desirable quality is that the performance of the algorithm should be easy to analyze. We are currently developing such analyses and algorithms.

Another challenge is to operating the monitor in aviation environment. A major concern for external equipment on a commercial aircraft is susceptibility to electrical discharge due to precipitation static (p-static). These high energy discharges can destroy electronics. While much of the hardware on our DPA are analog components that should handle p-static, the monitor receiver is not. It can be placed inside the aircraft but that would necessitate an additional cable which is not desirable as it means placing another hole in and running another cable through the aircraft body. It also means longer installation times reducing revenue generation for the aircraft. The DPA would then not be just another replacement antenna. Hence, another next step is to figure out how the DPA can work in the presence of p-static.

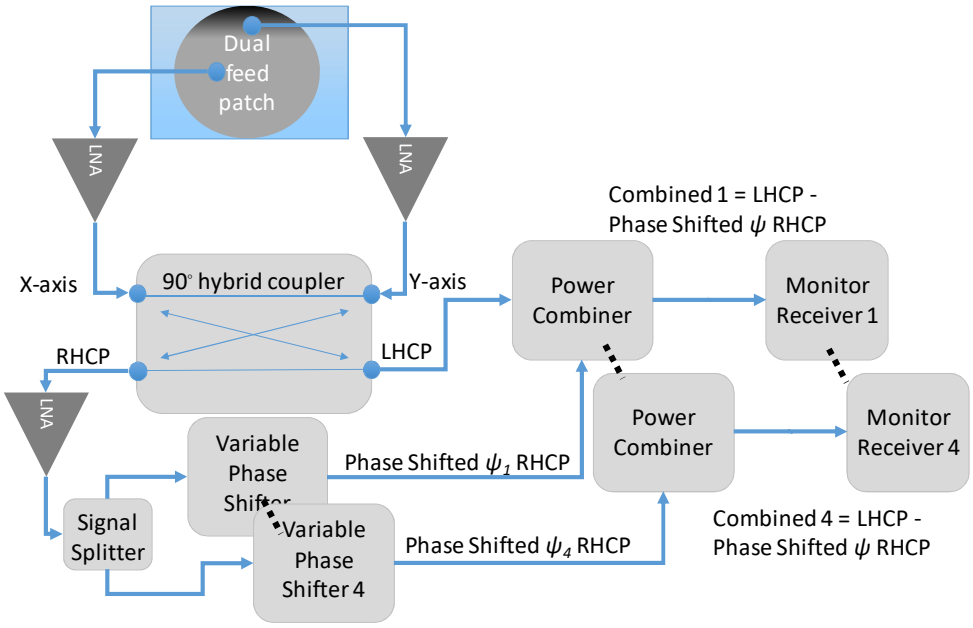

Figure 24. DPA Architecture for Faster Spoof Detection using additional hardware

CONCLUSIONS

GNSS spoof detection is becoming increasing important due to the significant of GNSS in safety and economically critical applications. The dual polarization antenna can be a potentially powerful and effective part of a spoof detection solution. Its design, as demonstrated by the Stanford prototype, has the form factor of a GNSS patch antenna making it more suitable size-wise than array antenna technologies. It is also ITAR compliant and can provide useful benefits including interference mitigation and even static heading using direction of arrival.

This paper examines the ability of the DPA concept to determine direction of arrival, particularly for GNSS spoof detection. It discusses the hardware and software processing aspects of the DPA-based direction of arrival determination. Then we demonstrate its performance in field tests with GNSS signals and with spoofing/jamming scenarios. It is able to find the DOA of satellites, spoofers and even jammers under the right conditions.

It is important to remember that the DPA should not be a stand-alone solution. Complementary spoof detection techniques such as those using redundancy checking, C/No and automatic gain control (AGC) are relatively easy to implement and the combination can make the GNSS receiver very difficult to attack. This is because each detection method examines on a specific and different signature left by a spoofing attack. Different spoofing attacks leave different types and levels of residual signatures. So having multiple techniques can help more comprehensively detect and more positively confirm the presence of spoofing. Multiple forms of detection is beneficial as that they can be layered to build an overall detection system that is sensitive to many attacks while not triggering during normal conditions. The DPA can form an important part of this overall system.

ACKNOWLEDGMENTS

The authors thank the Federal Aviation Administration (FAA) and the Stanford Center for Position Navigation and Time (SCPNT) for sponsoring this research

REFERENCES

[1] L. Huang, Q. Yang, “GPS Spoofing, Low cost GPS Simulator,” DEFCON 23, August 2015

[2] C. Sebastian, “Getting lost near the Kremlin? Russia could be 'GPS spoofing',” CNN Tech, Dec 2 2016, http://money.cnn.com/2016/12/02/technology/kremlin-gps-signals/index.html

[3] D. Goward, “Mass GPS Spoofing Attack in Black Sea?”, The Maritime Executive, July 11 2017, http://maritime-executive.com/editorials/mass-gps-spoofing-attack-in-black-sea

[4] Mark L. Psiaki, "Developing Defenses Against Jamming & Spoofing of Civilian GNSS Receivers," Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2014), Tampa, Florida, September 2014, pp. 2005-2012.

[5] M. L. Psiaki and T. E. Humphreys, "GNSS Spoofing and Detection," Proceedings of the IEEE, 2016.

[6] C. Günther, "A Survey of Spoofing and Counter-Measures," NAVIGATION, Journal of The Institute of Navigation, vol. 61, no. 3, Fall 2014, pp. 159-177, 2014.

[7] Y.-H. Chen, F. Rothmaier, D. Akos, S. Lo, P. Enge, "Towards a Practical Single Element Null Steering Antenna," Proceedings of the Institute of Navigation International Technical Meeting, Monterrey, CA, January 2017

[8] 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

[9] Yu Hsuan Chen, Sherman Lo, Per Enge, David Whelan, Iridium Authentication," Proceedings of the Institute of Navigation Joint Navigation Conference, Orlando, FL, June 2015 (Presentation Only)

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

[11] T.E. Humphreys, J.A. Bhatti, D.P. Shepard, and K.D. Wesson, "The Texas Spoofing Test Battery: Toward a Standard for Evaluating GPS Signal Authentication Techniques," Proceedings of the Institute of Navigation GNSS, Nashville, TN, 2012.

[12] Michael Jones, “Spoofing in the Black Sea: What really happened?” GPSWorld, October 11, 2017. http://gpsworld.com/spoofing-in-the-black-sea-what-really-happened/

[13] Henrik Lied, “GPS freaking out? Maybe you’re too close to Putin,” NRK beta, September 18 2017, https://nrkbeta.no/2017/09/18/gps-freaking-out-maybe-youre-too-close-to-putin/

[14] Esteban Garbin Manfredini, Dennis M. Akos, Yu-Hsuan Chen, Sherman Lo, Todd Walter, and Per Enge, “Effective GPS Spoofing Detection Utilizing Metrics from Commercial Receivers,” Proceedings of the Institute of Navigation International Technical Meeting, Reston, VA January 2018

[15] J. Gross, T. E. Humphreys, “GNSS Spoofing, Jamming, and Multipath Interference Classification using a Maximum-Likelihood Multi-Tap Multipath Estimator,” Proceedings of the Institute of Navigation International Technical Meeting, Monterey, CA January 2017.

[16] D. M. Akos, "Who's Afraid of the Spoofer? GPS/GNSS Spoofing Detection via Automatic Gain Control (AGC)", NAVIGATION, Journal of The Institute of Navigation, Vol. 59, No. 4, Winter 2012, pp. 281-290.