NON-COHERENT SPOOFING DETECTION WITH REAL-WORLD SPOOFING ATTACKS

M. Dorn[1], J. Dampf [2], T. Pany[3], W. Bär[4], J. Winkel[4], L. Mervart[5], J. A. Avila-Rodriguez[6], R. Ioannides[6]

[1] IGASPIN GmbH
Reininghausstraße 13a, 8020 Graz, Austria
Email: m.dorn@igaspin.at

[2] Trimble Terrasat GmbH
Haringstraße 19, 85635 Höhenkirchen-Siegertsbrunn, Germany
Email: juergen_dampf@trimble.com

[3] Universität der Bundeswehr München
Werner-Heisenberg-Weg 39, 85577 Neubiberg, Germany
Email: thomas.pany@unibw.de

[4] IFEN GmbH
Alte Gruber Straße, 85586 Poing, Germany
Email: w.baer@ifen.com, j.winkel@ifen.com

[5] Department of Geomatics, Czech Technical University
Thakurova 7, 166 29 Praha 6, Czech Republic
Email: mervart@fsv.cvut.cz

[6] ESA/ESTEC
Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
Email: jose.angel.avila.rodriguez@esa.int, rigas.ioannides@esa.int

ABSTRACT

This article demonstrates that a synthetic aperture antenna can reliably detect and mitigate even sophisticated spoofing attacks rendered against Global Navigation Satellite Systems. The direction-of-arrival is a reliable metric to discriminate spoofing signals from line-of-sight (LOS) signals and to also localize one or more spoofers with high angular resolution of two degrees. A special focus is given to the detection of non-coherent spoofers, whose signals are not perfectly aligned with the true line-of-sight signals of the satellites, and to the detection of spoofers that are transmitting incorrect data bits.

1. INTRODUCTION

Today many applications rely on GNSS (Global Navigation Satellite Systems) and the number is continuously growing. Some of these applications also incorporate GNSS reference station data to improve their navigation solution. Misleading or degrading a GNSS navigation solution can have serious harmful impacts, especially when thinking about Safety-of-Life services. GNSS spoofing is an intentional attack on a GNSS receiver to mislead or degrade the navigation solution. Spoofing is considered as a serious threat, especially when spoofing GNSS reference stations that distribute their degraded or falsified correction data to many GNSS users.

Whereas the position of a reference station (and its time) is typically well known and cannot be spoofed, a sophisticated spoofing attack may induce multipath like effects or ionospheric-like effects on the measured pseudoranges and carrier phases. This attack will degrade the performance of the reference station and the service relying upon it. These spoofing signals do not require a high signal power and thus may be well below the line-of-sight signal power. They are thus very difficult to detect as standard methods like signal-quality-monitoring, C/N_0 monitoring, or a time series analysis still see the line-of-sight signal as the main contribution (see [2]). Direction of arrival (DoA) estimation, addressed in this article, efficiently detects these attacks by making use of a synthetic antenna aperture and advanced detection and mitigation techniques.

A GNSS signal spoofer can be realized with various degrees of fidelity. In the simplest case, a GNSS signal is recorded and played back using commercial record and replay systems. In that case, one may also speak of a meaconing attack and the target receiver will see the position of the recorded signal. More sophistication is achieved if a GNSS RF simulator transmits a GNSS signal over air. This already allows for inducing an arbitrary position and time on the target receiver. Linking the spoofed position and time to the true position and time (in order to make the attack less obvious) requires further technology. Whereas the position link is easily established, if the true position of the target is known, time requires synchronizing the signal generator to the true GNSS time and frequency. This requires that a dedicated GNSS receiver provide a pulse-per-second (PPS) output to the spoofer signal generator. Even more sophistication is required if the spoofer attempts to broadcast an identical navigation message as the satellites. This will render the spoofing signal even less distinguishable from the true signal. As the message needs to be broadcast in real-time by the spoofer, it is necessary to predict the message, as a data link from the data message capturing receiver to the spoofer will always have some latency. [6] shows that producing a perfectly aligned spoofing signal with correctly predicted data bits and without an additional time offset and drift due to imperfections of the spoofer clock is not an easy task. For that reason, the authors assume that in future non-coherent spoofing signals will arise. The detection, mitigation, and eventually the localization of non-coherent spoofers will get more and more important.

This paper summarizes experiences gained with software simulations and real-world spoofing attacks. Section 2 describes the realized algorithm for spoofing detection of coherent and non-coherent spoofers. The realization is based on direction-of-arrival discrimination and is using a rotating GNSS antenna employing synthetic aperture processing and an adaptive beamforming algorithm. This GNSS receiver/antenna system not only increases the resilience of GNSS reference networks, which are otherwise very vulnerable against sophisticated spoofing attacks, but also allows to localize the spoofer with high accuracy. Section 3 describes the test data. The real-world spoofing attacks were conducted using a modified GNSS radio-frequency (RF) signal generator. Section 4 summarizes the gained results with special focus on the detection of non-coherent spoofers. Finally, section 5 concludes the paper.

2. SPOOFING DETECTION ALGORITHM

2.1. Spoofing Detection Via DoA Discrimination

Spoofing signals can be easily distinguished from true GNSS signals if the DoA can be estimated. Spoofing signals will most likely come from a ground based transmitter (thus arriving at a low elevation to the target receiver) and the DoA will be identical for all signals. DoA is of course different for each true satellite signal.

DoA estimation can be done with a proper GNSS receiver plus antenna, provided that multiple antenna elements are used within a phased array system (see [2]). An alternative approach is to use a synthetic aperture GNSS antenna exploiting the antenna motion to combine GNSS signals received at different spatial locations to optimize a certain performance criterion. Like phased array antennas, the synthetic aperture GNSS antenna allows to form a certain antenna gain pattern and can thus be used to eliminate the effect of spoofing signals. Synthetic aperture antennas have so far received only limited attention from the GNSS community. Proof-of-concepts have been shown conducted by [3]. [4] investigated several signal processing options for synthetic aperture antennas.

The work presented here uses a rotating GNSS antenna (similar to [3]), but with an updated mechanical design rendering it water and ice proof (see Figure 1). The antenna motion is measured precisely with a magnetic sensor allowing determination of the antenna position with sub-millimeter precision at every instant. The antenna rotates at a rate of one hertz and has a rotation radius of 50 centimeters. The rotation plane is horizontally aligned. A rotating antenna is mechanically relatively easy to realize and all mechanical components can be chosen for long-term operation without any maintenance. An RF slip ring is needed to connect the GNSS antenna.

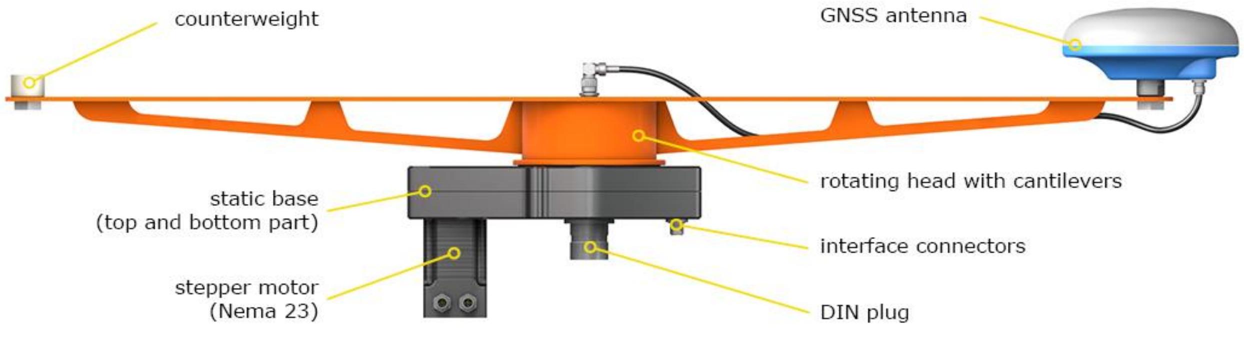

Figure 1. The rotating GNSS antenna (see Manufacturers for more information)

2.2. Principles of Synthetic Aperture Processing

The basic operating principle of the chosen synthetic aperture system is shown in Figure 2. It can be seen as a variant of a vector tracking receiver. If the receiver has a PVT solution available, the receiver predicts this solution for the next beamforming interval (e.g., duration of one second) and uses this prediction to compute replica signals. The replica signals are then correlated against the received GNSS signal from the rotating antenna. The correlation time interval is short (e.g., four milliseconds) and in this case 250 correlation values are obtained for each received GNSS satellite signal over one rotation. The rotating antenna is therefore equivalent to a phased array antenna with 250 elements.

The correlation values are collected for satellites and all code phase offsets (e.g., early, prompt and late). Then the impact of the satellite motion and the receiver clock drift and jitter is removed. The receiver clock has a nontrivial impact on the correlation values and using more stable oscillators (e.g., atomic frequency standards) considerably simplifies the receiver clock estimation efforts.

Once those effects are removed, it can be shown that the correlation values can be treated as though they were received at the same instant. Consequently, the whole theory for phased array systems can be employed. Digital beamforming and null steering techniques can be employed, allowing an update of the synthetic array weight vector per the time-varying signals’ conditions, and thus adjusting the radiation pattern of the antenna array dynamically, at each instant. It can be a maximization process, such as the maximization of the signal-to-noise ratio, or of the signal-to-interference-and-noise ratio; or it can be a minimization process, such as the minimization of an error between a model and the actual signals (Minimum Mean Square Error (MMSE) algorithm), or of the variance of the beamformer output (Linearly Constrained Minimum Variance (LCMV) algorithm or Minimum Variance Distortion-less Response (MVDR) algorithm).

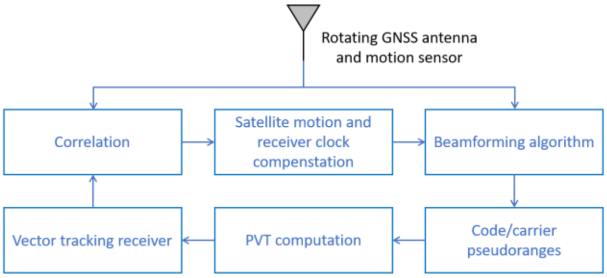
Rotating GNSS antenna and motion sensor
Correlation
Satellite motion and receiver clock compensation
Beamforming algorithm
Vector tracking receiver
PVT computation
Code/carrier pseudoranges

Figure 2. Work flow of synthetic aperture processing

The beamforming algorithm produces combined correlation values eventually exploiting the spatial diversity. Those correlation values form the basis for the generated code and carrier pseudoranges. It is important to consider distortion-less response algorithms, as they ensure that the beamforming does not introduce any biases in the code or carrier pseudoranges.

For our tests, an adaptive beamforming algorithm was selected, as shown in Figure 3. The algorithm is tailored to handle spoofing signals. Being an engineering solution, it first eliminates the LOS signals from the compensated correlation values by applying suitable Null operators. This can be done to high precision, as the DoA of the LOS signals is known. In the next step, the received signal power is estimated as a function of the DoA. This is done on a grid of elevation and azimuth values with a grid resolution of one degree. It should be noted that the raw beam width of the synthetic aperture antenna is on the order of 10 degrees, due to the selected diameter of one meter and wavelength of 19.03 centimeters.

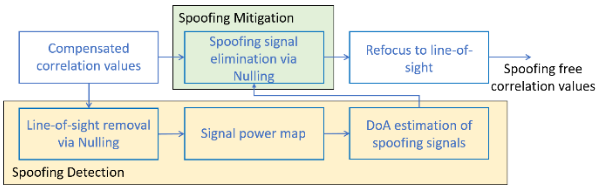
Spoofing Mitigation
Compensated correlation values
Spoofing signal elimination via Nulling
Refocus to line-of-sight
Spoofing free correlation values
Line-of-sight removal via Nulling
Signal power map
DoA estimation of spoofing signals
Spoofing Detection

Figure 3. Chosen beamforming algorithm with DoA estimation and Nulling

In the case where no spoofing signal is present (and no strong specular multipath reflection exists), the estimated received signal power (as a function of elevation and azimuth) is noise-like. In the case where a spoofing signal is present, it clearly shows up as a peak in this map (see Figure 4) and its DoA can be retrieved.

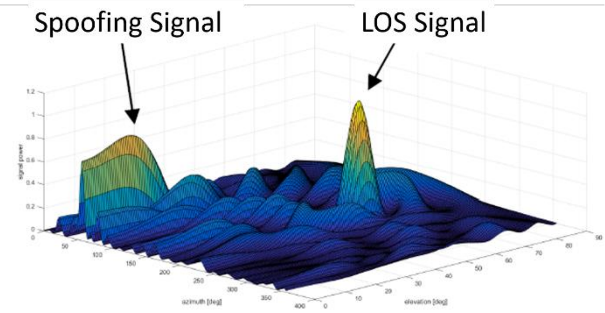

Figure 4: Signal power map including LOS and spoofing signal

The positions of the peaks are used to identify the DoA of the spoofing signals, which themselves are used to construct a Null operator to eliminate the spoofing signals from the compensated correlation values. After the spoofing signals have been eliminated, it is reasonable to assume that only the LOS is present and, by focusing the synthetic aperture antenna gain towards the LOS, optimal correlation values are obtained.

A characteristic of the chosen method is that spoofing signals are treated independently of their power. In other words, a weak spoofer is treated the same as a strong spoofer (provided the weak spoofer is detected). In contrast, an MVDR beamformer will react more adaptively on varying signal strengths.

2.3. Non-Coherent Spoofing Detection

One major drawback of the algorithm presented in section 2.2 is that it does not consider non-coherent spoofers, because they will not show up as local maximum in the signal power map of a prompt correlator. A non-coherent spoofer sends a not perfectly aligned spoofing signal compared to the true line-of-sight signal from the satellite. When not using a sophisticated time synchronization algorithm (see section 3.2) but the internal clock of the spoofer, it is likely that the spoofing signal contains an additional spoofing clock bias and a clock drift. Furthermore, it is possible that the transmitted data bits are not correctly predicted and thus the spoofing signal does not correlate with the true signal. Both cases are very likely, hence it is worth to extend the spoofing detection algorithm for non-coherent spoofers.

The effect of unknown data bits can be mitigated by squaring the correlator values. The problem of this approach is that at the antenna side a superposition of both, the true line-of-sight signal a and the spoofing signal b, arrives. As a result, the squared correlator values including a spoofing signal would lead to unwanted and unpredictable term which makes it difficult to separate the spoofing signal from noise:

$$(a+b)^2 = a^2 + 2ab + b^2 \qquad (1)$$

To avoid this effect, the influence of the line-of-sight signal is eliminated first from the correlator values (including the true data bits) using the Nulling operator. The remaining correlator values $\bar{P}_{sp}$ include the spoofing signal only:

$$\bar{P}_{sp}(t_k) = d_{sp}(t_k) a_{sp}(t_k) \exp\{2\pi i \frac{\Delta\rho_{\mu}^{sp}(t_k)}{\lambda}\} \quad (2)$$

The remaining correlator values consists of the unknown spoofing data bits d_{sp} (binary values ± 1), the amplitude of the spoofing signal a_{sp}, the micro-trajectory of the rotating antenna projected onto the unit vector pointing towards the spoofer $\Delta\rho_{\mu}^{sp}$, and the wavelength λ. By squaring equation 2, the unknown data bits of the spoofing signal are eliminated. Afterwards, the algorithm from section 2.2 can be conducted.

When the spoofing signal includes an unknown clock drift, the prompt correlator is blind to the spoofing signal and thus the spoofer cannot be detected. This clock drift generates a shift of the spoofing correlation peak in the Doppler and/or in the code/phase domain. The approach to detect this kind of spoofers is to search for the spoofer outside of the prompt correlator. To detect shifts in the code/phase domain, the signal power maps presented in section 2.2 can be calculated for different correlators (e.g., early or late correlator instead of prompt correlator). For this paper, a set of 21 correlators were selected in order to consider real time capability. To detect shifts in the Doppler domain, different Doppler shifts were introduced to the correlator values before calculating the signal power maps. With this approach, calculating signal power maps for different code phase offsets and Doppler shifts, the receiver is able to detect also non-coherent spoofers.

3. TEST DATA

3.1. Simulation of Spoofing Signals

In order to test the effect of wrong transmitted data bits of the spoofer on the software-based receiver working in synthetic aperture mode, a spoofing signal was simulated in MATLAB. The simulation combines the correlation values of a “true” line-of-sight signal with correlation values of a spoofer. The direction of both signals in terms of azimuth and elevation can be set by the user. Additionally, the software allows the simulation of noise and of an arbitrary antenna motion pattern. In this paper, the antenna motion is selected to be a circular motion with a radius of 0.5 meters and a rotating velocity of 1 rotation per second. These settings are equal to the real-world test setup presented in section 3.2.

To test unknown data bits of a spoofer, arbitrary bits are introduced to the spoofing signal before combining it with the “true” line-of-sight signal. This was conducted by multiplying the simulated correlation values of the spoofing signal with random binary numbers (+1 or -1). Even if completely random bits are not very likely in real-world scenarios, this setting is chosen to demonstrate the effect of unknown data bits on the target receiver.

3.2. Real-World Spoofing Test Setup

The spoofing setup used within this work for the real-world tests consists of an RF constellation simulator operated in a dedicated spoofing mode. The simulator is frequency synchronized via a rubidium atomic clock. Time synchronization to the true GNSS signals is achieved via a separate GNSS receiver. In general, the setup is similar to the one used in [1], but in this case a field-programmable-gate-array (FPGA) based constellation simulator has been used to generate the signals. Figure 5 shows the principal setup as a block diagram.

The spoofer calibration GNSS receiver delivers demodulated navigation data symbols. Those symbols are collected over a certain time and are then predicted for GPS C/A to allow real-time transmission of the true symbols. The Galileo spoofing was done on the pilot (E1C) only, and in this case no prediction is necessary. The spoofing mode allows for the application of position/velocity and time/time drift offsets to the true target position, velocity, and time (PVT).

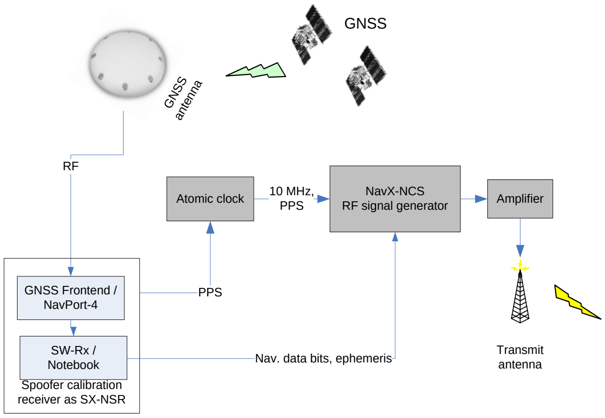
GNSS
GNSS antenna
RF
Atomic clock
10 MHz, PPS
NavX-NCS RF signal generator
Amplifier
GNSS Frontend / NavPort-4
PPS
SW-Rx / Notebook
Spoofer calibration receiver as SX-NSR
Nav. data bits, ephemeris
Transmit antenna

Figure 5: Spoofing signal generation setup (block diagram)

The setup was installed in a 19-inch rack in the laboratory with a 20-meter RF cable to the transmit antenna on the roof. The complete setup with all RF cables (signal-in-space (SIS) antenna to transmit antenna) was calibrated with a test receiver connected to the RF output of the signal simulator for the exact delay between the PPS of the rubidium clock and the PPS of the test receiver receiving the spoofing signal. The determined offset was configured in the spoofing mode setup of the RF simulator for compensation. To further compensate for the free space loss, 55 dB amplifiers were connected to the RF output to provide margin in addition to the simulator internal amplifier.

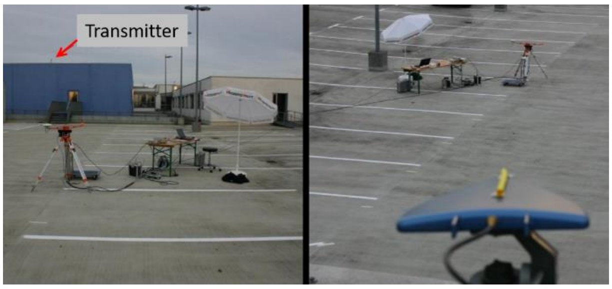

Figure 6: Test area on the parking deck with view from the receiver under attack up to the transmitter on the roof (left) and from the transmit antenna down to the test receiver (right)

The tests were performed on the IFEN premises in Poing, Germany. Respective transmission permission was granted and proper measures ensured that the spoofing signal was weak enough. The transmit antenna was installed on the roof pointing to the receivers under tests (one static and one rotating antenna receiver) placed on the parking deck. On the other side of the roof, outside the effects of the spoofing signal, a second static and a second rotating antenna receiver used as reference were installed and were running throughout the experimentations. Figure 6 shows the setup with views from and to the parking deck.

The setup was used in a test campaign lasting several days to perform the following spoofing attacks:

- Position spoofing by introducing a velocity after initial multipath spoofing to take over the receiver.
- Time spoofing by introducing a time drift after initial multipath spoofing to take over the receiver.
- Multipath spoofing without any offset to the truth position and time.

In order to test the detection algorithm of non-coherent spoofers, the test data of the time spoofing attack was used. For more details and results on the position spoofing attack and the multipath spoofing attack, as well as for detailed information on the angular resolution of the spoofing signal detection and the occurring fading effects, see [6].

4. RESULTS

4.1. Unknown Data Bits

To test the implemented algorithms for a spoofer transmitting not the correct data bits, signals arriving at a rotating antenna were simulated (see section 3.1). The simulated rotation has a circular motion frequency of 1 hertz with a radius of 0.5 metres. The correlation values were simulated with 4 milliseconds of coherent integration time, and a beamforming interval of 1 second was chosen (equals one rotation of the antenna). The satellite was simulated with an azimuth of 50°, an elevation of 75°, and at a distance of 20.000 kilometers. The spoofer was simulated with an azimuth of 100°, an

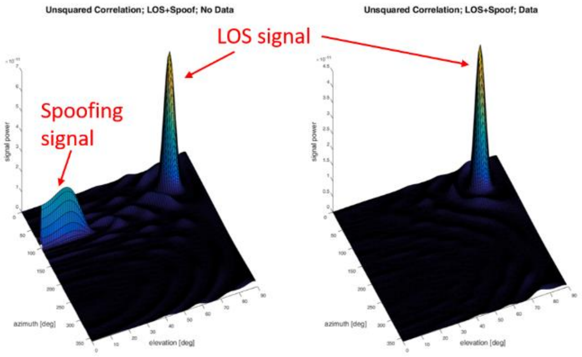

Figure 7: Signal power maps using the unsquared approach; Simulated spoofing signal transmitting the correct data bits (left) and simulated spoofing signal transmitting random data bits (right)

elevation of 10°, and with 70% of the signal power of the line-of-sight signal. The settings were chosen in a way that they are as realistic as possible and similar to the real-world spoofing attack of section 4.2.

When a spoofer transmits correct data bits (or if a pilot signal is used and no data bits are on the spoofing signal), the line-of-sight signal and the spoofing signal can be clearly separated in the signal power map using the unsquared algorithm (see Figure 7). When the spoofer transmits completely random data bits, the spoofing signal does not correlate with the replica signal at all, hence the spoofer cannot be detected by the unsquared approach (see Figure 7). Note that the more correct data bits are transmitted, the more the spoofing signal correlates with the replica signal. Completely random data bits are not very likely in real-world scenarios, but they are used in this paper to investigate the impact of unknown data bits on the signal power maps.

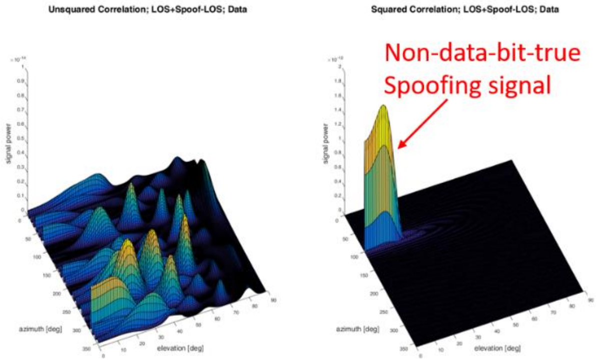

Figure 8: Signal power maps with eliminated LOS signals and a spoofing signal transmitting random data bits; No spoofing detection with the unsquared approach (left), spoofer is visible with the squared approach (right)

Figure 8 shows the signal power maps of the unsquared and the squared approach with eliminated line-of-sight signal via Nulling. This elimination must be done to enable the squared approach as described in section 2.3. The left plot shows again that the spoofer does not correlate with the replica signal using the unsquared approach. Using the squared approach, the impact of the unknown data bits of the spoofer on the correlation values can be eliminated and the spoofer is again detectable in the signal power map (see Figure 8).

4.2. Real-World Spoofing Attacks

For the time spoofing attack, the spoofer takes over the tracking loops of the receiver under attack and manipulates the receiver time by inducing a time drift in the spoofing signal. The effect of this time spoofing is shown on

- a conventional receiver with typical frequency, phase, and delay lock loops; and on a
- receiver with the synthetic aperture antenna.

The goal of this scenario was to capture the victim receiver's tracking loops and shift the receiver time more than 26.5 microseconds away. This threshold is given as an example by [5] of success for a timing attack against phasor measurement units (PMU) in electric power control systems. Figure 9 shows the receiver clock error and drift plots for the time spoofing attack. The upper plot refers to the conventional receiver and the lower one to the synthetic aperture receiver. The upper plot clearly demonstrates that it was possible to take over the control of the conventional receiver tracking loops and shift the receiver clock up to 400 microseconds away from the receiver's true clock error. For this scenario, the time spoofing started at 300 seconds with increasing time drift until the intended time drift of 1 nanosecond/second was reached and the time drift was kept constant for the whole spoofing period. The synthetic aperture receiver shown in the lower plot does not show any changes in the clock drift and remains at its true time solution.

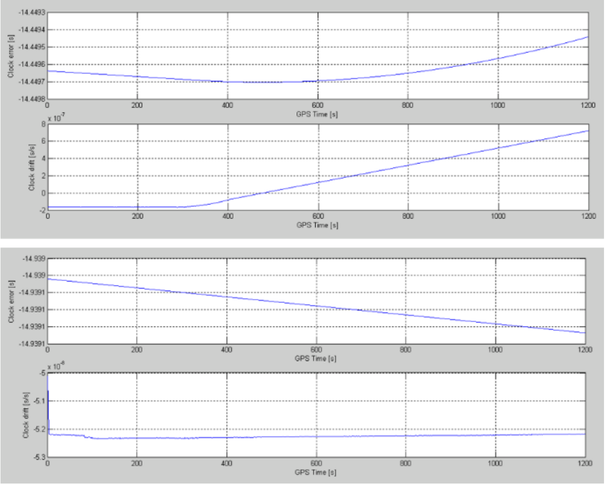

Figure 9: Upper plot shows the conventional receiver (receiver time is clearly affected by the spoofing attack); lower plot shows the rotating synthetic aperture antenna with applied spoofing mitigation techniques (clock remains almost stable)

As one can see, the time spoofing attack has a tremendous impact on conventional GNSS receivers. Nevertheless, this type of spoofing attacks cannot be detected in the prompt correlator all the time. Figure 10 shows the spoofing detection of the time spoofing attack at 2 different epochs. The upper plots show the beginning of the time spoofing attack, where the spoofing signal is perfectly aligned with the true line-of-sight signal in order to take over the tracking loops. At that time, the spoofer can be detected in the prompt correlator. Once the time spoofing attack has started, the correlation peak of the spoofing signal is drifting away in the Doppler and the code phase domain. The plot at the bottom of Figure 10 shows the spoofing detection at a time step, where the correlation peak of the spoofer is

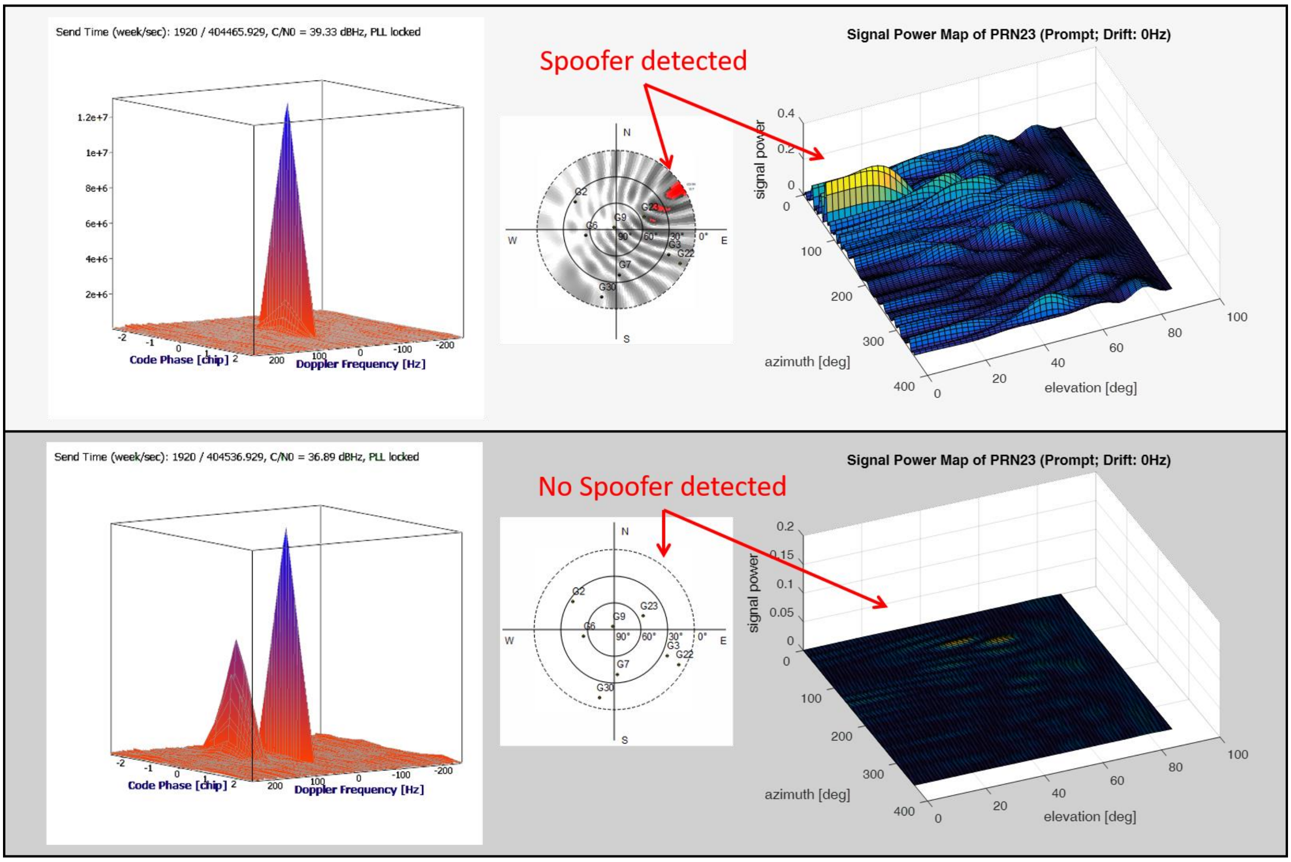

Figure 10: Results of the time spoofing attack: the upper plots show a perfectly aligned spoofing signal, which can be detected in the prompt correlator with the synthetic aperture antenna, the plot at the bottom show a spoofer with a significant offset in the Doppler and the code/phase domain, which cannot be detected in the prompt correlator

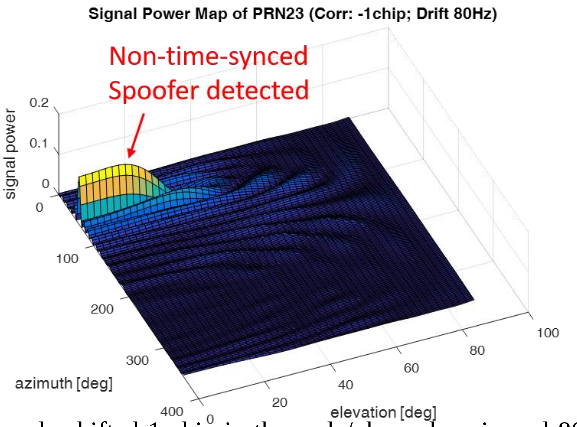

Figure 11: The signal power map of correlator shifted by 1 chip in the code phase domain, and with an introduced Doppler shift of 80Hz can detect the non-coherent spoofer

already shifted 1 chip in the code/phase domain and 80 hertz in the Doppler domain. Because of this shift, the spoofer cannot be detected in the prompt correlator anymore. This example is also similar to a spoofing attack of a non-coherent spoofer, where the correlation peak is also not aligned with the true line-of-sight signal, caused by imperfections in the spoofer clock.

With the extended spoofing detection algorithm for non-coherent spoofers presented in section 2.3, the detection of shifted spoofing signals in the Doppler and code phase domain is possible. The lower plots of Figure 10 show an example for this type of shifted spoofing signal. While the spoofer cannot be detected in the signal power map of the prompt correlator (see Figure 10), it is possible to detect this spoofer in the signal power map of another correlator (shifted by 1 chip) by introducing a drift of 80 hertz. Figure 11 shows the resulting signal power map, where this non-coherent spoofing signal is clearly visible. As a side-effect, the clock bias and the clock drift of the non-coherent spoofer can be estimated with this algorithm

5. SUMMARY AND OUTLOOK

By performing theoretical investigations, simulations, and real-world experimentation, it was demonstrated that a synthetic aperture antenna can reliably detect and mitigate even sophisticated spoofing attacks. The direction-of-arrival is a reliable metric to discriminate spoofing signals from line-of-sight signals and localize one or more spoofers with high angular resolution of two degrees.

Extensive real-world spoofing experiments have been conducted and the results obtained so far seem to confirm the theoretical expectations. Initial data processing shows that even sophisticated carrier phase based reference station data processing (e.g., for GNSS reference station networks) can be conducted during a (mitigated) spoofing attack. It can thus be expected that the synthetic aperture processing would represent an extremely robust solution for reference stations. In contrast, in all cases the conducted spoofing attacks caused the intended PVT degradation for a conventional GPS+Galileo receiver.

By performing the spoofing detection algorithm also for other correlators (not only for the prompt correlator), and by introducing different drifts and code phase offsets for the calculation of the signal power maps, the algorithms can be extended in order to detect non-coherent spoofers too. Results based on a real-world spoofing attack underline the strength of this algorithm. Additionally, results of simulations show that the squared approach is very promising to detect spoofers that transmit incorrect data bits.

6. ACKNOLEDGEMENT AND DISCLAIMER

Acknowledgement should go to Eva Bauer and Daniel Koch for their support during the measurement campaigns and data analysis for this work. All work has been performed within the Galileo Evolution Programme EGEP funded project SETI (No. EGEP-ID 89-1.11) of the European Space Agency (ESA). Some of the presented results are already an outlook for a research projected of IGASPIN called RADIAL2017. RADIAL2017 is funded by the FFG in Austria. The views expressed in this article are solely the opinions of the authors and do not reflect those of the ESA.

Manufacturers:
The rotating GNSS antenna used in these experiments was designed by Blickwinkel Design and Development, Graz, Austria, www.blickwinkel.at.

7. REFERENCES

1. Humphreys, T.E., Ledvina, B.M., Psiaki, M.L., O'Hanlon, B.W. & Kintner, P.M., Jr. (2008). 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 GNSS 2008).* Savannah, GA, pp. 2314-2325.

2. Ioannides, R.T., Pany, T. & Gibbons, G (2016). Known Vulnerabilities of Global Navigation Satellite Systems: Status, and Potential Mitigation Techniques. *Proceedings of the IEEE.* Volume 104, Issue:6, pp. 1174-1194.

3. Lin, T., Broumandan, A., Nielsen, J., O'Driscoll, C. & Lachapelle, G. (2009). Robust Beamforming for GNSS Synthetic Antenna Arrays. *Proceedings of the 22nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS 2009).* Savannah, GA, pp. 387-401.

4. Pany, T., Falk, N., Riedl, B., Stöber, C., Winkel, J. & Ranner, H.-P. (2013). AGNSS Synthetic Aperture Processing with Artificial Antenna Motion. *Proceedings of the 26th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 20138).* Nashville, TN, pp. 3163-3171.

5. Shepard, D.P., Humphreys, T.E. & Fansler, A.A. (2012). Going Up Against Time – The Power Grid's Vulnerability to GPS Spoofing Attacks. *GPS World*, August 2012, pp. 34-38.

6. Dampf, J., Pany, T., Bär, W., Winkel, J., Stöber, C., Mervart, L., Avila-Rodriguez, J.A. & Ioannides, R. (2017). Real World Spoofing Trials and Mitigation via Direction of Arrival Discrimination. *InsideGNSS*, May/June 2017, pp. 20-30.