Real World Spoofing Trials and Mitigation

via Direction of Arrival Discrimination

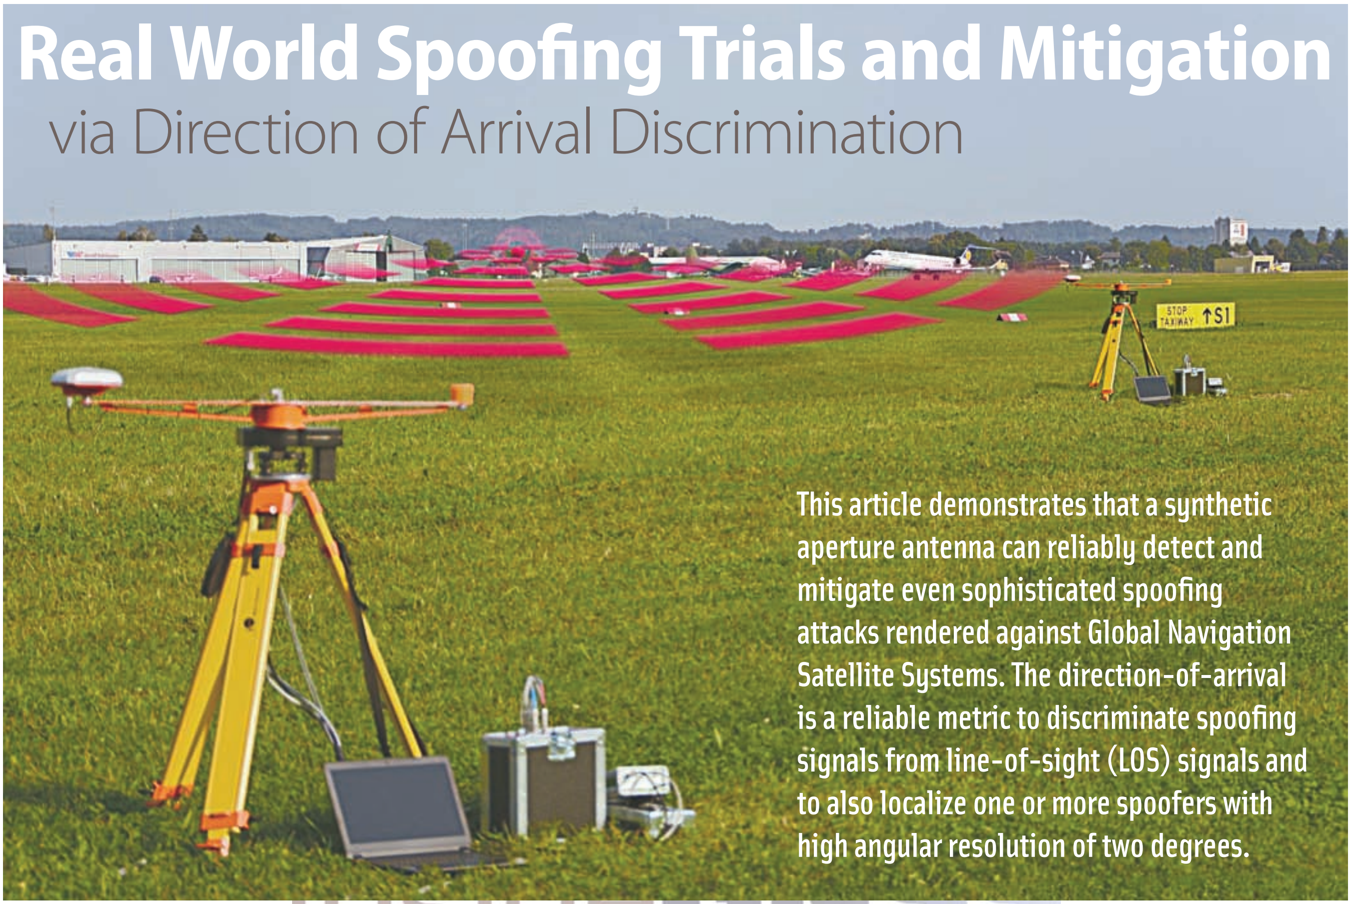

This article demonstrates that a synthetic aperture antenna can reliably detect and mitigate even sophisticated spoofing attacks rendered against Global Navigation Satellite Systems. The direction-of-arrival is a reliable metric to discriminate spoofing signals from line-of-sight (LOS) signals and to also localize one or more spoofers with high angular resolution of two degrees.

JÜRGEN DAMPF
IGASPIN GMBH

DR. THOMAS PANY
INSTITUTE OF SPACE TECHNOLOGY AND SPACE APPLICATIONS, UNIVERSITÄT DER BUNDESWEHR MÜNCHEN

WOLFGANG BÄR
IFEN GMBH

JÓN WINKEL
IFEN GMBH

LEOŠ MERVART
DEPARTMENT OF GEOMATICS, CZECH TECHNICAL UNIVERSITY

JOSÉ-ÁNGEL ÁVILA-RODRÍGUEZ
ESA/ESTEC, NOORDWIJK, THE NETHERLANDS

RIGAS IOANNIDES
ESA/ESTEC, NOORDWIJK, THE NETHERLANDS

Generation and transmission of faked GNSS signals – so-called spoofing – poses a major threat to GNSS. Spoofing has received considerable attention in recent years, but conclusive assessments or proven countermeasures have still not been found. This article summarizes experience gained while conducting real-world spoofing attacks, with one or two transmission antennas. They were conducted using a modified GNSS radio-frequency (RF) signal generator. A reliable countermeasure against spoofing is direction-of-arrival discrimination and this was realized 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 us to localize the spoofer with high accuracy. It also provides an excellent tool for studying GNSS signal reflections.

Today many applications rely on GNSS 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 R. T. Ioannides *et alia* in Additional Resources). 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.

We first describe our spoofing equipment and confirm that advanced spoofing attacks require considerable effort as a number of technical difficulties must be solved. The next section describes the rotating GNSS antenna plus receiver to detect and mitigate the spoofing attack. Finally, results from various spoofing attacks are presented and analyzed.

Degrees of Spoofer Fidelity and Predictability of the Navigation Message

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.

To better understand the predictability of the navigation message for GPS C/A and Galileo Open Service (OS), we note that the respective interface control documents specify the message structure for those data fields which are actually used for navigation. Those fields can be recorded by a dedicated receiver and can be predicted into the future. The prediction is valid as long as the content of the data fields (e.g., the Kepler elements) does not change. Changes after an upload from the ground control segment occur every few hours. There are other fields in the navigation message which are not specified in the interface control document. They are usually called spare or reserved fields. The extent to which these can be predicted has been assessed within a short experiment.

As the first and most important signal, the GPS C/A signal is considered. The GPS C/A message structure includes parity bits, telemetry, and handover words. A full-frame of the navigation message has a duration of 12.5 minutes and is subdivided into 25 pages. Each page has five subframes. It is straightforward to extrapolate the navigation data bits. Only the telemetry and handover word must be updated to reflect the current sent time, and this requires the recalculation of the parity bits. The content and definition of the reserved bits is not known. They are contained in Subframes 1, 4, and 5, with most of them in Subframe 4.

We tracked eight GPS C/A signals over 2,099 seconds starting on November 28, 2013 at 09:10:49 and logged the reserved bits. Tables 1 and 2 provide the statistics. The field "Bit" indicates the location of the bits within the subframe. The field "Cnt" is the number of occurrences of the specific pattern in that location. The field "Val" is the pattern in hexadecimal notation. The field "PRNs" lists the satellites that broadcast the listed pattern. The number of occurrences includes all satellites. First, Table 1 shows the occurrence within Subframe 1. For example, bits 91-113 showed mostly the content 0x1326fe for all PRNs over the whole duration, but for PRN9 the content 0x1726fe was also broadcast. For Subframe 5, the situation is similar (see Table 2). The bits in page 25 (=SV-Id 51) are different for two different sets of PRNs but did not change during the experiment.

The situation for Subframe 4 is more complex. This subframe has a different content for each page and this also applies for the reserved bits. The statistics for the whole duration show that most of the reserved bits repeat, but exceptions occur. For example, in the page with SV-Id 57 or 62, bits 271-292 change. If the duration of the statistics is shortened, then the reserved bits remain constant. We conclude that the GPS C/A navigation message can be predicted with a very good likelihood to guess the correct navigation data bits. Occasional changes (when ephemeris data is uploaded or changes occur in the reserved bits), however, cannot be foreseen.

As a second important case, the Galileo I/NAV message as used for E1 OS has been experimentally analyzed for broadcast status as of November 2013. The structure of the I/NAV message is more complex than that of the GPS C/A navigation message. The message period is 720 seconds (called a frame) and is subdivided into subframes with

```
Bit: 91-113, Cnt:  44, Val: 0x1726fe, PRNs:  9
Bit: 91-113, Cnt: 513, Val: 0x1326fe, PRNs:  3 6 9 16 18 19 22 27
Bit:121-144, Cnt: 557, Val: 0x5cd12e, PRNs:  3 6 9 16 18 19 22 27
Bit:151-174, Cnt: 557, Val: 0xf871fd, PRNs:  3 6 9 16 18 19 22 27
Bit:181-196, Cnt: 557, Val: 0x  1b15, PRNs:  3 6 9 16 18 19 22 27
```

Table 1 Statistics of Reserved Bits in GPS NAV Subframe 1

```
SV-Id: 51, Bit:271-276, Cnt:   4, Val: 0x     0, PRNs:  3 6 9 27
SV-Id: 51, Bit:271-276, Cnt:   4, Val: 0x     5, PRNs:  16 18 19 22
SV-Id: 51, Bit:277-292, Cnt:   4, Val: 0x     0, PRNs:  3 6 9 27
SV-Id: 51, Bit:277-292, Cnt:   4, Val: 0x  5555, PRNs:  16 18 19 22
```

Table 2 Statistics of Reserved Bits in GPS NAV Subframe 5

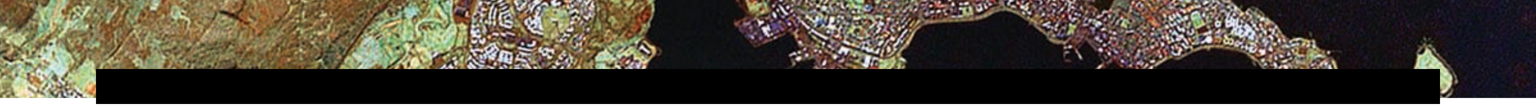

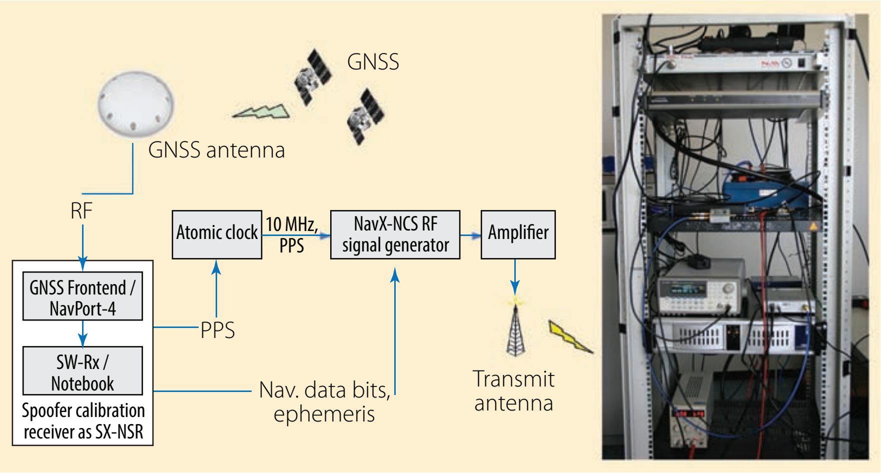
GNSS
GNSS antenna
RF
GNSS Frontend / NavPort-4
SW-Rx / Notebook
Spoofer calibration receiver as SX-NSR
Atomic clock
10 MHz, PPS
NavX-NCS RF signal generator
Amplifier
PPS
Nav. data bits, ephemeris
Transmit antenna

FIGURE 1 Spoofing signal generation setup (block diagram) and 19 inch rack installation

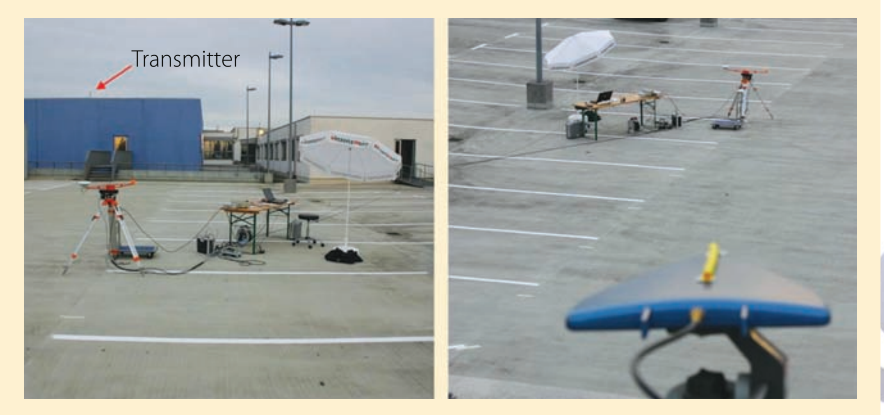

FIGURE 2 Test area on the parking deck with view from the receiver under attack up to the transmitter on the roof (left) and from the transmit antenna down to the test receivers (right)

```
Bit:  8- 31, Cnt: 1960, Val: 0x555555, PRNs:  11 12 19 20
Bit: 32- 55, Cnt: 1960, Val: 0x555555, PRNs:  11 12 19 20
Bit: 56- 79, Cnt: 1960, Val: 0x555555, PRNs:  11 12 19 20
Bit: 78- 93, Cnt: 1960, Val: 0x  5555, PRNs:  11 12 19 20
```

Table 3 Statistics of Word 0 in Galileo E1B I/NAV (November 2013)

30-second duration and pages with two-second duration. Each page has an even and odd part, with a duration of one second. The overhead of the messages includes the sync pattern, tail bits, and a cyclic-redundancy-check (CRC). The message is convolutional encoded and interleaved. The odd page part contains reserved bits, spare bits, and search-and-rescue (SAR) bits, which are potentially not predictable. Furthermore, the data content is not defined for Word type 0. The presence of alert pages complicates the structure and predictability.

Using exactly the same experiment as before, the Galileo I/NAV message content was analyzed. Results are obviously very preliminary as the data content of I/NAV may and will most likely change during the later phase of Galileo operation. Table 3 shows that Word 0 was broadcast frequently but always had the same data content of alternating zeros and ones. Furthermore, we found that the 40 bits of the "Reserved 1" field of the odd page contained 0x0, the 24 bits (22-bit SAR plus 2 spare bits) contained 0xaaaaa9, and the 8 bits of the "Reserved 2" field contained 0xfd. It should also be noted that during the experiment, no alert page was broadcast.

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 by T. E. Humphreys *et alia*, but in our case a field-programmable-gate-array (FPGA) based constellation simulator has been used to generate the signals. Figure 1 shows the principal setup as a block diagram and its realization.

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).

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 decibel amplifiers were connected to the RF output to provide margin in addition to the simulator internal amplifier.

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 2 shows the setup

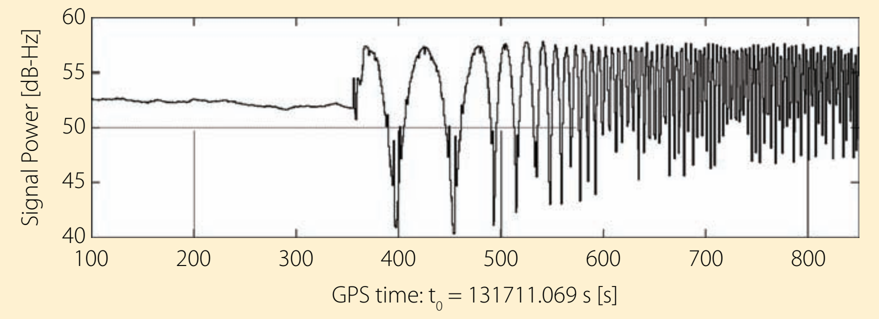

FIGURE 3 Observed fading effects (between spoofing and satellite signals) on the static reference receiver on GPS C/A PRN1 and multipath spoofing starting after 355 seconds

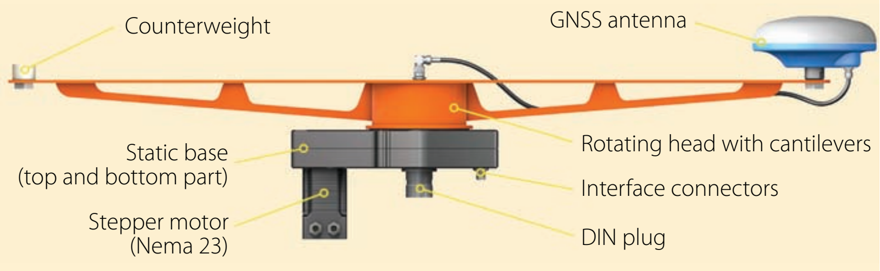

FIGURE 4 The rotating GNSS antenna seen from the side (see Manufacturers for more information)

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.

Furthermore, with an additional transmit antenna, the azimuth angular resolution for spoofer signal detection was determined. Therefore, the spoofing signal was split at the RF level and transmitted over two transmit antennas with an angle of around 30 degrees to the test receiver. While the first transmit antenna was kept static, the other transmit antenna was gradually repositioned and moved step by step nearer to the static first transmitter until they were both side by side.

At the beginning of each measurement day a brief calibration and verification of the test setup was performed to achieve the same signal propagation conditions for the whole measurement campaign. This was necessary due to severe fading effects. Fading effects in spoofing occur among all involved signals. They include the direct satellite and spoofing signal as well as reflections of these signals from the ground. These effects occurred especially on the static reference receiver with significant C/N_0 variations for nearly identical placement of the antennas. By moving the reference receiver antenna slightly, these effects have been increased or decreased. **Figure 3** shows the fading effect for the multipath spoofing scenario on the static reference receiver. Due to misalignment of the phase of the spoofing signal with respect to the direct LOS signal, signal reflections occur, especially for low elevation transmitted signals. Constructive and destructive signal effects can be observed as shown in the figure and are caused by the increasing drift of the spoofer oscillator. In order to reduce the fading effects (between the spoofing signal and its ground reflection) a ground plane with approximately 40-centimeter diameter was placed below the receiving antenna. This reduced the impact of ground multipath from the spoofing signal and decreased the fading amplitude; overall the experiments were then more repeatable.

Spoofing Mitigation Via Direction of Arrival 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 R. T. Ioannides *et alia*). 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 us 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 T. Lin *et alia* and the work by T. Pany *et alia* investigated several signal processing options for synthetic aperture antennas.

The work presented here uses a rotating GNSS antenna (similar to T. Pany *et alia*), but with an updated mechanical design rendering it water and ice proof (see **Figure 4**). 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.

Operation Principle of a Synthetic Aperture Antenna

The basic operating principle of the chosen synthetic aperture system is shown in **Figure 5**. It can be viewed as a vari-

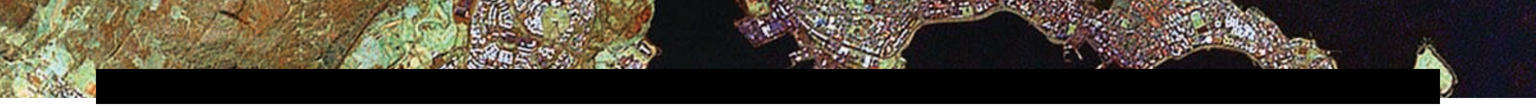

ant of a vector tracking receiver. If the receiver has a PVT solution available, the receiver predicts this solution for the next beamforming interval (e.g., duration of one second) and uses this prediction to compute replica signals. The replica signals are then correlated against the received GNSS 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).

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 6**. 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.

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 later sections for real-world data) and its DoA can be retrieved.

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. Furthermore, the implemented algorithms all require either pilot signals or known navigation data bits. Estimation of unknown navigation data symbols (or bits) for the LOS or for the spoofer signals is currently not considered.

Multipath Spoofing Attack

Spoofing detection and mitigation experiments were performed on the parking deck of the IFEN premises. All spoofing experiments were performed on L1 GPS C/A and Galileo E1 OS pilot. As a first case, we analyze a multipath spoofing scenario, where the spoofer

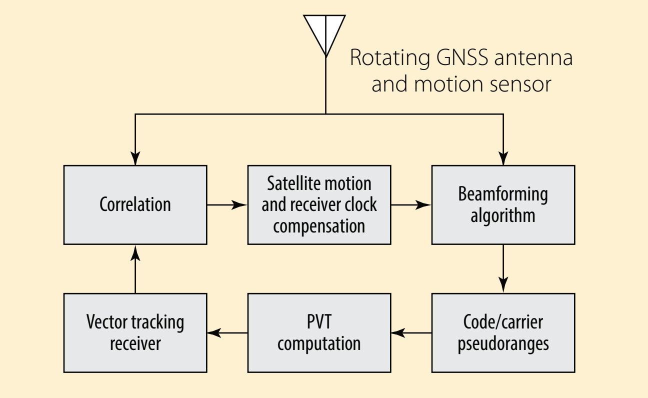
Rotating GNSS antenna and motion sensor
Correlation
Satellite motion and receiver clock compensation
Beamforming algorithm
Vector tracking receiver
PVT computation
Code/carrier pseudoranges

FIGURE 5 Block diagram of synthetic aperture processing

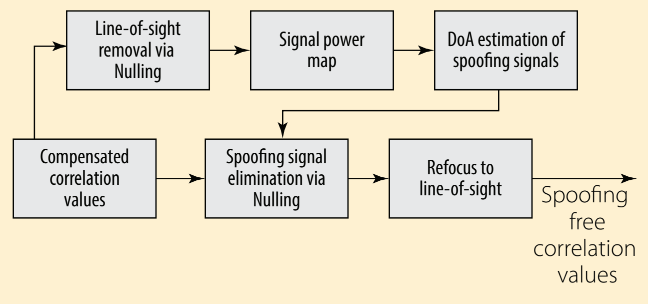
Line-of-sight removal via Nulling
Signal power map
DoA estimation of spoofing signals
Compensated correlation values
Spoofing signal elimination via Nulling
Refocus to line-of-sight
Spoofing free correlation values

FIGURE 6 Chosen beamforming algorithm with DoA estimation and Nulling

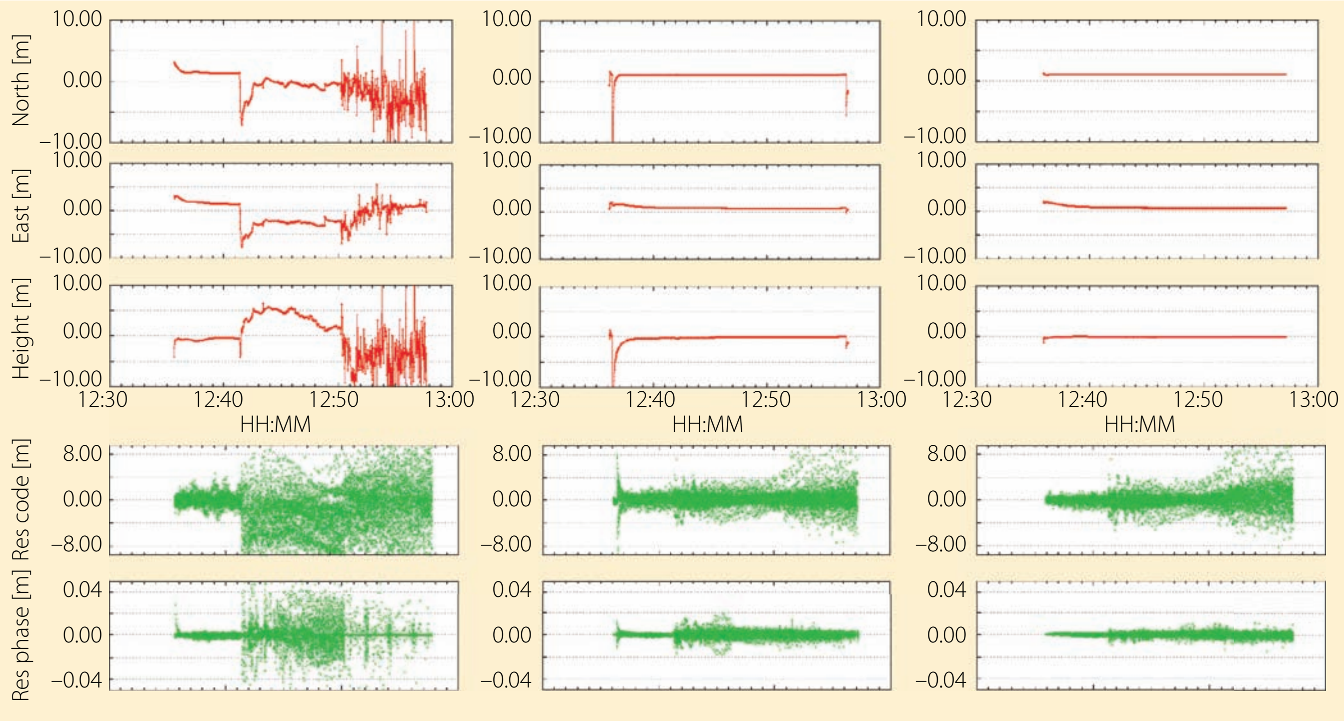

FIGURE 7 Position solution (red) and code and phase residuals (green) for GPS and Galileo during the multipath spoofing starting at 12:43; left column shows the conventional receiver; middle column shows synthetic aperture processing without Nulling; right column shows synthetic aperture processing with Nulling

transmits an identical signal as the signal-in-space without any position or time deviation. The effect of multipath 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 degrade the PVT solution of the receiver by the introduced multipath effects. **Figure 7** shows a baseline processing between the rover receiver and the reference station receiver. The baseline processing made use of a dedicated analysis tool for static baseline processing. This tool calculates the position solution and the code and carrier residuals. The spoofing attack starts at around 12:43. The lower two plots of the conventional receiver (left column) immediately show increased phase and code residuals with the start of the multipath spoofing (Note: Phase residuals above 0.05 meters or code residuals above 10 meters were directly eliminated by the processing and are not shown here). Over the complete time period, the measurements show cycle slips for almost every epoch and signal. In effect, the data from the static receiver is completely useless during the multipath spoofing attack and the position solution also degrades. In contrast, the code and phase residuals of the receiver with synthetic aperture antenna (with and without Nulling) also slightly degrade after the start of the spoofing attack, but the phase residuals still are in a reasonable range and the observations show almost no cycle slips in the processing. Code and carrier measurements have a degraded accuracy but can still be used for positioning. We also note a slightly higher accuracy if Nulling is applied (compared to synthetic aperture processing without Nulling) and conclude that focusing the antenna beam towards the satellite already eliminates the bulk of the spoofing signal energy.

Position Spoofing Attack

As a second case, we analyze a position spoofing scenario where the spoofer takes over the tracking loops of the receiver under attack and moves the position solution eastwards. The effect of position 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 position solution eastwards. **Figure 8** shows two position scatter plots, with the left plot referring to the conventional receiver and the right one to the synthetic aperture receiver. The left plot clearly demonstrates that it was possible to take over the control of the conventional receiver tracking loops and shift the position over 1.5 kilometers away from the receiver's true position, with "a" referring to the true position when tracking the LOS without spoofing. The second circle labeled as "b" refers to the start of the spoofing attack and it shows increased position residuals indicated symbolically by a larger circle diameter. It is expected that this increased variance is caused by signal fading effects due to overlapping of the true GNSS signal and falsified direct and surface reflected spoofing signals during signal propagation. The introduced position drift was stopped after about 1.5 kilometers offset at label "c." The synthetic aper-

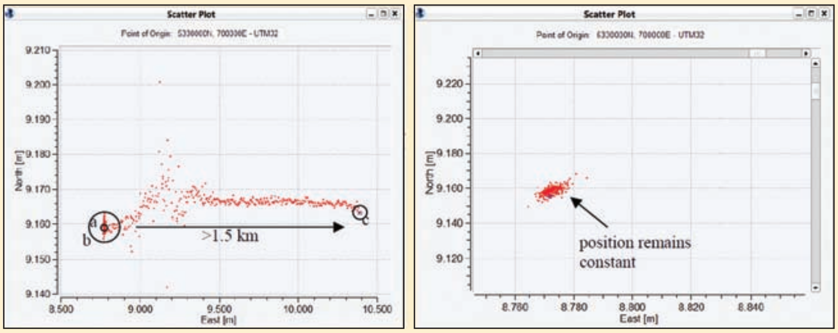

FIGURE 8 Scatter plot of the position spoofing scenario driving the position eastwards with 4 m/s over 1.5 km; left plot shows the standard tracking with a static antenna and the right plot shows the rotating synthetic aperture antenna with applied spoofing mitigation techniques

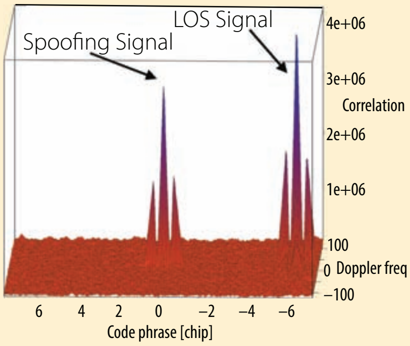

FIGURE 9 Doppler-delay map of the conventional receiver under a position spoofing attack, showing a Galileo E1 pilot signal; spoofed PVT with a constant position offset

ture receiver shown on the right of Figure 8 detects and mitigates the spoofing attack and prevents the victim receiver from locking onto the falsified signal. The synthetic aperture receiver remains at the true PVT solution, as indicated in the right plot.

The initial PVT of the spoofing signal coincides with the attacked receiver's PVT solution to take over the tracking loops smoothly, which means that the correlation function of the spoofer is exactly located in the LOS signal correlation function. **Figure 9** shows the multicorrelator output of the conventional receiver in the middle of the spoofing attack. At that instant we have a constant position displacement of several hundred meters to the true PVT. Figure 9 further verifies that a spoofing signal is actually present, which results in clearly separated correlation functions in the code phase direction.

The above described beamforming method (spoofing detection and Nulling) allows us to estimate the signal power coming from a certain DoA by projecting the received signal onto the expected phase signature and integration over the beamforming interval. Based on this, signal power maps spanned over azimuth and elevation can be derived (see **Figure 10**). The upper plots show the received signal power with an adjustable greyscale in a typical satellite sky plot, while in this case +10 decibels to LOS relates to black and –25 decibels to LOS relates to white. The left plots in this figure correspond to spoofing detection where a spoofing signal is still present and the right plots show the same signal after

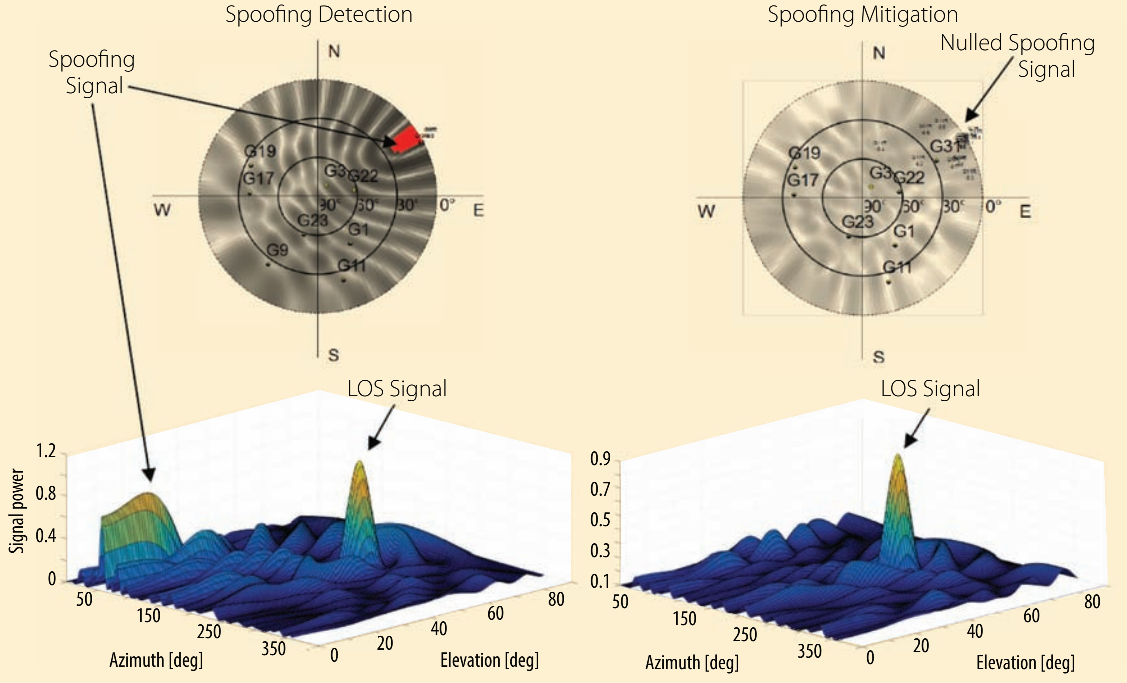

FIGURE 10 Normalized signal power plotted over DoA during a spoofing attack showing the LOS signal and the spoofing signal at about an azimuth=56° and elevation=20° in the left plot and to the right the same map after Nulling the spoofing signal; maps based on GPS C/A PRN 23

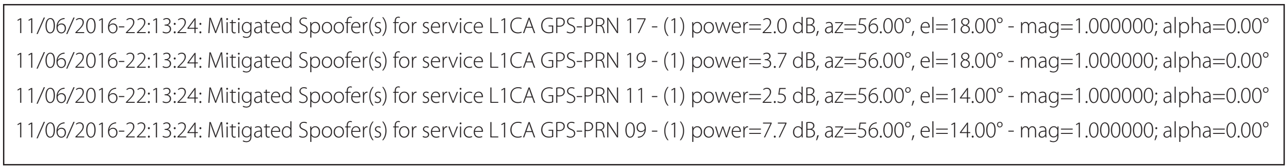

```
11/06/2016-22:13:24: Mitigated Spoofer(s) for service L1CA GPS-PRN 17 - (1) power=2.0 dB, az=56.00°, el=18.00° - mag=1.000000; alpha=0.00°
11/06/2016-22:13:24: Mitigated Spoofer(s) for service L1CA GPS-PRN 19 - (1) power=3.7 dB, az=56.00°, el=18.00° - mag=1.000000; alpha=0.00°
11/06/2016-22:13:24: Mitigated Spoofer(s) for service L1CA GPS-PRN 11 - (1) power=2.5 dB, az=56.00°, el=14.00° - mag=1.000000; alpha=0.00°
11/06/2016-22:13:24: Mitigated Spoofer(s) for service L1CA GPS-PRN 09 - (1) power=7.7 dB, az=56.00°, el=14.00° - mag=1.000000; alpha=0.00°
```

FIGURE 11 Text log of the spoofer detection and mitigation output during processing showing the time, GNSS system/service/PRN, power referred to LOS, estimated azimuth and elevation as well as applied phase correction parameters

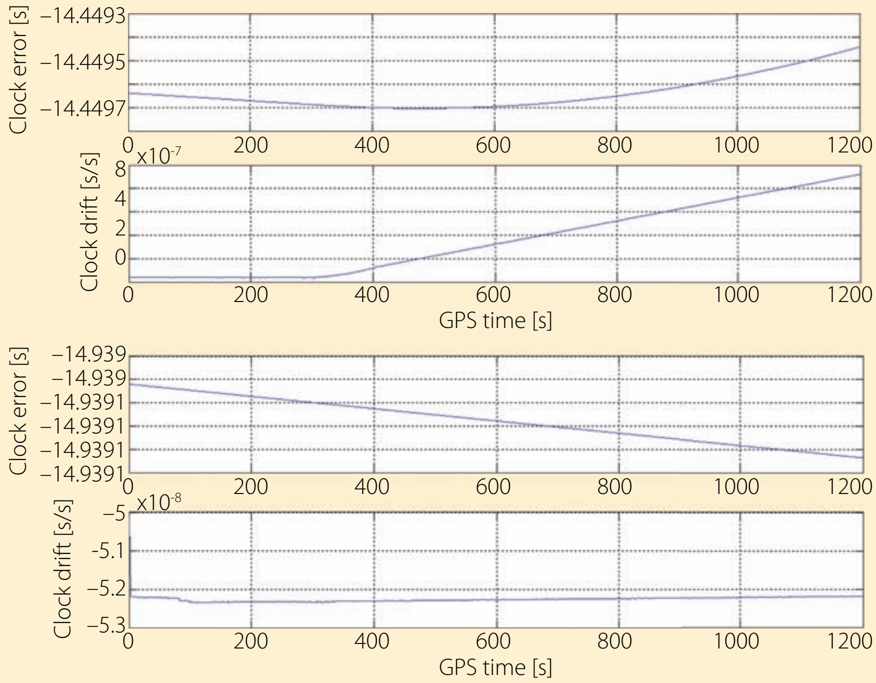

FIGURE 12 Analysis plot of the computed receiver time in single point positioning mode; upper plot shows the conventional receiver which clearly shows the introduced time drift resulting in change of receiver time; lower plot shows the rotating synthetic aperture antenna with applied spoofing mitigation techniques, which shows almost stable clock drift and no changes in receiver time computation

elimination (Nulling) of the spoofer. The left lower plot clearly shows two peaks, with the right one corresponding to azimuth and elevation of the LOS signal, and the left one to the spoofer location (indicated in red in the corresponding sky plot). The LOS signal is eliminated via Nulling in the upper plots but retained in the lower plots.

If the spoofing signal exhibits a certain threshold compared to the LOS signal, it is decided that a spoofing signal is present. In such a case the exceeding threshold is marked red in the sky plot as shown in the upper left plot of Figure 10 and azimuth and elevation angle are estimated to remove the spoofing signal by placing a Null into this spatial direction. The spoofing elimination result is shown in the right plot of the same figure, where only the LOS component remains. All spoofing detection information is written in real-time to a file and to the status window of the software receiver, as seen in **Figure 11**. The processing of different PRNs results in virtually identical DoAs as all originate from the same spoofing antenna.

Time Spoofing Attack

For the third case we analyze a time spoofing scenario, where 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 D.P. Shepherd *et alia* of success for a timing attack against phasor measurement units (PMU) in electric power control systems. **Figure 12** 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.

Multiple Spoofers

A more sophisticated spoofing attack may involve multiple transmission antennas. To test the ability of the rotating antenna to detect and mitigate transmissions from multiple spoofers, the same spoofing signal was distributed to two antennas via an RF splitter. The transmission antennas were located at 49 degrees and 68 degrees azimuth and both at about five degrees elevation. The synthetic aperture antenna was able to process this scenario and output the estimated signal power as a function of elevation and azimuth. An example plot is shown in **Figure 13**. It is very important to note that it is quite difficult to visualize signal power from multiple sources if those sources have

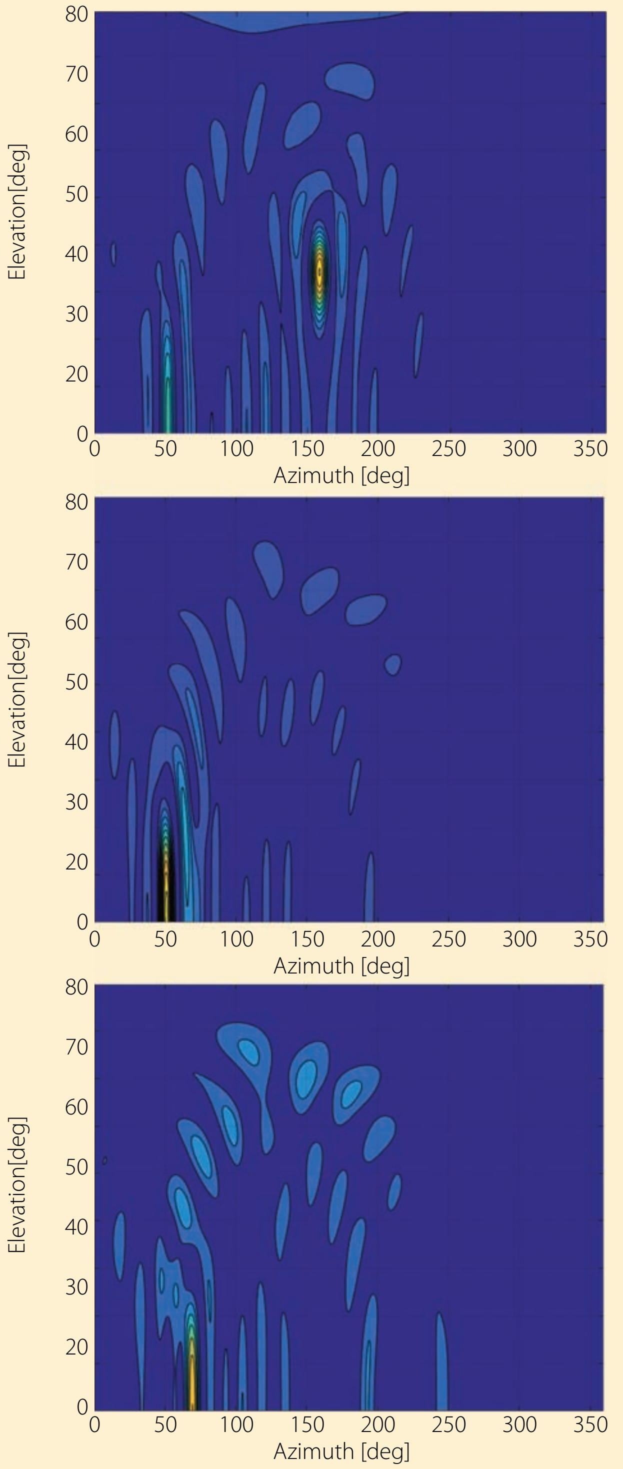

FIGURE 13 Signal power map for GPS PRN 27 under a dual antenna spoofing attack; upper: raw power; middle: after Nulling the LOS signal; lower: after Nulling the LOS and the strongest spoofing signal

Section No - Spoofer	Estimated Azimuth	True Azimuth
1 - Static Spoofer	69.9 deg	67.9 deg
1 - Moved Spoofer	50.6 deg	48.9 deg
2 - Static Spoofer	70.1 deg	67.9 deg
2 - Moved Spoofer	59.4 deg	58.7 deg
4 - Static Spoofer	-	67.9 deg
4 - Moved Spoofer	66.2 deg	67.6 deg

Table 4. Estimated and truth direction of arrival estimate,

significantly different powers. For example, if the combined signal power is plotted (upper plot of Figure 13), only the LOS signal from the satellite is clearly visible. After elimination (via Nulling), the stronger spoofing signal at 49 degrees azimuth is visible (lower left plot) and after further Nulling of the 49 degree spoofer, the spoofing signal at 68 degrees becomes visible (lower right plot). This demonstrates very well the ability of the synthetic aperture antenna to discriminate, localize, and eliminate multiple spoofing signals.

To verify the ability of the antenna to separate two spoofing signals, Figure 14 shows the estimated azimuth for both spoofers during a stepwise reduction of the azimuth difference between the spoofers. Each blue dot in the plot corresponds to a detected spoofing signal. It can be seen that both spoofers can be separated and detected well within the first two sections. The spoofer with the lower signal power is at about 70 degrees and the spoofers with the higher signal powers at about 50 and 59 degrees. Post-processing analysis showed that the 70 degree spoofer seems to appear significantly weaker compared to the moving spoofer, which is assumed to be caused from destructive signal multipath effects. From Section 3 on, it seems that the azimuth difference is too low to separate them via the DoA estimation algorithm used. It might be that the weaker spoofer is partially suppressed when applying the nulling and thus becomes invisible because the signal power drops below the noise floor. Nevertheless the azimuth of the strong spoofing signal is still reliably detected.

An interesting effect is marked in yellow in Figure 14. During the regions when one spoofing antenna is actually moved, the algorithm tends to detect the static spoofer. It is assumed that this effect comes from an improperly adjusted spoofing antenna during the movement, which leads to a significantly lower signal power making the algorithms briefly able to detect the static antenna at the beginning of Sections 3, 4, and 5.

Table 4 lists the estimated and true azimuth angles for all sections. All reference azimuth angles have been calculated from surveyed signal reception and transmission points. We conclude that the azimuth of the spoofer can be determined with an accuracy of around two degrees. Improvements in the signal processing (super resolution methods) and better

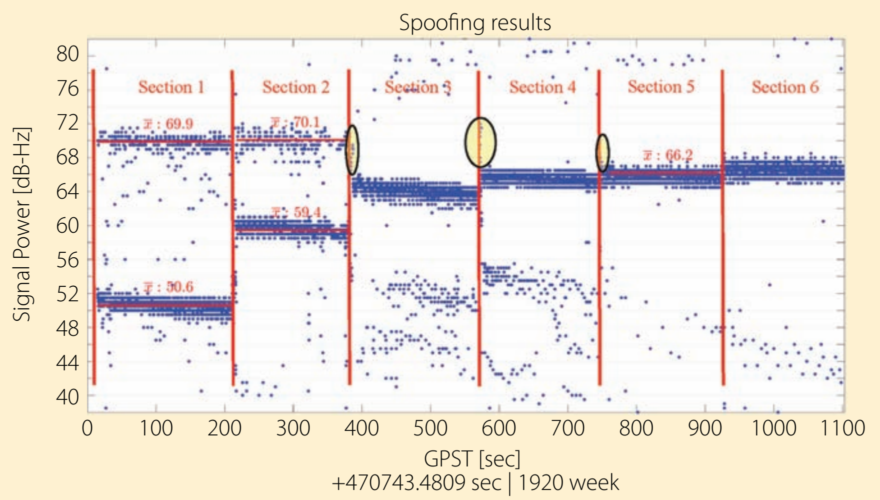

FIGURE 14 Detected azimuth over time for a static weaker spoofer and a stronger stepwise moving spoofer. Each blue dot corresponds to a detected spoofing signal

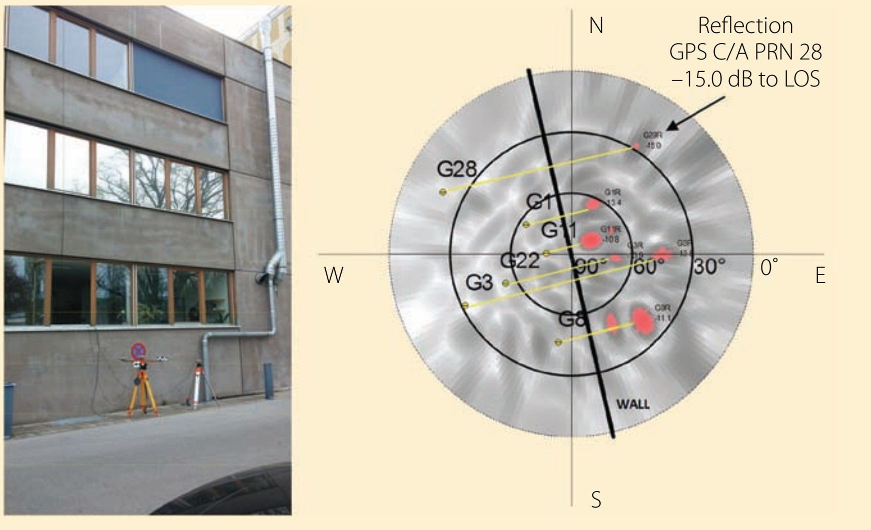

FIGURE 15 DoA based signal power map showing reflections of several satellites on a building wall

understanding of the signal propagation process may further increase the accuracy and are the subjects of ongoing research.

Multipath Analysis

Signal reflections on a wall, for example, can be seen as weaker and delayed spoofing signals. The DoA estimation algorithm was tested by detecting the DoA of multipath signals. Therefore, the rotating antenna was placed beside a building wall in order to detect multipath signal reflections. GPS C/A signals were analyzed and **Figure 15** shows the outcome of the verification experiment. The building wall is shown as a black line from (nearly) north to south within the sky plot. All tracked satellites which are visible in the sky plot are located in the west because the others are blocked by the wall. Within the signal power map, the LOS signal contribution was eliminated via Nulling. After Nulling, Figure 15 shows six clearly visible reflections. Due to the simple geometry, each reflection can be assigned to a GNSS satellite. The plot also shows the estimated signal strength in decibels with respect to the LOS signal and a yellow line outlines the corresponding satellite. No further analysis has been performed by the authors, but it is obvious that the synthetic aperture antenna provides a unique tool to analyze GNSS signal reflections.

Summary and Outlook

By performing theoretical investigations, simulations, and real-world experimentation, we 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 LOS signals and also 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.

Further sophistication of the synthetic aperture processing should in our view include methods to constrain the receiver clock during a spoofing attack and methods to handle spoofing signals with a broadcast message being different from the true message. The synthetic aperture antenna can also be used to study GNSS signal reflections as it can reliably estimate the DoA of multipath signals.

Acknowledgments 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). The views expressed in this article are solely the opinions of the authors and do not reflect those of the ESA.

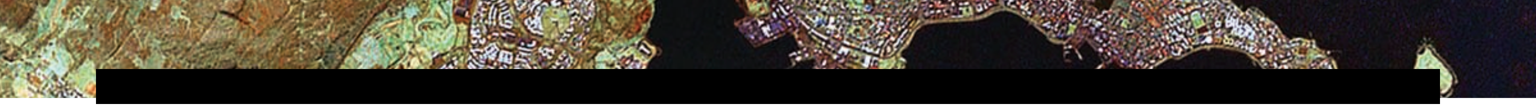

Manufacturers

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

Additional Resources

[1] Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O'Hanlon, B. W., and Kintner, P. M., Jr., "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, September 2008, pp. 2314-2325

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

[3] Lin, T., Broumandan, A., Nielsen, J., O'Driscoll, C., and Lachapelle, G., "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, September 2009, pp. 387-401

[4] Pany, T., Falk, N., Riedl, B., Stöber, C., Winkel, J., and Ranner, H.-P., "GNSS 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+ 2013),* Nashville, TN, September 2013, pp. 3163-3171

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

Authors

Jürgen Dampf received his M.Sc. in aviation engineering at the University of Applied Sciences in Graz. Since 2012 he has worked for IFEN GmbH as a system engineer emphasizing on GNSS reflectometry, GNSS/INS integration and embedded software receivers. He joined IGASPIN GmbH in 2016 as System Engineer and later became CTO, while his GNSS activities focus on spoofing signal generation, spoofing detection, mitigation and localization and GNSS signal processing techniques. He works as a lecturer at the University of Applied Sciences Graz and is a PhD candidate at the TU Graz, where he focuses on direct position estimation (DPE) and bayesian filtering techniques.

Prof. Thomas Pany is with the Universität der Bundeswehr München at the faculty of aerospace engineering where he teaches satellite navigation. His research includes all aspects of navigation ranging from deep space navigation over new algorithms and assembly code optimization. Currently he focuses on GNSS signal processing for Galileo second generation, GNSS receiver design and GNSS/INS/LiDAR/camera fusion. To support this activities, he is developing a modular GNSS test bed for advanced navigation research. Previously he worked for IFEN GmbH and IGASPIN GmbH and is the architect of the ipexSR and SX3 software receiver. He has around 200 publications including patents and one monography.

Wolfgang Bär has been head of mobile solutions department at IFEN GmbH since 2014. He joined IFEN in 2006 as a system engineer. Previously he was a research associate at the Institute for Geoinformatics and Remote Sensing of the University of Osnabrueck where he received his Ph.D.

Dr. Jón Ó. Winkel has been the head of receiver technology at IFEN GmbH since 2001. He studied physics at universities in Hamburg and Regensburg and received a Ph.D. (Dr.-Ing.) from the University FAF Munich, where his studies focused on GNSS modeling and simulations.

Leoš Mervart received his first Ph.D. from the Astronomical Institute, University of Bern and his second Ph.D. from the TU Prague. In 2002 he was appointed professor of geodesy at the TU Prague where he currently leads the Department of Geomatics. Mervart is working with GPS Solutions Inc. on the development of the RTNet software and with BKG on the development of the BKG Ntrip Client (BNC) software.

José-Ángel Ávila-Rodríguez is is GNSS Evolutions Signal and Security Principal Engineer at the GNSS Evolution Program and Strategy Division, Strategy & Program Department of the European Space Agency. Between 2003 and 2010 he was research associate at the Institute of Geodesy and Navigation at the University of the Federal Armed Forces Munich. Ávila-Rodríguez studied at the Technical Universities of Vienna, Austria, and Madrid, Spain, where he received a Master's degree in electrical engineering. He received a Master's degree in economics from the Spanish UNED University and a Ph.D. in aerospace engineering in signal design from the University of the Federal Armed Forces Munich. During his career, he has been a key contributor to the international interoperability and compatibility efforts leading to the Galileo signal plan, supporting the Galileo program in numerous working groups of the European Space Agency, the European Commission, and the Galileo Joint Undertaking. He studied at the Technical Universities of Madrid, Spain, and Vienna, Austria, and has a Ph.D. in signal design and M.S. in electrical engineering. His major areas of interest include the Galileo signal structure, GNSS receiver design and performance, and Galileo codes. He was recipient of the 2008 Parkinson Award and the 2009 ION Early Achievement Award, both from the U.S. Institute of Navigation.

Dr. Rigas Ioannides works at the TEC-ETN section in the RF Payload Systems Division at ESA-ESTEC in support of radionavigation activities and the Galileo project. His main research interests include GNSS signal design, signal processing techniques for stand-alone and integrated GNSS architectures, authentication and anti-jamming techniques at system and user level for GNSS applications, and GNSS integrity concepts. Ioannides holds a Ph.D. in trans-ionospheric propagation effects on GNSS signals, and an M.Sc. degree in communications and real-time electronic systems from the University of Bradford.

Em. Univ.-Prof. Dr.-Ing. habil. Dr. h.c. Guenter W. Hein is Professor Emeritus of Excellence at the University FAF Munich. He was ESA Head of EGNOS & GNSS Evolution Programme Dept. between 2008 and 2014, in charge of development of the 2nd generation of EGNOS and Galileo. Prof. Hein is still organizing the ESA/JRC International Summerschool on GNSS. He is the founder of the annual Munich Satellite Navigation Summit. Prof. Hein has more than 300 scientific and technical papers published, carried out more than 200 research projects and educated more than 70 Ph. D.´s. In 2002 he received the prestigious Johannes Kepler Award for *"sustained and significant contributions to satellite navigation"* of the US Institute of Navigation, the highest worldwide award in navigation given only to one individual each year. G. Hein in 2011 became a Fellow of the US ION. The Technical University of Prague honored his achievements in satellite navigation with a *Doctor honoris causa* in Jan. 2013. He has been a member of the Executive Board of Munich Aerospace since 2016. IG