Field Test Validation of Single-Element Antenna with Anti-Jam and Spoof Detection

Emily McMilin, Yu-Hsuan Chen, David S. De Lorenzo, Sherman Lo, Dennis Akos, Per Enge
Stanford University, USA

BIOGRAPHY

Emily McMilin is a PhD candidate under Professor Per Enge in the Stanford GPS Research Laboratory. She completed her Bachelor of Science at Stanford in Symbolic Systems and her Master's Degree in Electrical and Computer Engineering at the University of Victoria in British Columbia. Prior to returning to Stanford for her PhD in Electrical Engineering, Emily was an Antenna Engineer at Apple for 2.5 years, and is currently a contractual RF Engineer with Facebook's Connectivity Lab.

Yu-Hsuan Chen is a research associate at the Stanford GPS Laboratory. He received his Ph.D. in electrical engineering from National Cheng Kung University, Taiwan. His research interests include real-time GNSS software receiver, antenna array processing and APNT.

David S. De Lorenzo is a Principal Research Engineer at Polaris Wireless and a consulting Research Associate to the Stanford GPS Laboratory. His current research is in adaptive signal processing, software-defined radios, and navigation system security and integrity. He received the Ph.D. degree in Aeronautics and Astronautics from Stanford University and previously has worked for Lockheed Martin and for the Intel Corporation.

Sherman Lo is currently a senior research engineer at the Stanford GPS Laboratory. He is the Associate Investigator for the Stanford University efforts on the FAA evaluation of alternative position navigation and timing (APNT) systems for aviation. He received the Ph.D. in Aeronautics and Astronautics from Stanford University.

Dennis Akos completed the Ph.D. degree in Electrical Engineering at Ohio University within the Avionics Engineering Center. He has since served as a faculty member with Lule Technical University, Sweden, and then as a researcher with the GPS Laboratory at Stanford University. Currently he is a faculty member with the Aerospace Engineering Sciences Department at the University of Colorado, Boulder and maintains visiting appointments at Stanford and Lule Technical University.

Per Enge is a Professor of Aeronautics and Astronautics at Stanford University, where he is the Kleiner-Perkins Professor in the School of Engineering. He directs the GPS Research Laboratory, which develops satellite navigation systems based on the Global Positioning System (GPS). He has been involved in the development of WAAS and LAAS for the FAA. He has received the Kepler, Thurlow and Burka Awards from the ION for his work. He is also a member of the National Academy of Engineering and a Fellow of the IEEE and the ION. He received his Ph.D. from the University of Illinois in 1983.

INTRODUCTION

As the utility of GPS continues to pervade (and upgrade) our daily activities, our increased reliance makes us more vulnerable to the inherent weaknesses of GPS. The incredible faintness and unencrypted nature of the GPS signal exposes it to both unintentional and intentional overwhelming, via mechanisms such as interference, jamming and spoofing (the broadcast of counterfeit GPS signals intended to deceive a GPS receiver).

However some may argue that despite the great risk, the probability of a jamming or spoofing event is thankfully low. The argument may continue, that correspondingly few resources should be dedicated to protect against a low probability event. We attempt to resolve the tension between these high risk, yet low probability scenarios by establishing antenna designs that provide protection while exploiting existing infrastructure and equipment, thus requiring minimal additional resource dedication.

In our previous work, we introduced a backward compatible single antenna design for GPS spoof detection [4] and

anti-jam [5] for aviation applications. The proposed design required no additional signal processing blocks when in use with a standard GPS receiver and fit into the form-factor of a standard GPS antenna. In [4], we tested the proposed technique by combining simulated data with measured data to enable hardware-in-the-loop experiments, and in [5] we based our conclusions upon mathematical theory and simulation alone.

This paper reports the results of a recent field trial where we were able to run hardware-in-the-loop tests to evaluate the performance of both our anti-jam and spoof detection techniques in response to jamming and spoofing attacks. In this paper we intend to validate our prior claims of achieving greater than 10 dB jam suppression and reliable spoof-detection when these threatening signals originate from below the horizon of the GPS antenna, making aerial platforms an ideal application.

PRIOR WORK

As discussed in [5], most physical layer jam suppression is achieved by multiple antennas, connected to multiple radio front-ends and digitizers [6]. Although efforts to miniaturize the size of these array systems have proved successful while still providing impressive jam suppression [7], none yet have achieved form-factor compliance with the aviation ARINC 73 antenna dimension standards of 4.7 inch by 3 inch surface area by 0.73 inch height (11.938 mm x 7.62 mm x 1.854 mm). Furthermore, most multi-antenna arrays require additional receiver hardware, calibration and computational complexity.

A single antenna design has been recently introduced [8] that can achieve robust jamming mitigation (as well as spoof detection). However this design currently suffers a constant loss of $C/N_0 \approx 6$ dB. Furthermore, this technique requires a MIMO (Multiple Input and Multiple Output) receiver that must undergo periodic self-calibration to maintain phase coherency between the two radio frequency (RF) paths that connect the single antenna to the two receiver inputs. The work in [8] builds upon prior work introducing an interference suppression unit (ISU) [9], that when placed between a single GPS antenna and the receiver can provide impressive jam suppression. However, the idea published in 1998, was never further developed in later publications, perhaps due to implementation complexity. Nonetheless, the developments achieved in [8] and [9] are very promising and we draw from the polarization mismatch technique described in these papers as motivation for our work here.

Theoretical Background

We all know that electromagnetic waves can propagate through both free space, such as the space between the GPS satellites and our antenna, and along conductive structures, such as the coaxial cables that deliver the electromagnetic wave from the antenna to our receivers. However, less obvious is that certain mediums and geometries only support types of electromagnetic fields. The waves that travel from the satellites to our antenna take the form of transverse electromagnetic plane waves. In the case of GNSS, the electromagnetic plane waves are right hand circularly polarized (RHCP). An RHCP wave can be decomposed two orthogonal electric field components (which we can call an x-axis field and a y-axis field for some arbitrary coordinate system in the plane parallel to the plane wave). These two field components are not only orthogonal is space, but also in time, with the x-axis field lagging the y-axis field by 90°.

When an RHCP wave is directly incident upon an RHCP antenna, the two orthogonal electric field components will excite both feeds on the antenna, with a portion of the energy lagging by 90° in time. Contrarily, when an RHCP wave is directly incident upon the conductive ground plane, the magnetic component of the electromagnetic wave will induce surface currents along the ground plane. Despite the considerable losses endured in this transmission mechanism, some of these surface currents will travel along the body of the ground plane until they reach the antenna, at which time they will induce a potential difference between the ground plane and the patch antenna. This potential difference causes an electric field, similar to that produced by the RHCP wave directly incident at the antenna. However, in the case of the surface currents, there will be only one temporal component to the electrical field. Put another way: no component of the energy will be lagging another component by 90°. Thus, although we may select a coordinate systems that decomposes the surface current's electric field into two spatially orthogonal components, there will only be one temporal component to the surface current's electrical field. For this reason, the electric field induced by a surface current is similar to that induced by a vertically polarized (VP) electromagnetic plane wave, and thus we refer to these signals VP. This transformation of an incident waveform into a VP signal is true for any arbitrarily polarized incident energy [5].

One significance of a vertically polarized field in GNSS, is that we can be quite confident that it did not originate directly from a GNSS satellite (note that some low elevation GNSS satellites waveforms can appear largely VP to a patch antenna, as will be addressed later). Specifically, for an antenna atop a large cylindrical ground plane (such as an aircraft), any signals that reach the antenna due to the propagation of surface currents will presumably do so be-

cause a direct path to the antenna is blocked by the ground plane, and thus must have originated from beneath the horizon of the antenna. Thus, for the remainder of this paper, we will assume VP fields are due to waveforms that only originate from elevation angles below the horizon of the antenna. Waveforms that originate from satellites above the horizon will include at least some RHCP energy, and thus will not be VP.

Another significance of a VP signal is that it can be linearly decomposed into an RHCP signal and a left hand circularly polarized (LHCP) signal, with both signals having equal magnitude and phase coherency. In fact these two signals are exactly what we see at the two output ports of a 90° hybrid coupler, when presented with a vertically polarized signal at its input ports [5].

Basic Design

Fig. 1 shows a high-level schematic diagram of the anti-jam and spoof detection technique tested in this paper. As indicated in the figure, all the components are housed inside the antenna assembly, with only a single coaxial cable connecting the antenna assembly to the GPS receiver. We can see that 90° hybrid coupler has two output ports, labeled as "RHCP" and "LHCP". Note that even genuine signals from high elevation GNSS satellites will deposit a small amount of LHCP energy at port 4. This LHCP energy can derive from many sources, such as atmospheric effects and antenna imperfections. Additionally, GNSS multi-path components will have changed from RHCP to LHCP after a single bounce.

To mitigate these and other negative effects, the LHCP energy is generally deposited directly into a 50Ω load, as shown in green switch state displayed in Fig. 1. When the switches are in this state, the block diagram represents a standard GNSS antenna. Furthermore this state will serve as the model for the "normal mode" signals later plotted in the paper. We will later detail how the other switch states in the antenna implement anti-jam and spoof detection.

EXPERIMENTAL SET-UP

Our prior work relied on signals generated in software and simulation as a proxy for actual GNSS spoofer and jamming signals. In this paper we will only present results from actual signals captured and recorded in field trial experiments. This next section will detail both our hardware set-up in the field and our field test methodology.

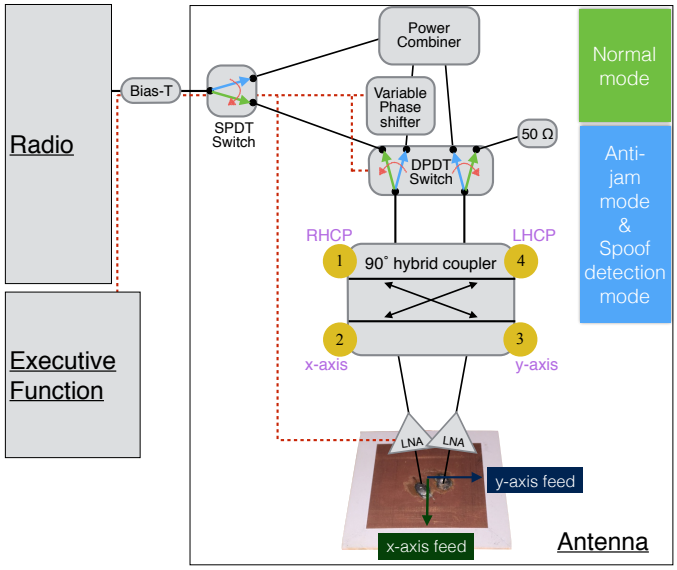

Fig. 1: A high-level schematic diagram of a normal GPS antenna, as well as the anti-jam and spoof detection technique tested in this paper. When the switches are in the green state, the block diagram represents a standard GNSS antenna, and in this state will serve as the model for the "normal mode" signals later plotted in the paper.

Hardware under test

As we note above, both the anti-jam and the spoof-detection mechanisms are designed to be compatible with standard off-the-shelf GPS receivers. Thus, only one cable, namely a coaxial cable should connect the antenna to the receiver. However, referring to Fig. 2, we can see two RF cables connecting the antenna to the Universal Software Radio Peripheral (USRP) B210, which is the software defined radio hardware platform [2] we used to capture raw signals during our field trial. This arrangement facilitated post-processing of the raw I/Q samples captured by the B210 for later playback to a GPS software defined receiver (SDR). In this paper, we utilized the Stanford GPS SDR [1] to serve as our reference GPS receiver, which will be detailed in later sections. Comparing Fig. 2 to Fig. 1 we can see that the similarity between the two diagrams ends immediately after the LNAs. Specifically, ports 2 and 3 of the 90° hybrid coupler are replaced by the input ports of the B210. Thus, in this paper we will use software to emulate all the components after ports 2 and 3.

This section and later sections will detail how the components in Fig. 1 are represented in software, such that when we "playback" the processed signals for a software defined GNSS receiver, the C/N_0 results represent those we expect to see in an actual hardware implementation. We will highlight any areas where our software implementation may deviate from expected results given by commodity hardware.

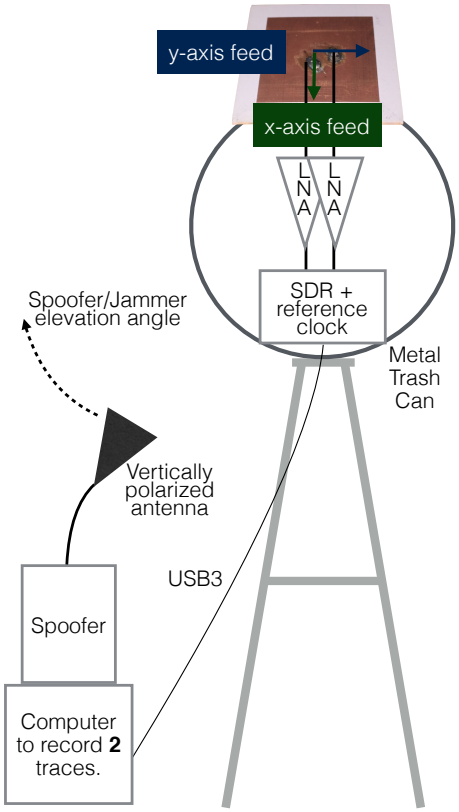

Fig. 2: Block diagram of our hardware setup in the field trial. Note the antenna on top with two RF cables passing through LNAs and then connecting to the Universal Software Radio Peripheral (USRP) B210, which is also connected to an external clock and a laptop for storing the recorded signals. Comparing to Fig. 1 we can see the that the similarity between the two diagrams ends immediately after the LNAs. Specifically, ports 2 and 3 of the 90° hybrid coupler are replaced by the input ports of the B210.

For this field trial we designed and fabricated a simple GPS patch antenna that we expect to consistent with existing patch antennas, and that meets ARINC 743 form-factor constraints. The antenna, shown in Fig. 3, is a 40mm by 40mm substrate with a 30mm by 30mm square copper patch on top. The substrate, at 1.28 mm thick, is a single layer of Rogers RO3010 material with a dielectric constant of 10.2. The high dielectric constant permits a relatively small form-factor antenna half wavelength resonant antenna in the *medium* at 1.575 GHz. Specifically, the wavelength in the medium, $\lambda_m = \frac{\lambda_{fs}}{\sqrt{\epsilon_r}} \approx 60$mm, where λ_{fs} is the GPS L1 free space wavelength of about 190mm and ϵ_r is the dielectric constant of 10.2. A single layer of substrate was preferred for ease of prototyping, however a slightly thicker substrate will generally improve antenna efficiency [11]. The patch antenna has two coaxial feeds, separated spatially by 90 degrees. We select a coordinate system such that we call one of the feeds the *x-axis feed* and the other the *y-axis feed*.

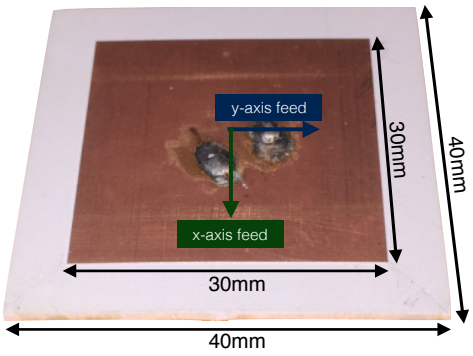

Fig. 3: The antenna we designed and prototyped for this field trial is a 40mm by 40mm Rogers RO3010 material substrate with a 30mm by 30mm square copper patch on top. The substrate, at 1.28 mm thick, is a single layer of the Rogers material, that has a dielectric constant of 10.2.

This antenna was designed to be mounted on a large conductive body, such as the fuselage of an airplane. In the case of this field trial the antenna was mounted on a metal trash can that served to emulate the cross-section of an airplane fuselage. Specifically, the antenna was affixed to a standard 31 gallon galvanized steel trash can with a 533 mm (21 inch) diameter (at its widest dimension) and a 685 mm (27 inch) length. Although we one day hope to replace the trashcan with an actual aircraft, for the time being, any relatively large ground plane will serve to validate our technique. Furthermore, we expect improved results as the ground plane increases in size [5].

The return loss for both feeds are shown in Fig. 4, showing greater than 10 dB return loss from about 1570 MHz to 1580 MHz. Because the antenna was tuned for a large ground plane, this measurement was done with the antenna mounted on the trash can. Note the out of band return loss of about 1.5 dB is due to the one-way insertion loss of 0.75 dB in the cables feeding from the antenna into the interior of the trash can.

Two coaxial cables connect the ports of the patch antenna with two parallel LNAs. We again note that a future integrated solution would include the LNAs and other discrete components within a form-factor compliant antenna assembly, such that only one coaxial cable will connect the antenna to the GNSS receiver. The LNAs used in this field trial provided about 13 dB of gain and have a 0.5 dB noise figure [3]. Scattering parameter measurements of each LNA's S21 show good magnitude and phase parity (within measurement error) to one another. Two more coaxial cables depart the LNAs and feed into the B210, which is a two port radio with two phase coherent signal paths. Signals on these two parallel paths were sampled at a rate of 5 Msps. Both receive channels have a common internal clock, however we have instead elected to use an external 10 MHz Rubidium oscillator for improved clock stability

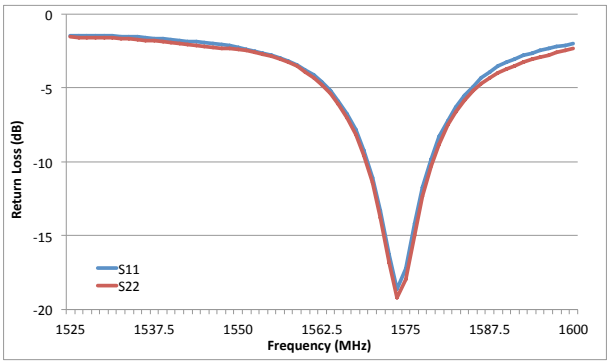

Fig. 4: The return loss for both feeds of our prototype antenna mounted on the trash can, showing greater than 10 dB return loss from about 1570 MHz to 1580 MHz, and an out of band return loss of about 1.5 dB due to the one-way cable insertion loss of 0.75 dB.

of less than 2 picoseconds over 100 second interval. Note that this clock synchronization between the two feeds will not be necessary in a hardware implementation of this design, as it will only have a single feed. Finally the B210 is connected via USB3 cable to a laptop to record the raw I/Q samples that are captured from each stream. We did not perform any real time processing on the recorded samples. Thus, during the field measurement we were blind to the performance of the PNT solution and the degradation that the jamming and/or spoofing may be inflicting upon the solution. In later sections we will revisit these two recorded streams of data and the post processing steps performed upon the raw data we recorded in the field.

Field Measurement

The hardware setup implemented in the field is shown in Fig. 5. Except for the addition of the DC power supply to the LNAs and the AC power supplies to the clock and B210, the image in Fig. 5 matches the block diagram shown in Fig. 2.

The field measurements were conducted during the Joint Interagency Field Experimentation (JIFX) program on 12 May 2015 at Camp Roberts, California, USA. A Google Maps image of the test site is shown in Fig. 6. Superimposed on the map's image is the approximate location and orientation of our test antenna, as well as a sky plot of the overhead GPS satellites during the test. The sky map includes a dotted red line to indicate the direction of the jamming signal (parallel to the y-axis feed of the antenna) and yellow line to indicate the axis orthogonal to it (parallel to the x-axis feed of the antenna). It will later be noteworthy that satellite 17 is almost directly overhead, satellite 28 is in the direction of spoofed signal, satellites 15 and 30 are approximately orthogonal to the direction of the spoofed signal, and satellites 6 and 7 are at a low elevation angle

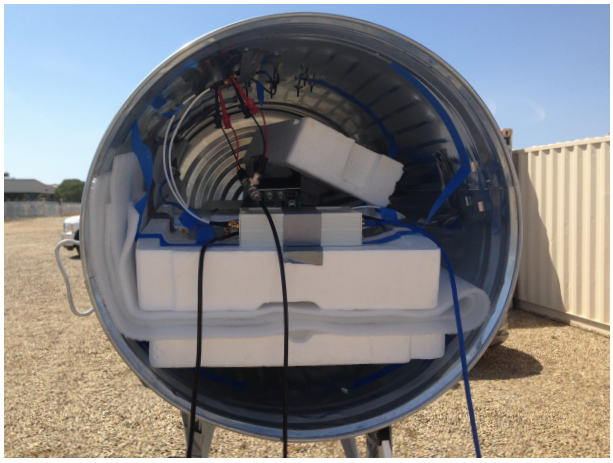

Fig. 5: The actual hardware setup implemented in the field. With the exception of the addition of the DC power supply to the LNAs and the AC power supplies to the clock and B210, this picture matches the block diagram shown in Fig. 2

(6° and 15° respectively).

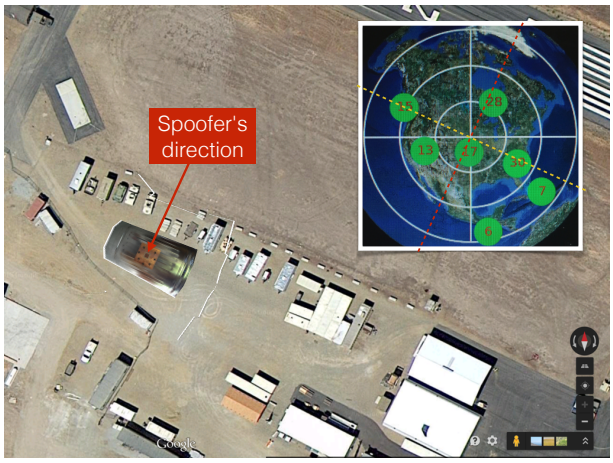

Fig. 6: A Google Maps image of the test site labeled with our approximate location and orientation of our test antenna relative to north and relative to the spoofing signal. Also included is a sky plot of the overhead GPS satellites during the test. The sky map includes a dotted red line to indicate the direction of the jamming signal and yellow line to indicate the axis orthogonal to it.

Members from the Joint Vulnerability Assessment Branch (JVAB) were on site to support our field trial. Fig. 7 shows the Spectracom GPS emulator used by JVAB to generate the spoofed GPS signal shown in Fig. 8. This signal was generated in band at a center frequency of 1.575 GHz with -65 dBm of power. To minimize impact upon other nearby experiments, JVAB spoofed for a position in China, thus generating signals for a largely non-overlapping set of GPS satellites and with different dopplers than the genuine overhead satellites. Specifically, the spoofed signals have about 300 Hz higher doppler than the genuine signals.

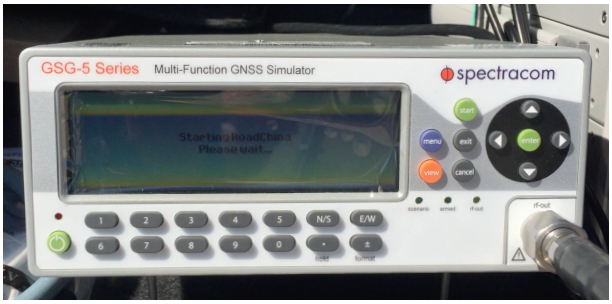

Fig. 7: Spectracom GPS emulator used by the Joint Vulnerability Assessment Branch (JVAB) to generate the spoofed GPS signal shown in Fig. 8

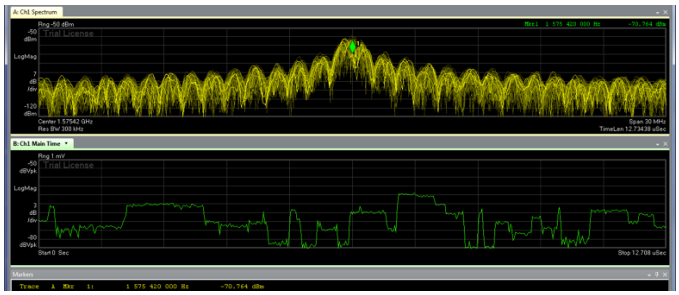

Fig. 8: The spoofed signal generated by JVAB, at a center frequency of 1.575 GHz with -65 dBm of power. To minimize impact upon other nearby experiments, JVAB spoofed for a position in China, thus generating signals for a largely non-overlapping set of GPS satellites with different dopplers than the genuine overhead satellites

The spoofed signal was fed to a 7 dBi vertically polarized horn antenna which can be seen in the left of Fig. 9. This figure also shows the last stage of our measurement process. Prior to conducting the measurement, we affixed the antenna plus trashcan assembly a top a small ladder, just under 2 m high. Referring to Table 1, we began with measurement of the overhead genuine GPS satellite signals, without any spoofing signals present. After about 120 seconds the spoofer signal was turned on. At about 145 seconds we placed the horn antenna into its first position near the ground and at about 1.2 m from the GPS antenna, with an elevation angle of about -75° relative to the horizon of the antenna (15° above nadir), and an azimuthal angle of 90° relative to the x-axis feed of the antenna. As indicated in the table, we held the spoofing horn static for about 30 sec to 40 sec time increments, before increasing the elevation angle of the horn by about 25 degrees (relative to the GPS antenna). After 290 seconds we completed the experiment. Regrettably, we did not vary the polarization of the spoofing signal, so we saved this for future work.

Fig. 9: David De Lorenzo holding the horn at the horizon of the GPS antenna, in the last stage of our measurement process, with Yu-Hsuan Chen capturing signals on the B210. The GPS antenna is affixed to the top of the metal trash can (which serves to emulate an airplane fuselage), which itself is mounted on top a ladder.

time (sec)	action (elevation angle below the horizon)
0	start recording GPS signals
120	start spoofing signal
145	hold antenna at -75°
185	hold antenna at -50°
215	hold antenna at -25°
250	hold antenna at 0°
290	stop spoofing signal

Table 1: The test procedure for capturing the genuine and spoofing/jamming signals in the field trial

In addition to this spoofing experiment, we conducted two other experiments with JVAB's support. In both experiments a carrier tone generated at 1.575 GHz was transmitted and the elevation angle of the horn was varied in a manner similar to the spoofing experiment explained above. In one of the jamming experiments we varied the azimuthal angle of the horn relative to the GPS antenna and in the other experiment we varied the power level of the jamming signal. Unfortunately, a spectral analysis of the signals captured during this experiments did not indicate the presence of the jamming signal and we saw no degradation in C/N_0 of the acquired satellite signals as we turned on and increased the jamming signal strength.

Fortunately, we were able to test both the anti-jam and the spoof-detection techniques with the signals recorded during the spoofing experiment. As mentioned previously, the spoofed signals were for a largely non-overlapping set of satellites with dopplers very distinct from the overhead satellites. Thus, the Stanford GPS SDR was able to acquire and attempt to maintain lock on the genuine overhead satel-

lites, while the spoofed signals from the non-genuine satellites appeared as a jamming signal. This is the method we used during playback of our recorded signals when post-processing in accordance with the anti-jamming technique. Contrarily, by sweeping doppler in our GPS SDR, we were able we locate and acquire the spoofed signals. The Stanford GPS SDR can track both the genuine and spoofed signals and attempt to maintain lock on both while the higher power spoofing signal is transmitted. This is the method we used during playback of our recorded signals when post-processing in accordance with the spoofing technique, as we will discuss more in the next section.

RESULTS AND ANALYSIS

We revisit Fig. 1 to see a block level diagram of the components that we implemented in software. Although, we would expect that the antenna will generally remain in "normal mode" or the *green* state, various events may trigger a "change of state", as we will discuss shortly.

The "90° hybrid coupler" block, "variable phase shifter" block, "Power combiner" block, and the switches were all implemented in C++ assuming idealized performance of each block component. In fact, in the actual hardware implementation we do expect signal impairments, primarily insertion loss to be higher in the anti-jam signal path than the normal mode signal path. However, the inclusion of the two parallel LNAs immediately after the antenna's feed ports can reduce the increases in noise figure to less than 1 dB[5]. By weighting and combining the raw I/Q values as both recorded streams were played-back into the Stanford GPS SDR, we implemented the 90° hybrid coupler, thus outputting the parallel RHCP and LHCP signals (as indicated in Fig. 1). This RHCP signal serves as the normal mode stream. In the case of the anti-jam and spoof detection mode streams, we inserted a variable phase shift value to the RHCP signal, and finally combined the phase shifted RHCP signal with the LHCP signal. Finally the resulting streams (such as the normal mode stream, and the anti-jam stream or the spoof-detection stream) are read into the Stanford GPS SDR to solve for position.

The Stanford GPS SDR can run up to 5 streams in parallel, tracking a total of 60 channels (or 12 channels for each stream). The software starts by reading data from disk and then stores the data into a 2 second-long queue. Every 1 millisecond data is processed in 5 working threads. Each thread serves the 12 channels for which is executed functions including the software correlator, signal acquisition/tracking and message decoding. These functions have the most computational complexity, so we distribute these channels to multiple threads to save processing time. Every 100 milliseconds, another thread takes measurement from all the tracked channels to solve for the receiver's position. For the signal tracking, the coherent integration time is 20 milliseconds. The C/N_0 output rate is 2.5 Hz (or computed every 400 milliseconds). This section will further detail how we used the Stanford GPS SDR to implement the anti-jam and spoof detection techniques, and analyze the results.

Anti-Jam

First we will assume that change of state to anti-jam mode has occurred. We can see that when the switches in Fig. 1 are in the *blue* anti-jam state, both port 1 and 4 of the 90° hybrid coupler end up meeting in the power combiner (such as a Wilkinson power combiner). However, the signal path coming from port 1 (labeled "RHCP") first passes through the variable phase shifter en-route to the power combiner. As indicated in the figure by the dotted red control lines, the optimal value of the phase insertion added by the variable phase shifter will be determined by the *executive function* block. Note that the behavior of the variable phase shifter is the primary implementation difference between anti-jam mode and spoof detection mode. To maintain backward compatibility with existing GPS receivers, the only reporting framework used will be the receiver's own C/N_0 information. Thus, during a jamming scenario, the figure of merit used to qualify an "optimal" phase shift could be determined by a minimization of the C/N_0 degradation suffered by the satellite signals that are being jammed.

In practice, the optimal value can be hand tuned, or more likely implemented with a power minimization algorithm steered by the automatic gain control in a standard receiver [5]. In the former scenario, a "change of state" from normal mode to anti-jam mode might be triggered by an operator seeing degradation in C/N_0 for some unknown reason. In the latter scenario, the "change of state" could be triggered by the AGC deviating from a pre-established nominal value to a lower gain value, in response to the increased "signal" level caused by the jammer. This scenario would likely require a firmware upgrade in the receiver (unless the receiver already reports AGC deviations via a serial output).

For this paper, we have the benefit of not needing to change states, but instead we can monitor both states in parallel. In an actual implementation, we would expect the C/N_0 result to jump from normal mode to anti-jam mode once the switch is flipped. For these field trial results we determined the optimal phase shift value by referring to the theory developed in our prior paper [5]. Specifically, given the geometry of our antenna and the direction of the jamming signal, the calculated optimal phase shift is -90°.

Using 1° step size, we swept a range of phase shift values around our calculated optimal value and determined qualitatively that a phase shift of -73° produces the most

optimal results, using the figure or merit discussed above. Fig. 10 compares the Stanford GPS SDR's results for the direct "Normal mode" stream (in green) and the "Anti-jam mode" stream (in red), when the variable phase shifter has been steered to a value of $-73°$. This variable phase shift value will steer a null in the direction of the jamming signal: at an azimuthal angle of approximately $25°$ east of north, or $90°$ from the x-axis feed of the GPS antenna.

Referring back to Table 1, we can see that the drop in the green normal mode C/N_0 is correlated with the increase in the elevation angle of the horn transmitting the jamming signal. When the jamming signal is incident upon the "fuselage" at a lower elevation angle, it must propagate along the ground plane for a longer distance before it reaches the antenna, and thus is further attenuated. However, as the jammer increases its elevation angle up to the horizon of the antenna, the *effective* signal strength of the jammer increases despite no change in the transmission power level. Consequently later in the signal recording we are more likely to see a loss of lock on satellite signals. This is particularly the case for the lower elevation satellites (6, 7, and 15) which already had a lower initial normal mode C/N_0 prior to the introduction of the jamming signal, and satellite 28 which is aligned with the direction of the spoofer.

Now turning to the red anti-jam C/N_0 traces, we see jam suppression ranging from about 10 dB to greater than 20 dB. Anti-jam performance for the high elevation satellites (13, 17 and 30, but excluding satellite 28) increases to around 10 dB of jam suppression. While jam suppression of the lower elevations satellites is generally 20 dB or better, avoiding a loss of lock for several satellites (when compared to the normal mode performance).

By coincidence, the spoofing signals are originating from the same direction as satellite 28, thus a radiation pattern null has been formed along a line in the azimuthal plane that is parallel to satellite 28 (indicated by the purple arrow on the sky plot in Fig. 10). Collaterally, a slight radiation pattern peak has been formed along the line orthogonal to the direction of the spoofer in the azimuthal plane (indicated by the black dotted line on the sky plot in Fig. 10). This black dotted line happens to run between satellites 15 and 7. Thus, in satellite 28 we see a slight reduction in anti-jam mode C/N_0 as compared to normal mode C/N_0 before the jamming signal has begun to degrade normal mode C/N_0 (and we also see more dramatic jam suppression as the effective signal strength of the jammer increases). Contrarily, in low elevation satellites 15 and 7, we see a slight increase in the anti-jam mode C/N_0 as compared to normal mode C/N_0 even before the jamming signal has begun. This superior performance of anti-jam mode C/N_0 as compared to normal mode C/N_0 continues for satellites 15 and 7 as the effective jamming signal strength increases, because of the compounded effects of the peak steered toward these two satellites and the null steered toward the jammer.

The deviation in the optimal phase shift from the calculated value of $-90°$ to the empirically measured value of by $-73°$, corresponds to an azimuthal angle offset of half the difference, or $\frac{1}{2}(-73° - -90°) = 8.5°$ in azimuth. This geometric deviation could easily be caused by our horn antenna placement being off by several degrees and by the dominant y-axis field in the prototype antenna being off by several more degrees.

Note that there are periodic dips (approximately every 18 seconds) in the C/N_0 traces for both the RHCP and the anti-jam traces. These dips appear to be caused by an unknown source of interference, as they appear throughout our experiment, prior to turning on the spoofing and jamming signals.

Spoof detection

We will now consider a change of state to spoof detection. Referring back to Fig. 1, and focusing again on when the switches are in the *blue* spoof detection state. As we mention above, the implementation difference between spoof-detection and anti-jam modes comes down to the behavior of the variable phase shifter. Also unlike anti-jam mode, the "change of state" into spoof detection mode will most likely be triggered according to some deterministic schedule and last for some predefined time period.

The reason for these distinctions between the implementation of anti-jam verse spoof detection is because we assume that the goal of the spoofer is to evade detection, and consequentially there would be no undetectable trigger. Rather, the antenna would preemptively scan for the presence of spoofers, following some predefined schedule.

For this field trial, our spoof detection implementation differs slightly from what we first introduced in [4]. Specifically, our later work on anti-jam in [5] calculated the phase dependency of the RHCP and LHCP signals vs azimuthal angle, motivating the inclusion of the variable phase shifter component. However, unlike our anti-jam implementation introduced in [5] and described in the subsection above, we can not rely on using a power minimization algorithm for the determination of the optimal phase value. This is again because the evasive and sophisticated spoofer would attempt to avoid triggering any detectable AGC response, as could be achieved by very cautiously just-slightly over-powering the genuine GPS signals at the receiver under attack.

Thus, similar to our discussion above, we are forced to deterministically steer the variable phase shifter through all

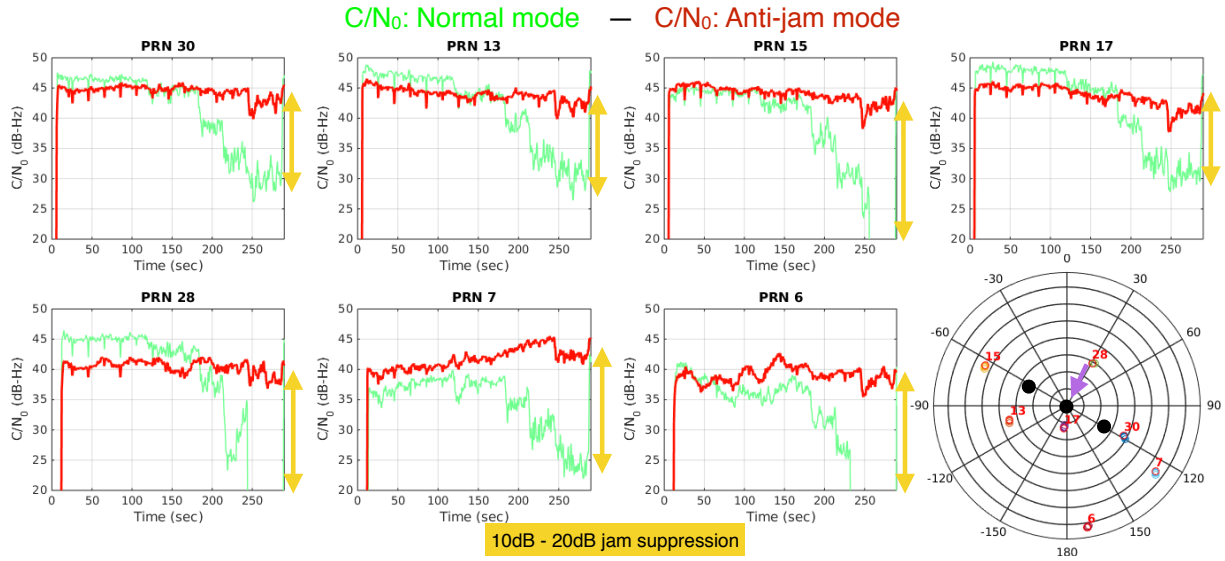

Fig. 10: Comparison of the Stanford GPS SDR's results for the direct "normal mode" stream (in green) and the "anti-jam mode" stream (in red), when the variable phase shifter has been steered to an optimal value of -73°. This variable phase shift value of -73° will steer a null in the direction of the jamming signal resulting in jam suppression ranging from about 10 dB to greater than 20 dB. Anti-jam performance for the high elevation satellites (13, 17 and 30, but excluding satellite 28) increases to around 10 dB of jam suppression. While jam suppression of the lower elevations satellites is generally 20 dB or better, avoiding a loss of lock for several satellites (when compared to the normal mode performance).

360° to provide visibility in every possible direction of attack. Note our assumption of spoofed signals originating from below the horizon requires that we only scan along the azimuth.

As we cycle though all 360° of the phase shifter, we expect to see large oscillatory swings in C/N_0 from the spoofed signals. All of the genuine GPS signals will also have unique phase shift values at which the RHCP and LHCP components of their signal are in phase and out of phase. However, we do not expect to see such large oscillatory behavior in C/N_0 for the genuine signals because the magnitude of the RHCP signal is generally much greater than that of the LHCP signal.

We elected to cycle through the 360° of the phase shifter in 10° steps, waiting 800 milliseconds at each step. The motivation for waiting 800 milliseconds is that we want to outlast the C/N_0 output rate at each phase shift value in order to avoid a smearing of the results over time. Recall that in the case of the Stanford GPS SDR, the coherent integration time is 20 milliseconds and the C/N_0 output rate is 2.5 Hz (or computed every 400 milliseconds). We thus select a detection mode duration that is less than 20 milliseconds and a pause at each phase shift value for longer than 400 milliseconds. Specifically, with the 800 milliseconds interval every 10°, we complete a full 360° revolution of phase shifter values every 28.8 seconds.

Fig. 11 shows the results of implementing both the spoof detection process and the normal mode process in parallel. Recall that the spoofer signal here is the same signal that was previously serving as a jammer. However this time, we have adjusted the Stanford GPS SDR to track both the genuine signals and the spoofer signals. Also recall that in a normal GPS receiver, we would expect to see only one stream (either the green trace or the red trace) at one time. Similar as before, we see the spoofer signal begin at around 120 sec, and appear as a "genuine" signal to the normal mode stream, except for the slightly high C/N_0 values (which have no effect on our detection mechanism). Looking at the spoof detection mode stream, we can see a clear distinction between the genuine and the spoofed signals in the C/N_0 oscillatory behavior.

As indicated above, the amplitude of the oscillatory swing in C/N_0 that we see in Fig. 11 is a loose function of the elevation angle of the energy source. Waveforms from GPS satellites above the horizon will have a larger RHCP component than LHCP component. However, waveforms that originate from below the horizon of an antenna on a large ground plane will be predominately VP, thus having near equal magnitude RHCP and LHCP signals. As we constructively and destructively combine these two signals of near equal magnitude, we expect to see large swing in the resultant C/N_0. However, as we constructively and destructively combine the large RHCP signal with a much smaller LHCP signal, we expect only a small ripple about the nominal RHCP C/N_0 value. Specifically, the depth of the C/N_0 ripple in dB-Hz for the nth satellite (or spoofer)

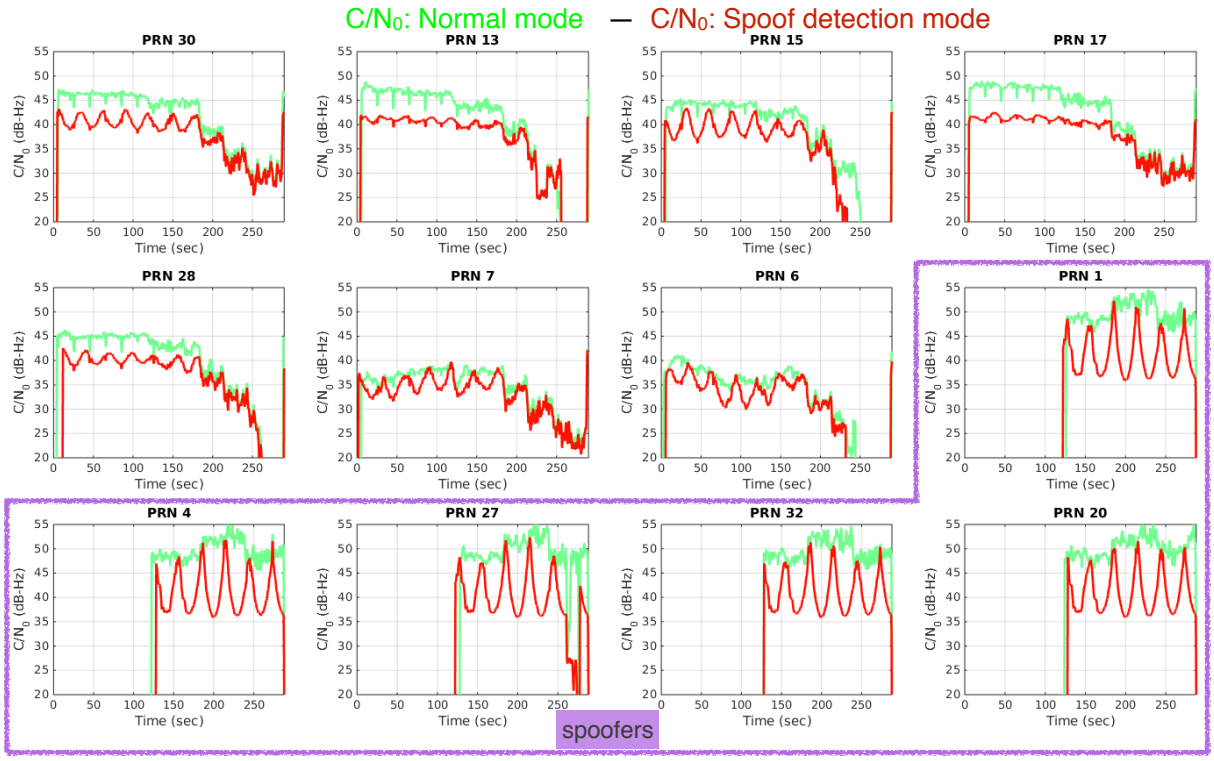

Fig. 11: Comparison of the Stanford GPS SDR's results for the direct "normal mode" stream (in green) and the "spoof detection" stream (in red), when the variable phase shifter has been rotated through 360° every 28.8 seconds. As expected, we see a quite low amplitude swing for the higher elevation satellites (17, 13, 28 and 30), a larger amplitude swing for the lower elevation satellites (6, 7, 15), and the largest amplitude swing for the spoofed satellites (1, 4, 20, 27, 32). It is also noteworthy that each unique satellites has it's own swing amplitude and offset in the time domain of where the peaks and troughs fall during the 28.8 second cycle. However, all the satellites spoofed from a single location will share the same amplitude and time domain offset with one another.

can be calculated for the ideal implementation as:

$$\begin{aligned} R_n &= Constructive_n - Destructive_n \\ &= 10\log\left(g_{RHCP}(\theta_n) + g_{LHCP}(\theta_n)\right) \\ &\quad - 10\log\left(g_{RHCP}(\theta_n) - g_{LHCP}(\theta_n)\right) \end{aligned} \tag{1}$$

Where θ_n is the elevation angle of nth satellite being tracked, and the antenna gain g is shown in lower case to indicate that we are specifying the linear representation of the term, instead of its dB representation.

In this field trial results, we were surprised to see greater C/N_0 oscillatory behavior than was expected, for all signals (both genuine and spoofed). We believe this is an artifact of the Power Ratio Method algorithm used for the SDR's C/N_0 calculation, as later field trials using standard off the shelf GPS receivers produced C/N_0 oscillatory behavior that more closely matched theory. To offset this effect in the SDR, we implemented rapid switching between normal mode and spoof detection mode, with a 5% duty cycle, such that only 1 millisecond out of every 20 milliseconds is spent with the switches in spoof detection mode. The remainder of the time the switches return to normal mode.

As expected we see a quite low amplitude swing for the higher elevation satellites (17, 13, 28 and 30), a larger amplitude swing for the lower elevation satellites (6, 7, 15), and the largest amplitude swing for the spoofed satellites (1, 4, 20, 27, 32). It is also noteworthy that each unique satellite has its own swing amplitude and offset in the time domain of where the peaks and troughs fall during the 28.8 second cycle. However, all the satellites spoofed from a single location will share the same amplitude and time domain offset with one another.

CONCLUSION

This paper has reported the results of a recent field trial where we were able to run hardware-in-the-loop tests to evaluate the performance of both our anti-jam and spoof detection techniques in response to jamming and spoofing attacks. This paper helps to validate our previous work, in which we introduced a backward compatible single antenna design for GPS spoof detection [4] and anti-jam [5] for aviation applications.

Specifically we have validated our prior claims of achiev-

ing greater than 10 dB jam suppression and reliable spoof-detection when these threatening signals originate from below the horizon of the GPS antenna.

Our next steps include a full hardware implementation of both the anti-jam and the spoof detection designs.

ACKNOWLEDGMENTS

The research conducted for this paper took place at the Stanford University Global Positioning System Research Laboratory with funding from the WAAS program office under FAA Cooperative Agreement 12-G-003. Thank you to Abiud Jimenez and his colleagues at JVAB for their generous support during our field trial.

REFERENCES

[1] Y.-H. Chen, J-C. Juang, J. Seo, S. Lo, D. Akos, D. S. De Lorenzo, P. Enge, "Design and Implementation of Real-Time a Software Radio for Anti-Interference GPS/WAAS Sensors," *Sensors* No. 12, pp. 13417-40, 2012.

[2] Ettus Research, "USRP B210 (Board Only)", `http://www.ettus.com/product/details/UB210-KIT`

[3] Minicircuits, "Ultra Low Noise Amplifier", `http://www.minicircuits.com/pdfs/ZX60-1614LN.pdf`

[4] E. McMilin, D. S. De Lorenzo, T. Walter, T. H. Lee, P. Enge, "Single Antenna GPS Spoof Detection that is Simple, Static, Instantaneous and Backward Compatible for Aerial Applications," Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2014), Tampa, FL, September 2014, pp. 2233-2242.

[5] E. McMilin, D. S. De Lorenzo, D. Akos, S. Caizzone, A. Konovaltsev, T. H. Lee, P. Enge, "GPS Anti-Jam: A Simple Method of Single Antenna Null-Steering for Aerial Applications," Proceedings of the ION 2015 Pacific PNT Meeting, Honolulu, Hawaii, April 2015, pp. 470-483.

[6] Y.-H. Chen, S. Lo, D. Akos, D. S. De Lorenzo, P. Enge, "Validation of a Controlled Reception Pattern Antenna (CRPA) Receiver Built From Inexpensive General-purpose Elements During Several Live-jamming Test Campaigns," *Proceedings of the 2013 International Technical Meeting of The Institute of Navigation*, San Diego, California, January 2013, pp. 154-163.

[7] A. Konovaltsev, S. Caizzone, M. Cuntz, M. Meurer, "Autonomous Spoofing Detection and Mitigation with a Miniaturized Adaptive Antenna Array," *Proceedings of the 27th International Technical Meeting of The Satellite Division of the Institute of Navigation*, Tampa, FL, September 2014, pp. 2853-2861

[8] 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.

[9] M. W. Rosen, M. S. Braasch, "Low-Cost GPS Interference Mitigation Using Single Aperture Cancellation Techniques," *Proceedings of the 1998 National Technical Meeting of The Institute of Navigation*, Long Beach, CA, January 1998, pp. 47-58.

[10] ANSYS, "HFSS", `http://www.ansys.com/Products/Simulation+Technology/Electronics/Signal+Integrity/ANSYS+HFSS`

[11] B. Rama Rao, W. Kunysz, R. L. Fante and K. F. McDonald, *GPS/GNSS Antennas*, Artech House, 2013

[12] D. M. Pozar, *Microwave Engineering, 4th Ed.*, John Wiley & Sons, 2012, pp. 343.

[13] W. L. Stutzman, *Polarization in Electromagnetic Systems*, Artech House, 1993, pp. 24.

[14] M.D. Zoltowski, A.S. Gecan, "Advanced Adaptive Null Steering Concepts for GPS", *Military Communications Conference*, IEEE , vol.3, no., pp.1214-1218 vol.3, 8 Nov 1995

[15] A. Konovaltsev, M. Cuntz, C. Haettich, M. Meurer, "Performance Analysis of Joint Multi-Antenna Spoofing Detection and Attitude Estimation," Proceedings of the 2013 International Technical Meeting of The Institute of Navigation, San Diego, California, January 2013, pp. 864-872.

[16] J. Nielsen, A. Broumandan,G. Lachapelle, "GNSS Spoofing Detection for Single Antenna Handheld Receivers", NAVIGATION, Journal of The Institute of Navigation, Vol. 58, No. 4, Winter 2011-2012, pp. 335-344.

[17] M. L. Psiaki, S. P. Powell, B. W. O'Hanlon, "GNSS Spoofing Detection using High-Frequency Antenna Motion and Carrier-Phase Data," Proceedings of the 26th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2013), Nashville, TN, September 2013, pp. 2949-2991.

[18] F. N. Bauregger, T. Walter, D. Akos, P. Enge, “A Novel Dual Patch Anti Jam GPS Antenna,” Proceedings of the 58th Annual Meeting of The Institute of Navigation and CIGTF 21st Guidance Test Symposium, Albuquerque, NM, June 2002, pp. 516-522.