

Authentication by polarization: a powerful anti-spoofing method

Wim De Wilde, Jean-Marie Sleewaegen, Bruno Bougard, Gert Cuypers, [Septentrio](#)
Alexander Popugaev, Markus Landman, Christopher Schirmer; *Fraunhofer IIS*
Daniel Egea-Roca, José A. López-Salcedo, Gonzalo Seco-Granados, *Dept. of Telecommunication and Systems Engineering, IEEC-CERES, Universitat Autònoma de Barcelona*

BIOGRAPHIES

Wim De Wilde (M.Sc. in EE) joined Septentrio in 2002. He works as a system architect, with focus on RF and digital signal processing.

Dr. Jean-Marie Sleewaegen (M.Sc. and Ph.D in EE) is system architect at Septentrio, where he is responsible for GNSS signal processing, system design and technology development.

Dr. Bruno Bougard (M.Sc. and Ph.D in EE) is R&D director at Septentrio, in charge of research, development, and engineering.

Dr. Gert Cuypers (PhD. and M.Sc. in EE) joined Septentrio in 2010. He works as a hardware system engineer, with focus on RF and antennas.

Dr. Alexander Popugaev (M.Sc. and Ph.D in EE) joined Fraunhofer IIS in 2004 and is currently Chief Scientist at the RF and SatCom Systems Department. His main research activities focus on the design of customer-specific GNSS antennas.

Dr. Markus Landmann (M.Sc. and Ph.D in EE) did his Ph.D. in 2008 and is now group leader Over-The-Air-Testing with Fraunhofer IIS and coordinating the research activities. His focus is on signal processing, over-the-air testing and channel parameter estimation.

Dr. Christopher Schirmer (M.Sc. and Ph.D in EE) did his Ph.D. in 2018 in the field of Over-The-Air-Testing and Wave-Field Synthesis. He is working on the development of test procedures for wireless communication devices for applications in mobile communications and GNSS

Dr. Daniel Egea-Roca (M.Sc. and Ph.D in EE) is a post-doctoral researcher at the Signal Processing for Navigation and Communications (SPCOMNAV) group, of the Universitat Autònoma de Barcelona (UAB) and his research is focused on statistical change detection techniques for signal-level GNSS integrity.

Dr. José A. López-Salcedo (M.Sc. and Ph.D in EE) is associate professor at the Department of Telecommunication and Systems Engineering, UAB, and member of the SPCOMNAV research group.

Dr. Gonzalo Seco-Granados (M.Sc. and Ph.D in EE) is an associate professor at the Department of Telecommunication and Systems Engineering, UAB and head of the SPCOMNAV research group.

ABSTRACT

This paper presents a method to detect and mitigate a spoofing attack by means of a dual polarized antenna. It exploits the similarity in polarization of spoofed satellites to identify spoofed satellites and copes with three major challenges. A first challenge is to avoid false alarms, which could be triggered by occasional polarization alignment of authentic satellites. The second challenge is the detection of spoofed signals out of a mix of spoofed and non-spoofed signals, as is the case in most practical spoofing attacks. The final challenge is to be able to work with spoofed signals from RHCP spoofing antennas operating from a higher elevation.

The technique was developed based on analysis of a large amount of experimental signal data recorded in spoofed and non-spoofed environments. The paper first describes the recording system, which uses a high-performance dual polarized antenna, optimized for low axial ratio. This connects to a multi-frequency multi-constellation receiver, supporting concurrent coherent tracking of the RHCP and LHCP signal components provided by the antenna.

We subsequently discuss the measurement campaign. It is rather straightforward to collect data in a variety of non-spoofed environments to build a database of scenarios which are supposed to yield a negative spoofing indication. This doesn't hold for spoofing scenarios, because of regulatory constraints. Therefore, the spoofing tests were done in a special anechoic chamber which can simulate both polarization and angle of arrival of satellite signals. This wave field synthesis (WFS) testbed was configured to create a mix of satellite signals, some of them emulating authentic signals and the other ones representing the spoofer. The WFS testbed was used to simulate an advanced matched power timing attack.

Finally, the paper discusses a new spoofing detection algorithm, based on the experimental data. We present an analysis of the spoofing classification performance, analyzing metrics for probability of false alarm and probability of detection.

INTRODUCTION

Global Navigation Satellite Systems (GNSS) have become a ubiquitous tool for our modern society for vital tasks such as transportation, civil engineering or precision agriculture. This breath has reached the realm of such safety-critical applications as time management of critical infrastructures and autonomous vehicles, in which GNSS is an essential tool nowadays. In this case, GNSS is used for synchronization of communication networks and power-plants, which is a key enabler for the inception of smart cities. Nevertheless, GNSS receivers are vulnerable to radio-frequency interference. In particular, spoofing is becoming one of the main concerns for the GNSS community. The main reason is that a spoofing threat can inject misleading information into these vital systems, potentially with catastrophic consequences.

These types of threats are increasingly becoming a worldwide concern, particularly due to the fact that some recent incidents are suspected to have been caused by spoofing attacks [1,2]. Moreover, with the new trend of safety-critical applications coming up, the motivation of spoofers is increasing, too. From an implementation point of view, spoofing can be done in a large variety of ways, ranging from simplistic to very sophisticated [3]. The former are expected to become a real threat in the coming years thanks to the very affordable software defined radios (SDR) that have entered the market the last few years. These SDRs can be configured to forge GPS signals using open source software. However, the forged signal of this class of spoofers has many artifacts, which can be exploited to easily detect the spoofing attack [4].

Notwithstanding, this doesn't hold for a reradiation attack where, instead of generating the signal, the spoofer captures the RF signal at one spot and re-transmits the very same signal with full level of detail. Most SDRs can be set-up to record and reradiate the signal with very significant changes in timing or position. This is hard to detect, particularly if the genuine signal has been blocked or jammed. Another class of attacks uses higher-end simulation equipment, which can simulate the GNSS signal with great level of detail. This enables an attack, in which the spoofer gradually takes control of the signal and pulls away the timing provided by a GPS receiver from the actual timing. This type of attack was demonstrated in [5], in which an affordable multi-frequency simulator was configured to spoof the timing reference of a power plant.

Many techniques to detect a spoofing attack have been published in the past [6]. Single antenna techniques exploit anomalies in the spoofed signal waveform, its timing or its navigation data. Many of them can be implemented on a standard receiver, but they are inherently vulnerable to some types of attacks, for instance reradiation attacks in which the authentic signals have been blocked. Multi-antenna techniques analyze the consistency of the spatial signature of the electromagnetic field of the signal. In a typical spoofing attack, a single spoofing source generates multiple satellite signals. Hence, these signals will all come from a single direction different from the actual satellite direction. This can be detected by analyzing differential carrier phase measurements between the antenna elements. Unfortunately, multi-antenna systems have an inherently large size. For this reason, dual polar antennas gained interest in recent years, as they provide additional information on the EM field of the signal while their dimensions and ranging performance can be identical to standard industrial grade GNSS antennas. This is because they can use the same radiating element.

Most industrial grade GNSS antennas achieve their right-hand circular polarization (RHCP) by combining two or four feeds from this element, which have a linear radiation pattern. The combination is done with an appropriate electronic phase rotation. Hence,

providing an output sensitive to left-hand circular polarization (LHCP) is a matter of implementing a second electronic phase combiner with associated signal conditioning circuitry rather than changing physical structures. These electronics can be miniaturized. Several publications address the spoofing detection problem with dual polar antennas. In [7] a method is discussed, which alternately provides an RHCP and LHCP antenna output to a receiver, measuring the difference in signal strength via the C/No reported by the receiver. The idea behind this approach is that a spoofer operating from negative elevation would show a much lower signal strength difference between the RHCP and LHCP antenna outputs than authentic signals from high elevation because of the cross-polar behavior of both radiation patterns. Another approach which is discussed in [Error! Reference source not found.] uses a ground plane to convert the polarization of any signal coming from ground level or below into a linear polarization and then combines the RHCP and LHCP antenna outputs with a variable phase shift to determine the presence of the spoofing signal and its angle of arrival. This method is constrained by the need of a large ground plane and is restricted to detecting spoofing signals from low elevation.

CONCEPT

A spoofer which attempts to take control of the receiver would emulate at least four satellites and probably even all satellites in view at the time of interest. All these signals would perceive the same transmission channel from the spoofer towards the target receiver. The authentic signals all perceive different channels. Hence, spoofing can be detected by identifying a set of satellites which have a common set of channel parameters, different from the other satellites. One of these parameters is polarization and this can be measured with a dual polarized antenna. Therefore, identifying a group of satellites with identical polarization could reveal a spoofing attack (Figure 1).

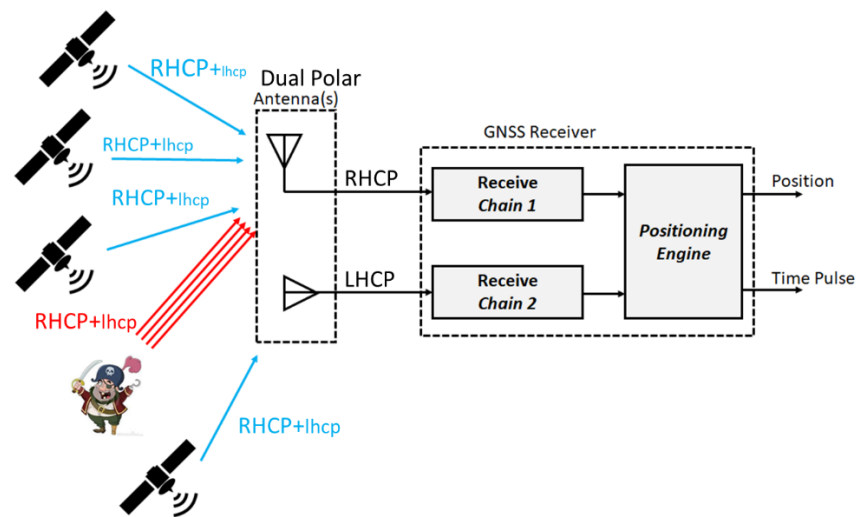


Figure 1: Spoofing Detection by Polarization

Authentic signals are transmitted with a very pure RHCP. However, the associated electromagnetic field (EM-field) at the receiver will usually be degraded because of reflections, resulting in a non-zero LHCP component. The signal from different satellites will usually come from different directions, so that the polarization of the EM-field is expected to differ for every satellite. The EM-field component related to the spoofer would also consist of RHCP and LHCP at the receiver antenna, because of reflections, spoofing antenna polarization, or a combination of both. In this case, though, the underlying EM-fields of the spoofed satellites are all expected to be received with the same polarization. The reason is that they all follow the same path (i.e. from spoofer to receiver). We will base our contribution on this behavior to detect the presence of spoofing with the use of a dual polar antenna.

There are of course many questions related to this approach. In a so-called intermediate attack, in which the spoofer gradually takes control of the signals, the authentic signals and spoofing signals would at least initially have a similar code phasing and would interfere with each other. The polarization vector which the receiver would measure would be a combination of the spoofing signal and the authentic signal, raising the question to which extent spoofing could still be detected in this case. Besides, one could wonder

if the polarization vector of authentic signals could not be sufficiently identical in some situations to trigger false alarms. To investigate this further, a large amount of experimental data was collected both in simulated and real environments.

These activities were done in the scope of the FANTASTIC project (Field Aware Navigation and Timing Authentication Sensor for Timing Infrastructure and Centimeter level positioning). This project included many sub-topics to improve robustness and performance of a PNT solution. The potential of dual polar antennas to improve positioning robustness was analyzed, considering environments with challenging obstacles [9] as well as spoofed environments.

Next, we will describe the hardware used to collect the experimental data.

DUAL POLARIZED ANTENNA

A novel, patent-pending dual circularly polarized antenna has been developed in the project. The achieved results are summarized in Figure 2.

The antenna supports all GNSS signals in the L band and has two RF outputs: RHCP and LHCP. The measured radiation patterns for both the LHCP and RHCP mode are shown in Figure 3. As can be seen, the plots for the co-pol and cross-pol components are almost identical and nearly optimal in terms of multipath suppression and ability to receive signals from satellites at low elevation angles. The antenna provides high polarization purity in all directions of interest with a cross polarization discrimination (XPD, in RHCP mode: difference between RHCP and LHCP gain) of at least 10 dB (axial ratio 5.6 dB), the value in the zenith is not worse than 16 dB (axial ratio 3 dB). It should be noted that a high XPD is very important for the approach described here because of the navigation performance targets, avoiding spill-over of LHCP multipath components into the RHCP line-of-sight signal.

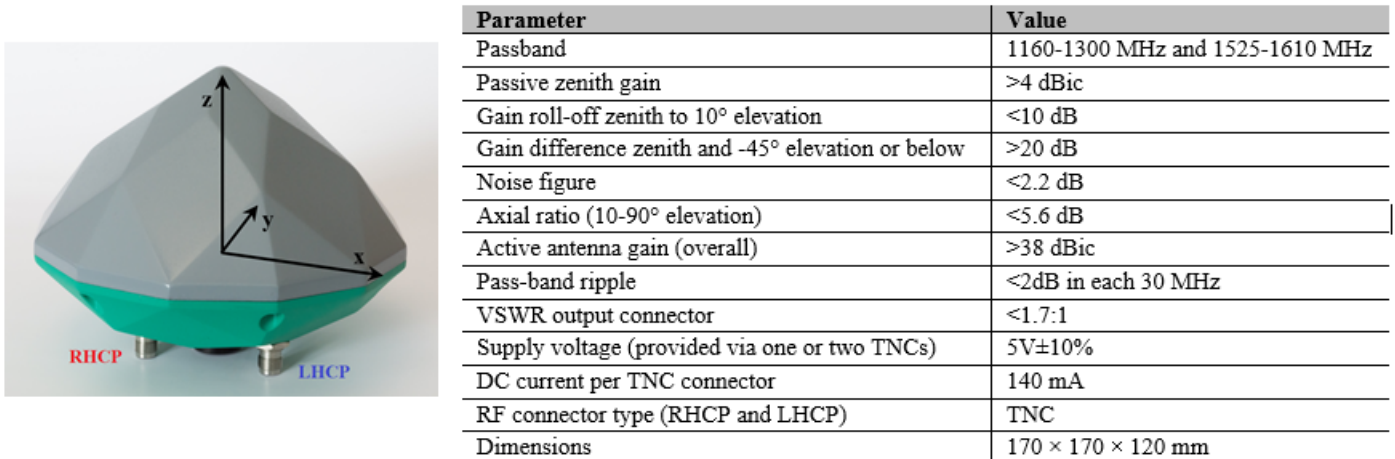


Figure 2: Dual Circularly Polarized Antenna Prototype and Its Main Characteristics




Figure 3: Measured Radiation Patterns (Normalized): LHCP Output (Left) and RHCP Output (Right)

RECEIVER

For data collection, a Septentrio AsteRx-U dual antenna multi-frequency receiver was used. This is a commercial L1/L2/E5/E6 receiver used in many industrial positioning and attitude applications. The RHCP output of the antenna was connected to the main input of the receiver, while the LHCP output connected to the auxiliary input of the receiver. The default software of the AsteRx-U receiver independently acquires and tracks satellite signals from each antenna input. This is a suitable approach for the normal attitude use case of the receiver. However, when used for dual-polarization applications the receiver would fail to permanently monitor the polarization of the signal, as the LHCP component could only be tracked if its C/No is sufficiently high.

Therefore, the receiver software was modified. Rather than having an independent tracking of the LHCP and RHCP components, the receiver only tracks the RHCP component and replicates the local code and carrier timing to correlators connecting to the LHCP signal (Figure 4). In this way the receiver synchronously gathers RHCP and LHCP correlation values, ensuring a permanent polarization monitoring of the signal. Both correlations are integrated coherently over a programmable time ΔT , yielding the following correlation output for satellite signal k at epoch T :

$$C_{R,k}(T) = \int_{T-\Delta T}^T RHCP(\tau) \cdot PRN_k(\tau) \cdot \exp(-j \cdot (\omega_0 \cdot \tau + \varphi_k(\tau))) \cdot \widehat{D}_k(\tau) \cdot d\tau,$$

$$C_{L,k}(T) = \int_{T-\Delta T}^T LHCP(\tau) \cdot PRN_k(\tau) \cdot \exp(-j \cdot (\omega_0 \cdot \tau + \varphi_k(\tau))) \cdot \widehat{D}_k(\tau) \cdot d\tau.$$

For some tests ΔT was set to 100 ms, for others to 1 second.

The PRN-code and local carrier phase φ_k are estimated from the RHCP signal, as well as the navigation data bit \widehat{D}_k . For pilotized signals, only the pilot component was used in the correlation process, making \widehat{D}_k irrelevant. This was done for all satellites in view

from GPS (L1/L2C), GALILEO (E1/E5b), GLONASS (L1CA/L2CA) and BeiDou (B1/B2). The resulting RHCP and LHCP correlations were logged on non-volatile memory in the receiver for post-processing, along with the usual raw GNSS data.

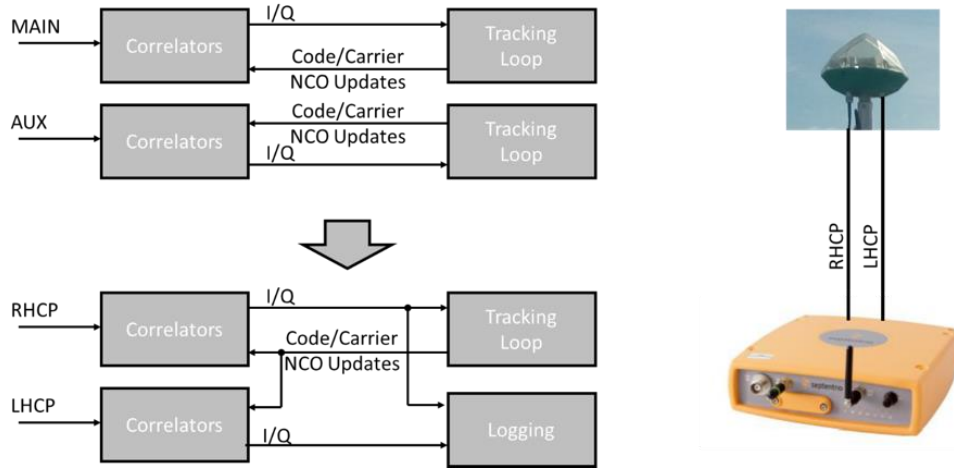


Figure 4: Receiver Modification for Permanent Polarization Vector Monitoring

SIMULATION OF A SPOOFED ENVIRONMENT

Radiated tests are required for the analysis of spoofing detection based on polarization. However, it is very hard to obtain formal regulatory approval for outdoor spoofing tests because of possible interference with air traffic. Instead, the spoofing tests were done in a special laboratory environment, called FORTE (Facility for Over-The-Air Research and Testing, [10]), operated by Fraunhofer IIS. It provides an anechoic chamber (Figure 6), which can emulate the spatial signature of radiation sources using the hardware depicted in Figure 5. This also offers a controlled environment, which drastically simplified the set-up of an advanced spoofing attack.

The hardware in the server racks of Figure 5 provides a connectivity of 12 x 32 hardware channels, i.e. 384 physical signal paths. Each of the 32 RF output signals consists of a superposition of each of the 12 input signals, convolved with arbitrary path weights. The path weights consist of calibration data, wave-field synthesis (WFS) weights, and propagation channel data. The hardware is capable of processing 80MHz wide data in a frequency range from 350MHz up to 6GHz. The digital signal processing is done in the FPGA-based DPSs (FDSP). They deliver the digital data to the S1000 digital-to-analogue up-converters (DAU). For WFS it is indispensable to have phase coherent signals, this is ensured by a common clock distribution of all devices. In our scenario, the 12 digital inputs are fed by a Spirent RFCS (radio frequency constellation simulator) with six genuine satellite signals, and six spoofer satellite signals.

The anechoic chamber setup (Figure 6) consists of three combined constellations for different purposes: a horizontal “ring” (blue) with 24 antennas, a hemispherical 16- and 32-antenna constellation (orange, green), respectively. The 24-antennas constellation is for two-dimensional experiments with linear polarization. The 16-antennas constellation is used for three-dimensional experiments with dual-linear-polarized antennas, and the 32-antennas configuration for three-dimensional experiments with linear polarization. For experiments realizing amongst others RHCP-polarized satellite signals, the 16-antennas configuration is used. With this constellation, we are not only able to emulate RHCP, but also LHCP and linear and elliptic polarized signals.

Before the actual measurement can be performed, a calibration procedure has to be conducted to compensate for free-space loss, non-ideal antennas, different cable lengths and analogue parts of the signal generators. Therefore, an electromagnetic field probe is used to measure the region where the test antenna is placed, called sweet spot. The probe is moved on a grid laid out inside the sweet spot to measure each antennas influence on each of the point positions in three spatial field components (x, y, z). At each of the point positions \mathbf{o} , the target field vectors $\mathbf{e}_{\text{tgt}}(f)$, which is produced by a wave/signal impinging from a certain incidence angle (azimuth, elevation), is defined. By pseudo-inversion of the transfer matrix $\mathbf{X}(f)$ to $\mathbf{X}(f)^+$, which is the measured data of M antennas times N points and multiplication with the target field vectors, the antenna weights $\mathbf{s}(f)$ are calculated, [11].

$$\begin{aligned} \mathbf{X}(f) \cdot \mathbf{s}(f) &= \mathbf{e}_{\text{tgt}}(f), \\ \mathbf{s}(f) &= \mathbf{X}(f)^+ \cdot \mathbf{e}_{\text{tgt}}(f). \end{aligned}$$

The target polarization is defined via $\mathbf{e}_{\text{tgt}}(f)$, which is a vector of field vectors $\mathbf{e}_{\text{tgt}}(f, \mathbf{o})$ that consist of three spatial field components. The calculations to obtain the antenna-/WFS-weights are performed for each spatial component (x, y, z) independently [12]. Wave-field synthesis allows not only to emulate signals impinging from the antenna positions, but due to coherent superposition of elementary waves also from in-between. Due to the short duration of the emulated spoofing scenario, the satellites were not spatially moved according to their orbits but fixed at their initial position within the scenario.




Figure 5: OTA hardware providing 12 x 32 phase coherent signal channels and a Spirent GSS9000 TS788 satellite emulator




Figure 6: Anechoic chamber with 3D WFS setup and test antenna (left), and constellation visualization (right)

SIMULATED ATTACK

The simulated spoofing attacks represented an intermediate matched power timing attack. In this type of attack, the spoofer transmits a signal which initially matches the authentic signal at the receiver antenna but has a slightly higher signal strength to get control of the tracking loops. The timing of the spoofing signal is then gradually delayed, pulling away the tracking loops from the authentic signals. This will be reflected in the timing provided by the receiver.

The authentic satellites were emulated with 6 simulator channels. The wave field synthesis was configured to simulate the angle of arrival of the signal and its right hand circular polarization. A different set of 6 simulator channels were allocated to represent the spoofer. They were enabled three minutes after the start of the test scenario, to study the transient from authentic to spoofed signal. The orbits were duplicated from the “genuine” channels, but the satellite clock was configured to diverge at a fixed rate of 5 ns/s. This is illustrated in Figure 7.

As in real GNSS systems, the authentic signals have a nearly fixed incident power level, regardless their elevation and azimuth. The receiver antenna on the other hand, shows large gain variations over elevation (Figure 3), resulting in a variety of C/No levels at receiver level. To remain unnoticed, the spoofer should attempt to only slightly overrule these signals, e.g. by creating a 2 dB higher C/No. The spoofer signals come from a single direction and hence a single antenna gain applies, usually different from the gain perceived by the authentic signal. Therefore, the power of the spoofing signals had to be adjusted, making up for the receiver antenna gain difference. This was done as follows. We ran first a scenario which just included the authentic signals. The associated C/Nos were registered. Afterwards the scenario was repeated, but this time the signals were transmitted by the spoofing antenna. This produced a different set of nearly fixed C/Nos. The C/No difference of between the “authentic” and spoofed satellites was then used to correct the spoofer transmit levels, adding 2 dB to overrule the authentic signal.




Figure 7: Spoofing Test Scenario

The tests were done for various spoofer antenna elevations, always using an RHCP spoofing antenna emulated with the WFS system. The pulse-per-second (PPS) signal from the receiver was monitored on an oscilloscope, using the simulator PPS signal as a trigger source. The simulation produced the desired effect: the PPS output from the receiver started drifting at a rate of 5 ns/s when the spoofer kicked in. When the spoofer was activated, the receiver reported a sudden C/No increase and a clock jump (Figure 8). It is however not possible to use these as a reliable spoofing indication, as an identical behavior could be triggered by respectively regular interference and mechanical effects.




Figure 8: GALILEO E1BC C/No jump and clock drift upon spoofing

POLARIZATION ANALYSIS

The receiver gathered the correlations $C_{R,k}(T)$ and $C_{L,k}(T)$ for each signal k in the test, with a 1 second integration interval ($\Delta T=1s$). The corresponding polarization vector $P_k(T) = C_{R,k}(T)/C_{L,k}(T)$ could be used to build up a diagram in which the horizontal coordinate corresponds to the magnitude $20 \cdot \log_{10}(\|P_k(T)\|)$ and the vertical coordinate to the phase $arg(P_k(T))$. If we now look at this diagram before the spoofer is present, we get Figure 9(a). Each dot in this figure represents the polarization vector of the indicated GALILEO satellite. We can see a large spread in phase and amplitude. This is set by the cross-polar properties of both the transmit antennas and the receive antenna, as multipath was absent in the simulation. Because the (simulated) authentic signals all come from different directions, the ratios are all uncorrelated.

After three minutes in the test, the spoofer radically changes the diagram. The polarization of all satellites is suddenly clustered in 3 dB by 20° area. Hence, the detection of a cluster like this could be indicative for a spoofing attack. Of course, we should consider that not necessarily all satellites are being spoofed in a practical scenario, resulting in a cluster of spoofed satellites as in Figure 9(b) superimposed on a random scattering of the remaining satellites.




Figure 9: Polarization diagram before (a) and during (b) spoofing

This may seem quite surprising, as the spoofing signal and authentic signal have nearly the same code phase and comparable signal strength. However, we should consider the 1 second integration time. The 5 ppb drift would result in 7.87 revolutions per second of the authentic signal component if the spoofed signal is tracked. This averages out during the 1-second integration, leaving only a residual perturbation with small amplitude.

OUTDOOR TESTS

We could wonder if the polarization diagram could not accidentally show a cluster in an authentic scenario, due to multipath or antenna peculiarities. For this we could piggyback on the large datasets recorded with the dual polar antenna in the scope of multipath mitigation [9]. These sets were recording in a variety of environments. The first environment was a very benign environment. The antenna was placed on a tripod on a farmer's field with dry soil and no objects anywhere nearby (Figure 10a). This minimizes the effect of multipath. The reported polarization will mainly be set by the cross-polar properties of the antenna. The resulting amplitude-phase scatter diagram was showing many dots, most of them with an RHCP/LHCP ratio above 15 dB. Figure 10b is showing this diagram for the GPS L1 C/A and GALILEO L1BC signals, with a bias of 70 for the GALILEO satellite IDs. The outdoor recordings were done with a receiver integration time of 100 ms, but this was further coherently integrated to 1 second to reduce thermal noise and study polarization dynamics. The dots are moving slowly but consistently, because of the satellite motion. Often dots get grouped in a small cluster, but this eventually falls apart because of this motion.

Another recording was done under dense tree foliage. The amplitude-phase diagram still shows many dots, but their dynamic behaviour is radically different. The dots move much more rapidly. Clusters are frequently occurring, but only last for a very short time. An example is in Figure 11, showing a large cluster of satellites at the left, which is completely gone 14 seconds later, as clear from the scatter plot at the right.

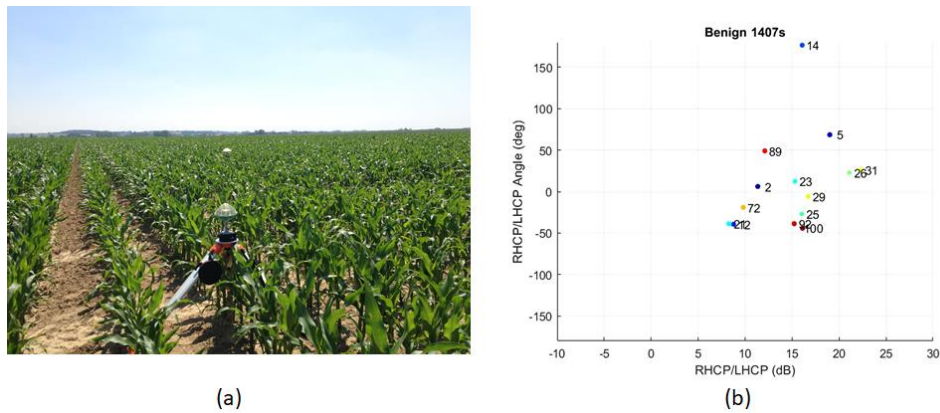


Figure 10: Benign Environment

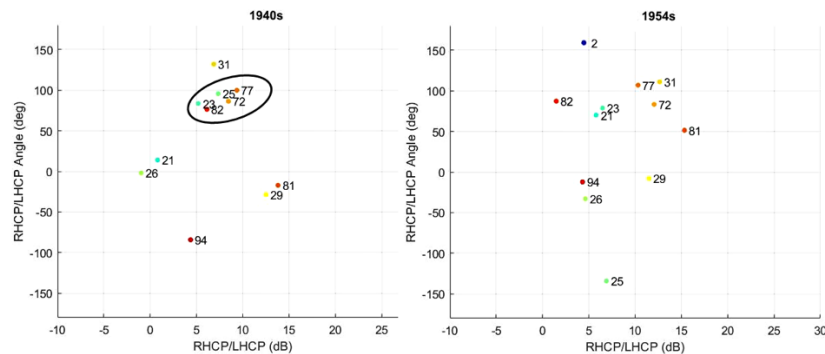


Figure 11: Polarization scatter diagram under tree canopy

The third recording was done in an urban canyon. This shows a limited number of satellites, and the dots in the scatter plot are moving faster than in the benign scenario, but much less nervously as in the foliage-scenario. The dots move in consistent patterns, relating to the systematic reflections on the walls of the buildings.

Finally, a dynamic test was done, collecting data during a two-hour ride of the test van in a mix of environments, including highways, forests, sub-urban and dense urban environments. As could be expected, this most often shows an polarization diagram with very rapidly moving points.

SPOOFING DETECTION ALGORITHMS

From the analysis of experimental data, we could see that spoofing can be detected by looking for point clusters in the polarization diagram, which are stable over time. This stability refers to the sets of satellites in the cluster rather than the position of the cluster in the diagram, as the spoofing channel may be time variant.

The polarization diagram could be directly scanned for increased densities to detect clusters. This can be done by counting the number of points in a sliding 2D window of fixed size, or by using a more intelligent clustering algorithm like k-means. However, the computational complexity of this would be quite high. Besides, the density which would need to be reached to consider a set of points as a cluster would be set by heuristics and would be sub-optimal as it doesn't consider noise statistics.

Instead, a more optimal detection algorithm was developed, which even allows to detect smaller clusters, with just two or three spoofed satellites, which would start challenging the standard receiver autonomous integrity monitoring (RAIM) algorithms. This is discussed in the next paragraph.

PROPOSED DETECTION ALGORITHM

Let $x_i(n) \in \mathbb{C}$ be the ratio $P_i = C_{R,i}/C_{L,i}$ at time instant n , then from the results previously shown we can define the following model for hypothesis testing

$$\begin{aligned} \mathcal{H}_0: & \quad E[x_i(n)] \neq E[x_j(n)] \quad \forall i \neq j \\ \mathcal{H}_1: & \quad E[x_i(n)] = E[x_j(n)] \quad \text{for some } i = j \end{aligned}$$

That is, in the absence of spoofing (\mathcal{H}_0) all the satellites have different polarization ratio (some of them may coincide in a short period of time), whereas when a spoofer is present (\mathcal{H}_1) those satellites that are spoofed have the same polarization ratio. Based on this model we can build up a distance measure between pairs of satellites i and j , namely $d_{ij}(k)$. Thinking about complex numbers (i.e. $x_i(n) \in \mathbb{C}$), the idea is that under \mathcal{H}_1 $d_{ij}(k)$ will be a short value, denoting likelihood between polarization ratio, whereas it will be a large value under \mathcal{H}_0 , denoting different polarization ratio. So, we can use the test $d_{ij}(k) \leq h$ to decide whether the pair of satellites ij is spoofed or not.

To define a proper distance measure, let us first define

$$c_{ij}(m) = \frac{\left\| \sum_{n=(m-1)N_{\text{coh}}+1}^{mN_{\text{coh}}} (x_i(n) - x_j(n)) \right\|^2}{N_{\text{coh}} \sigma_{ij}^2(m)},$$

where N_{coh} denotes the number of samples that are coherently added and

$$\sigma_{ij}^2(m) = \frac{1}{N_{\text{coh}}} \sum_{n=(m-1)N_{\text{coh}}+1}^{mN_{\text{coh}}} \|x_i(n) - x_j(n) - \mu_{ij}(m)\|^2, \quad (1)$$

$$\mu_{ij}(m) = \frac{1}{N_{\text{coh}}} \sum_{n=(m-1)N_{\text{coh}}+1}^{mN_{\text{coh}}} (x_i(n) - x_j(n)). \quad (2)$$

It is worth pointing out that $c_{ij}(m) \approx 1$ under \mathcal{H}_1 and $c_{ij}(m) > 1$ under \mathcal{H}_0 . The coherent integration is applied in order to avoid false clusters due to multipath or any other effect that may cause similar polarization between satellites under nominal conditions. As shown in Figure 11, sometimes under absence of spoofing some satellites have a similar polarization vector. Nevertheless, these polarization vectors rapidly depart from each other, usually within around 15 – 20 s. So, we might extend the integration time to around 20 s in order to avoid false clusters under nominal conditions, but at the expense of increasing the time we need to detect the spoofer.

If the antenna would rotate, the angle of the polarization vectors would change by twice the rotation rate due to phase windup. This could cancel the angular information in the coherent integration, leading to a less powerful discrimination. Therefore, we will limit

the coherent integration to two seconds. We would then average several samples of $c_{ij}(m)$ to extend the monitoring interval. We will refer to this process as non-coherent integration. Doing so, we get our distance measure given by

$$d_{ij}(k) = \frac{1}{N_{nc}} \sum_{m=(k-1)N_{nc}+1}^{kN_{nc}} c_{ij}(m), \quad (3)$$

with N_{nc} the number of samples of $c_{ij}(m)$ that are non-coherently averaged.

Figure 12 shows the behavior of the polarization ratio and the corresponding distance for a scenario with a spoofer appearing at sample 130. In the left plot we have the polarization ratio for 6 satellites, which is different for all the satellites (except for SVN1 and SVN6 that are close) under \mathcal{H}_0 (before sample 130). On the other hand, under \mathcal{H}_1 we see how the polarization ratio of all the satellites converges to a similar value after the cross-talk phase (between around 130 and 300 samples). This is because in this experiment all the satellites are spoofed. In the right plot we see how this behavior of the polarization ratio is translated into the expected behavior of the polarization distance between satellites (i.e. close to 1 under \mathcal{H}_1 and departs from 1 under \mathcal{H}_0). In this case, the set of 6 satellites rise 15 distances between pairs of satellites.

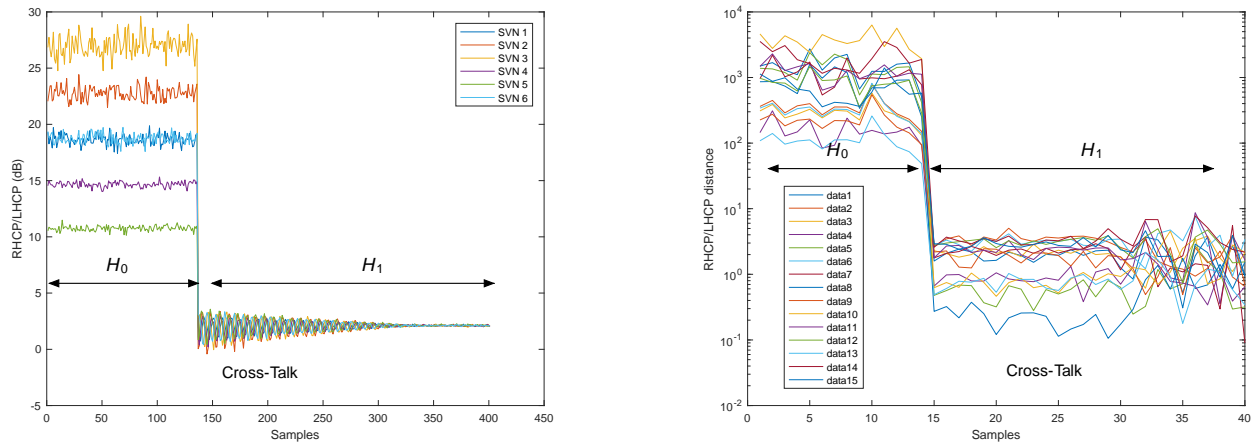


Figure 12 Behavior of the polarization ratio (left) and its corresponding RHCP/LHCP distance for non-coherent time of 10 s (right).

DETECTION METRIC

We can think of fixing the detection threshold h so that when $d_{ij}(k) \leq h$ a spoofing event is declared (recall that under a spoofing event $d_{ij}(k)$ is smaller than in the case of absence of spoofing). The spoofing attack can be mitigated in this case by removing satellites i and j from the positioning solution. Traditionally, h is fixed through the statistical distribution of the detector in order to have a fixed value of probabilities of false alarms or missed detection [13], formally

$$P_{FA} \doteq \Pr\{d_{ij}(k) \leq h | \mathcal{H}_0\} = F_0(h),$$

$$P_{MD} \doteq \Pr\{d_{ij}(k) > h | \mathcal{H}_1\} = 1 - F_1(h),$$

and then if we would like to have $P_{fa} = \alpha$ or $P_d = \beta$ we would fix the detection threshold as $h = F_0^{-1}(\alpha)$ or $h = F_1^{-1}(\beta)$, respectively, where $F_i(h)$ stands for the cumulative distribution function (cdf) of $d_{ij}(k)$ under hypothesis \mathcal{H}_i , with $i = 0,1$, evaluated at h . Unfortunately, the cdf of $d_{ij}(k)$ is unknown in a closed-form, so that we cannot use the traditional procedure to fix the detection threshold and we must rely on some alternative method. We propose to fix the detection threshold based on the mean μ_1 and variance σ_1^2 of $d_{ij}(k)$:

$$h = \mu_1 + \gamma\sigma_1$$

The design parameter γ is to be fixed experimentally.

This can be further simplified if we focus on \mathcal{H}_1 , in which case the mean equals 1.

$$h = 1 + \gamma\sigma_1 \quad (4)$$

We will use the experimental data gathered within the framework of the FANTASTIC project in order to characterize the statistical distribution of the polarization distance.

We are interested on analyzing the behavior of the proposed polarization distance (see (3)) as a function of the coherent integration time. The goal is to analyze the number of samples needed to satisfy the statistical characterization performed in the previous section. To do so, we show in Figure 13 the polarization distance obtained after analyzing one of the spoofing tests. Specifically, the test was composed of 6 satellites received under nominal conditions during the first 136 s, and then all the satellites are spoofed. We see in the figure how the polarization distances clearly show a change just when the spoofer appears. The results have been obtained with a non-coherent time of 10 s, this is why the plots show a change around sample 15 (the spoofing tests has a sampling rate of 1 Hz).

In the left plot, we show the results for $T_{\text{coh}} = 2$ s (i.e. $N_{\text{coh}} = 2$ samples). We see how the distance change when the spoofer appears, but the statistical characterization is not completely fulfilled. To see so, let us focus on the behavior of the distance after the cross-talk period (up to sample 110). We know from **Error! Reference source not found.** that $\mu_1 = 1$, but we can see how in the left plot the mean of the distance after the cross-talk period is greater than 1. Moreover, we can see some instabilities of the distance. The reason is that we are only using 2 samples to coherently integrate the distance, this does not seem enough to fulfill the expected statistical behavior. This is confirmed with the right plot of Figure 13, in which we have used $T_{\text{coh}} = 5$ s. In this case, we see how the distance is around the expected mean (i.e. 1) and it shows a more stable behavior. So, we can say that 5 samples are enough to have a good statistical matching, and then for a practical implementation it would be nice to average around 5 – 10 samples. It is worth pointing out that the outdoor scenarios were performed with a sampling rate equal to 10 Hz, and then we can safely use the maximum coherent integration time allowed to avoid phase cancellation (i.e. $T_{\text{coh}} = 2$ s). This will provide around 20 samples to coherently integrate, thus enough to have a good statistical matching.

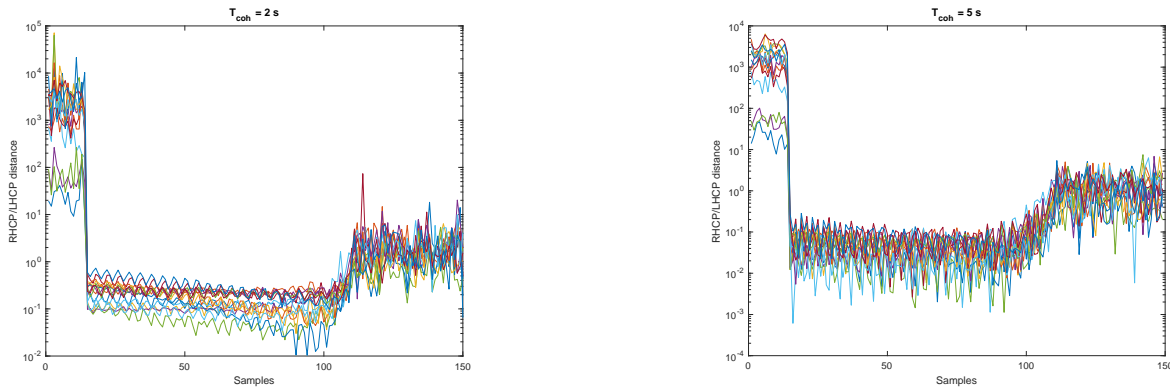


Figure 13 Detection metric for different coherent integration times (i.e. 2 s in the left and 5 s in the right) and $T_{\text{nc}} = 10$ s.

PROBABILITY OF FALSE ALARM

In this section, we analyze the data collection campaign including all the outdoor scenarios: Benign, Static Urban, Static Foliage and Dynamic. The goal here is to use these data to evaluate the statistical properties of the distance under nominal conditions (i.e. \mathcal{H}_0). More specifically, we want to analyze the probability of false alarms as a function of the non-coherent and coherent times. This is shown in Figure 14, which shows the probability of false alarm as a function of T_{nc} and for different T_{coh} 's in the range $[0.5, 2]$ s. The probability of false alarms is computed as the ratio of the number of spoofing events and the total number of epochs. A spoofing event is defined as those epochs in which $d_{ij}(k) \leq h$ for some pair of satellites ij . Since we are analyzing data under benign conditions, such epochs are false alarm events.

In the left plot of Figure 14 we show the results for the case of using $\gamma = 3$ in order to fix the threshold (see (4)). As expected, the probability of false alarms decreases as we increase both coherent and/or non-coherent time. The reason is that the longer we integrate the less probable is that two polarizations of different satellites match. Particularly, we see lineal behavior of the probability of false alarm as a function of the non-coherent time (in log-scale), for a fixed coherent time. The used data campaign was composed of

around 10^5 samples, this is why we cannot see values lower than 10^{-4} and the results for $T_{nc} > 4$ s are not very precise, and for $T_{nc} > 15$ s we do not get any false alarm event. Notwithstanding, the curves shown in Figure 14 are very useful to do a line fitting and be able to select the proper T_{nc} for a given T_{coh} and the desired value of false alarms. For instance, we see the curve of $T_{coh} = 2$ s decreases from $P_{FA} \sim 10^{-2}$ to $P_{FA} \sim 10^{-4}$ by increasing T_{nc} from 2 s to 4 s. So, if we would like to fix $P_{FA} \sim 10^{-6}$ we should use $T_{nc} \geq 6$ s, similarly if we use $T_{nc} = 10$ s we would get $P_{FA} \sim 10^{-10}$. In the right plot of Figure 14 we show similar results but using $\gamma = 10$. We see how the obtained values for the probability of false alarm are larger for a given $\{T_{coh}, T_{nc}\}$, as expected (see right plot of **Error! Reference source not found.**). In this case, it is for $T_{nc} > 20$ s when we do not get any false alarm event. This is shown also shown in Figure 16, which illustrates the behavior of the polarization distance in the Benign scenario. We use $T_{coh} = 2$ s and $\gamma = 10$ and two different values for T_{nc} . In the right plot we use $T_{nc} = 5$ s and we see some false alarm events (i.e. epochs in which the distance is below the threshold). On the other hand, the left plot shows how when using $T_{nc} = 20$ s we do not get any false alarm event.

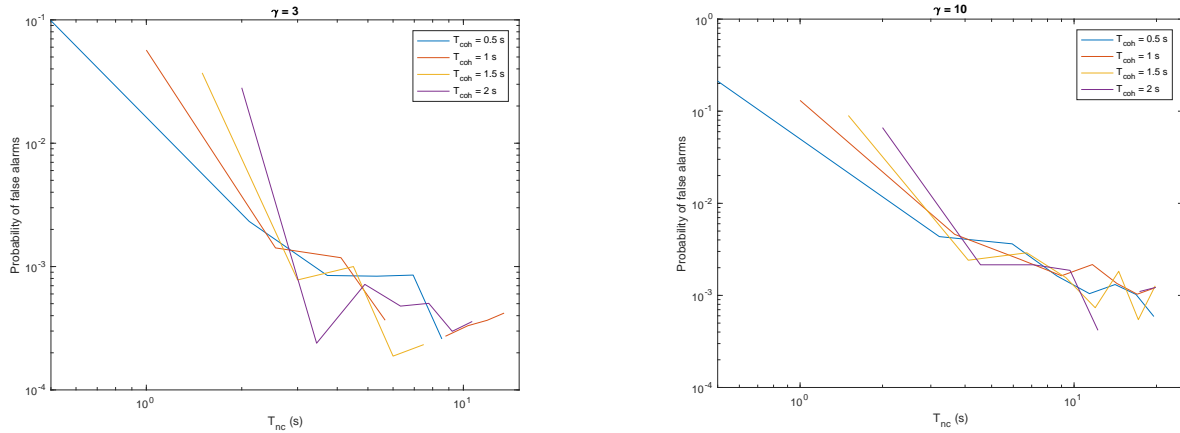


Figure 14 Probability of false alarm as a function of the non-coherent time for different coherent times. Different plots for different values of γ , used to fix the detection threshold.




Figure 15 Detection metric behavior in the Benign scenario for different non-coherent integration times and $T_{coh} = 2$ s. Polarization distance (solid lines) and detection threshold (dashed line) using $\gamma = 10$.

PROBABILITY OF DETECTION

We analyze the detection performance of the proposed distance measure for a $T_{coh} = 5$ s and $T_{nc} = 10$ s. Figure 16 shows the polarization distance for the spoofing test under analysis together with the detection threshold. In the left plot we show the results when using $\gamma = 3$. We see how when the spoofer appears, the distance measure is below the detection threshold most of the epochs.

We see some epochs after the cross-talk period that some distances are above the threshold, so they will not raise an alert. Nevertheless, the spoofer is detected all the epochs. The reason is that although some distance is above the threshold the rest are below so that they will raise an alert. For instance, imagine the case we have $d_{12} \leq h$ and $d_{13} \leq h$, but $d_{23} > h$, the last test will tell us that satellite 2 and 3 are not spoofed, but the two first tests will tell us that satellite 1, 2 and 3 are spoofed, so that they will be detected anyway. We can avoid these ambiguities by increasing the value of γ as shown in the right plot of Figure 16, which shows how for $\gamma = 10$ all the distances are below the threshold when the spoofer is present.

Finally, we should remark that the detection time can be shortened in case multiple satellite pairs simultaneously show a match, for a given probability of detection.

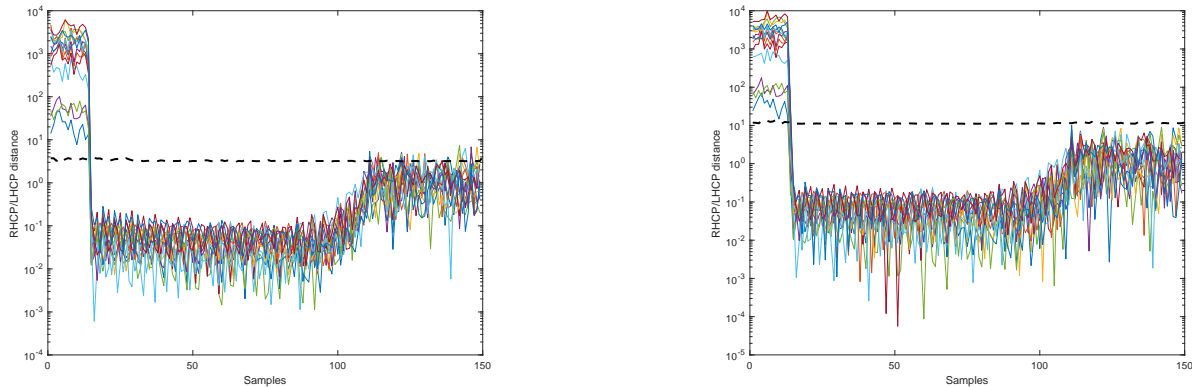


Figure 16 Analysis of the detection performance for different values of γ (i.e. $\gamma = 3$ left and $\gamma = 10$ right) for the spoofing test under analysis.

CONCLUSION

It was demonstrated that a spoofing attack could be revealed by comparing the complex polarization vector of multiple satellites as provided by a dual polarized antenna. The polarization of spoofed satellites would be grouped in a cluster which lasts over time. Authentic signals could also be grouped in a cluster, but these clusters are only stable for a short while. These observations were used to define a detection metric which operates on the polarizations of a pair of satellites. The metric can detect spoofed satellite pairs within a few tens of seconds, without causing false alarms on authentic signals.

ACKNOWLEDGMENTS

This work was partly supported by the European GNSS Agency (GSA) under the FANTASTIC project (GSA/GRANT/01/2016). The authors wish to thank Instituto Superiore Mario Boella (ISMB) and GMV, as part of the consortium that carried out the FANTASTIC project, for their valuable support.

REFERENCES

1. Daniel Shepard, Jahshan A. Bhatti, Todd E. Humphreys, "Drone hack: Spoofing attack demonstration on a civilian unmanned aerial vehicle", *GPS World*, Vol. 1, n° Dec., 2011
2. CNN, "Getting lost near the Kremlin? Russia could be GPS spoofing", CNN tech, Retrieved from: <http://money.cnn.com/2016/12/02/technology/kremlin-gps-signals/index.html>, Accessed: 18-6-2017, 2016.
3. Todd E. Humphreys, Brent M. Ledvina, Mark L. Psiaki, Brady W. Hanlon, Paul M. Kintner, "Assessing the spoofing threat: Development of a portable GPS civilian spoofer", *Proceedings of the 21st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION)*, pp. 2314-2325, 2008.
4. De Wilde, Wim et al, "Spoofing Threats: Reality Check, Impact and Cure," Proceedings of the 30th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2017), Portland, Oregon, September 2017, pp. 1289-1327.

5. lie, Iurie et al, "Spoofing of Electrical Power Grid: It's Easier Than You Think", Proceedings of the 30th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2017), Portland, Oregon, September 2017, pp. 1383-1408.
6. Ali Jafarnia-Jahromi et al, "GPS Vulnerability to Spoofing Threats and a Review of Antispoofing Techniques", *International Journal of Navigation and Observation*, Volume 2012, Article ID 127072
7. Emily McMilin, "Single Antenna GPS Spoof Detection that is Simple, Static, Instantaneous and Backwards Compatible for Aerial Applications", *Proceedings of the 27th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2014)*, Tampa (Fla.) September 2014, pp. 2233-2242
8. Chun Yang, Ananth Vadlamani, Andrey Soloviev, Michael Veth, Clark Taylor, Demonstrating Single Element Null Steering Antenna Direction Finding for Interference Detection, *Proceedings of ION ITM 2018*, Jan 2018, pp 240-259
9. GNSS Measurement Exclusion and Weighting with Dual Polarized Antenna: The FANTASTIC project, Daniel Egea-Roca et al, *ICL-GNSS.2018.8440897*
10. Fraunhofer IIS, "Facility for Over the Air Research and Testing (FORTE)", 2018, online: <https://www.iis.fraunhofer.de/en/profil/standorte/forte.html>.
11. C. Schirmer, "Over-The-Air Testing using Wave-Field Synthesis", *Ilmenau: Universitätsverlag Ilmenau*, 2018.
12. C. Schirmer, M. H. Landmann, W. A. T. Kotterman, M. Hein, R. S. Thom\ "a, G. Del Galdo und A. Heuberger, "3D wave-field synthesis for testing of radio devices", *Antennas and Propagation (EuCAP)*, 2014 8th European Conference on, 2014.
13. Steven M Kay, "Fundamentals of Statistical Signal Processing", Volume 2: *Detection Theory*, Prentice Hall Upper Saddle River, NJ, USA, 1998