Executive Summary Report

Authors: Dr. Robert Bensch, Ulrich Mittmann, Luka Sachße, Medeea Horvat,
Dr. Patrick Henkel, ANavS GmbH
Dr. Jan Wendel, Airbus Defence and Space GmbH

Date: 25.07.2023

Issue: 1.0

Table of Contents

1 Introduction

Over the past years the thread of intentional GNSS jamming and spoofing, aiming at falsifying the computed Position, Velocity and Time (PVT), has significantly increased since nowadays virtually anyone is able to perform GNSS attacks with low budget and open source code. Since today many sectors, including critical infrastructure, rely on Global Navigation Satellite Systems (GNSS) the necessity and relevance is given for the development of anti-spoofing techniques that allow to protect against GNSS attacks.

In this activity state-of-the-art and emerging technologies that can be used for PVT assurance at the receiver level have been consolidated and a flexible anti-spoofing testbed has been developed (Figure 1-1) that can cope with different sensor configurations, emulate/integrate multiple technologies and includes simulators for sensor data, network aiding ranges and various spoofing attacks. An innovative anti-spoofing sensor fusion based on machine learning (ML) techniques has been developed and integrated into the testbed that allows to detect and mitigate diverse spoofing attacks.

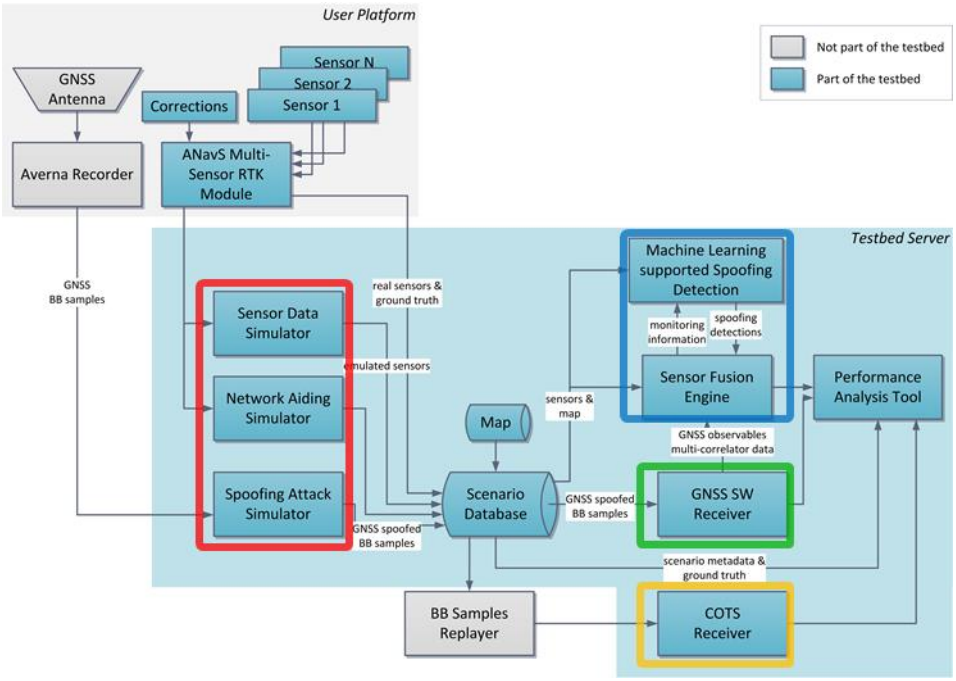

Figure 1-1: High-level Anti-Spoofing Testbed Architecture. Data simulations modules (red), User segment anti-spoofing techniques: ML-supported Sensor Fusion (blue), GNSS SW Receiver (green) and Custom-off-the-shelf (COTS) Receiver (yellow).

A data collection campaign has been conducted using a test vehicle equipped with an integrated sensor platform (ISP) from ANavS GmbH and a GNSS baseband sample recorder, provided by Airbus. Real sensor measurement, such as GNSS baseband samples, IMU, wheel odometry, camera and Lidar data have been recorded. GNSS baseband samples represent the raw radio

frequency (RF) signals that are received by a GNSS receiver and are used in this study to inject spoofing. The testbed and testbed modules have been implemented and validated and the machine learning based techniques have been trained followed by a performance evaluation covering different use cases and anti-spoofing techniques. The use cases comprise a static, low dynamic and high dynamic user, the anti-spoofing techniques compared are classical anti-spoofing techniques (implemented in the Airbus Software Defined Radio (SDR)), the developed machine learning supported sensor fusion and the anti-spoofing capabilities of a commercial-off-the-shelf (COTS) GNSS receiver.

2 Innovative Anti-Spoofing Testbed and Sensor Fusion

2.1 Anti-Spoofing Testbed

The developed anti-spoofing testbed with the architecture shown in Figure 1-1 offers a flexible framework for testing and developing anti-spoofing techniques and can cope with different sensor configurations and emulate/integrate different technologies, grouped in Figure 2-1.

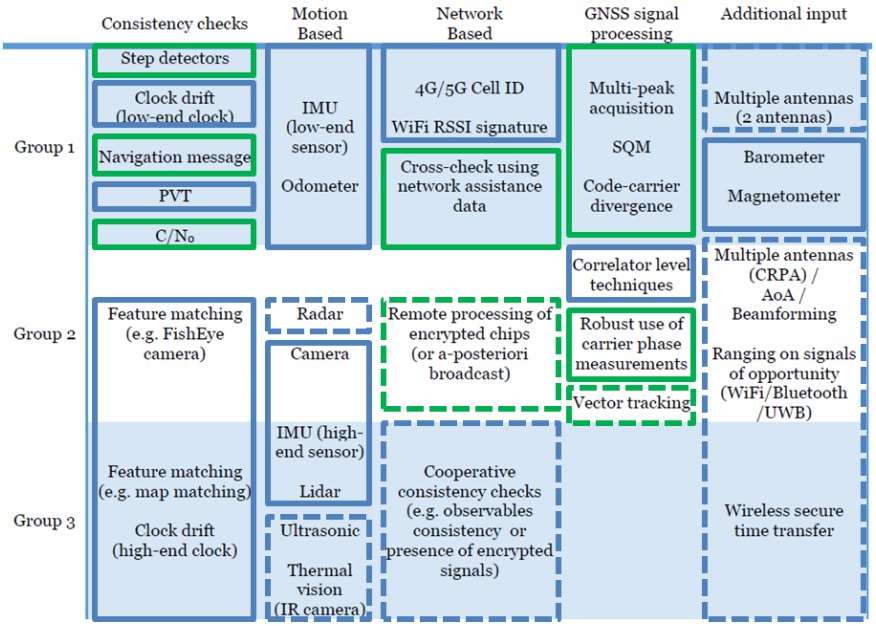

Figure 2-1: Categorization and allocation of anti-spoofing techniques in the testbed. Type of technology/information (columns) and level of maturity and costs (group 1-3). Allocation to the GNSS SW Receiver (green boxes) and the ML-supported Sensor Fusion (blue boxes). Techniques not fully emulated but flexibility for future extension is offered (dashed boxes).

The first part of the testbed architecture (left side), the data generation part, deals with real data acquisition and data simulation (red). The data is stored in a central file-based database with defined API interfaces and file formats (center). The user segment (right side) provides anti-spoofing techniques processing the data stored in the database. Integrated anti-spoofing techniques comprise classical anti-spoofing techniques implemented in the GNSS SW Receiver (Airbus SDR, in green), the developed machine learning supported sensor fusion engine (blue)

and a custom-off-the-shelf GNSS receiver (Ublox, in yellow). Real GNSS baseband samples have been recorded that are processed by the spoofing attack simulator to inject different types of spoofing attacks, uncoordinated and coordinated attacks. The GNSS SW Receiver processes the baseband samples in software and generates GNSS observables, such as pseudorange, carrier phase, doppler shift and C/N0 measurements. The GNSS observables are processed by the machine learning supported sensor fusion. The COTS receiver is feed directly be replaying the authentic and the spoofed baseband samples. A performance analysis tool has been implemented in MATLAB to evaluate the performance of the different techniques comparing their output against a reference solution that has been generated from additional sensor measurements acquired during data collection.

2.2 Anti-Spoofing Sensor Fusion and Machine Learning

Figure 2-2 shows the overall machine learning supported anti-spoofing sensor fusion architecture. The upper part contains the machine learning based spoofing detection, the lower part contains the anti-spoofing sensor fusion.

There are two approaches at the core of the machine learning based spoofing detection:

- A neural network (NN) based approach that takes multi-dimensional time series data, i.e. multicorrelator outputs, as input to predict whether a GNSS attack is taking place or not, at each point in time. The approach has two stages: An unsupervised feature encoding step based on recurrent autoencoders (RAE), that learns a compact low-dimensional representation from the input data without requiring labelled data (unsupervised). This step is followed by a supervised deep NN classifier, that is trained using labelled data to predict whether spoofing is taking place or not.
- A camera-based approach is used to detect spoofing on the one hand, and to provide a global position and heading information on the other hand. A precise geo-referenced road map is generated from camera images and precise position and heading information a priori. To detect spoofing, camera images are projected to bird's eye view (BEV), placed into the map using the current position and heading estimate and matched to the map using classical feature extraction and feature matching. This method is used to verify the position obtained from the GNSS receiver with camera image observations and an available road map of the environment.

The two approaches described above are used for spoofing detection and mitigation as follows:

- GNSS multi-correlator data time series data are used as input to train the NN-based approach. Correlator data as output of GNSS signal acquisition and tracking show a characteristic triangular pattern when authentic signals are tracked. The characteristics of these patterns change when spoofed signals are tracked, or show a specific transition when the spoofed signal takes over.
- The general applicability of the NN-based approach has been further demonstrated using sensor data as input. Time series of vehicle wheel odometry velocity (magnitude)

and velocity (magnitude) derived from a GNSS/INS based solution was used as input and to train the NN-based approach. Since the vehicle wheel odometry sensor cannot be attacked by GNSS spoofing any deviation of these measurements relative to a GNSS-based solution indicate GNSS spoofing.

- The camera-based approach has been used to compute a 2D geo-referenced pose that is independent of GNSS signals and is fused in the sensor fusion to support positioning. Based on a positioning solution computed from all available sensors excluding the vulnerable GNSS signals, BEV-projected camera images are placed into and matched against the map. The obtained relative transformation yields a corrected pose that is fused as additional measurement into the sensor fusion. This step assists spoofing mitigation.
- Furthermore, it has been demonstrated that camera image to map matching can provide important indications for a spoofing attack. If the user position has been attacked and manipulated, matching the camera image to the map fails and indicates spoofing. The position calculated from the GNSS signals could not be verified. Slighter drifts from the actual position can also be detected. Camera image matching provides a relative transformation that reflects the drift of the position computed from the GNSS signal, if this drift goes beyond the expected noise of the position estimate spoofing is possible.

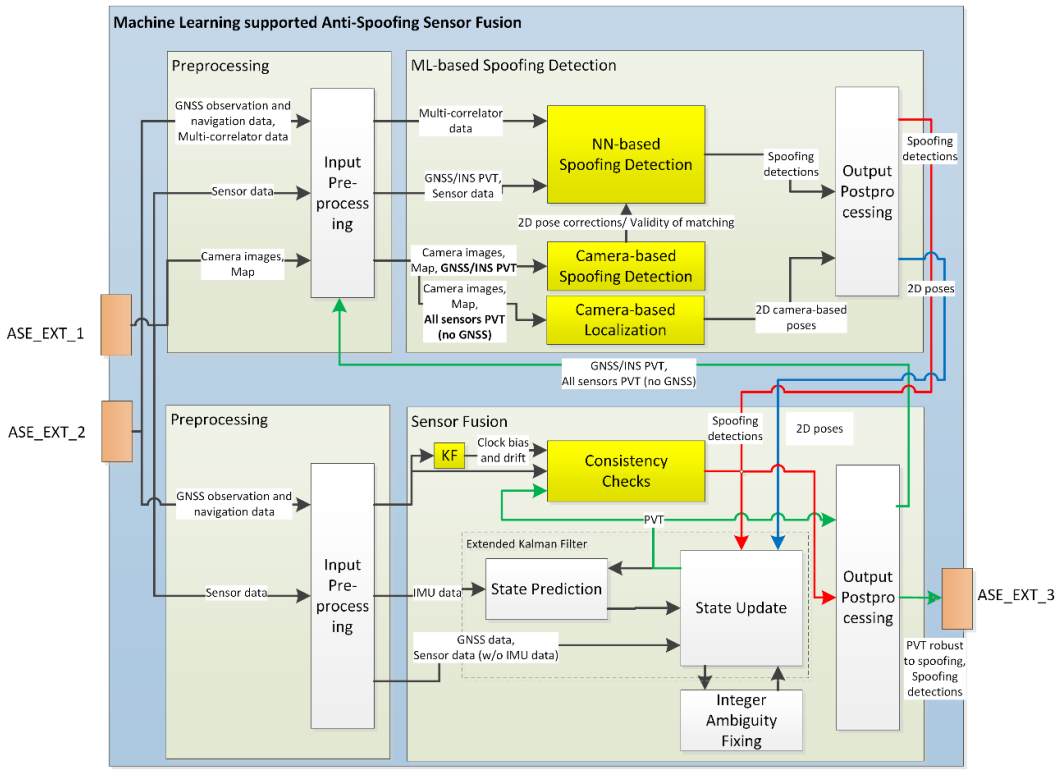

Figure 2-2: Machine learning (ML) supported anti-spoofing sensor fusion architecture. Top: ML-based spoofing detection. Bottom: Anti-spoofing sensor fusion.

The Sensor Fusion Engine is based on an Extended Kalman Filter (EKF) that performs the state prediction step using IMU measurements. All the other measurements are used to perform the state update. A proprietary and patented algorithm is used for integer ambiguity fixing that allows precise positioning at centimeter-level using carrier phase measurements.

To perform classical consistency checks, such as C/N0, clock bias and clock drift, PVT monitoring and navigation message checks, a Consistency Checks module has been implemented. A separate preceding Kalman Filter is used to estimate the receiver clock bias and drift. A file-based interface is provided to read signal specific spoofing detections, provided by the ML-based spoofing detection module. Fusion of spoofed signals is disabled in the Kalman Filter update step, such that only the authentic GNSS signals are used to compute the final user position. To support PVT calculation an interface is added to fuse additional pose information. 2D position and heading information computed by camera-image based map matching is fused using this interface to enhance the PVT solution.

3 Results

3.1 Data Collection and Data Simulation

Measurement Vehicle, Equipment, Location and Recorded Tracks

In the data collection campaign real sensor data including data to generate the reference trajectory have been recorded using a measurement vehicle, VW Golf 7, a sensor platform mounted on top of the vehicle and a baseband sample recorder, depicted in Figure 3-1.

Figure 3-1: Measurement vehicle (top left) with mounted data acquisition and reference system, the Integrated Sensor Platform (ISP) from ANavS GmbH for multi-GNSS, IMU, wheel odometry, camera and lidar data recording. Averna Recorder for BB-sample acquisition (top right), connected to one antenna of the ISP. Location of RTK initialization and static datasets acquisition (s2, s3) with open sky conditions (bottom).

Several datasets have been recorded in Munich-Obermenzing close to the interchange A8/A99 including urban areas and highway covering static, low-dynamic and high-dynamic user scenarios. The static scenarios have been recorded at the location shown in Figure 3-1 and marked in Figure 3-2. The three tracks that are marked in Figure 3-2 cover the dynamic scenarios.

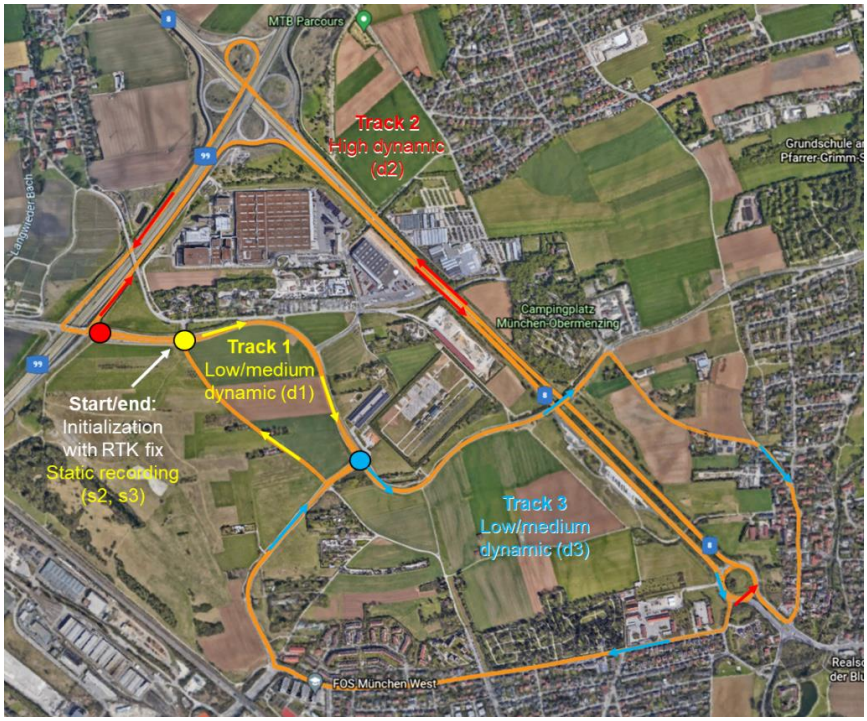

Figure 3-2: Overview of data collection in Munich-Obermenzing (interchange A8/A99). Dynamic use cases (tracks 1-3, dots mark start/end point). Initialization and static use cases (marked by white arrow). Another static use case (s1) was recorded at a different location, in Munich-Laim.

An overview of the recorded datasets and the injected simulated spoofing is given in Table 3-1. Each dataset covers at least 10min of recording to provide enough time for sensor fusion approaches to converge before injected spoofing attacks are activated. The dynamic datasets consist of three rounds of the respective track. This allows to split the data into training, validation and test set, for example to generate maps during training and using separate data for validation and testing for the camera-based approaches.

Use Case	Datasets	Duration (min:sec)	Simulated Spoofing Attacks
Static	s1, s2, s3	10:35, 10:46, 09:49	Timing attack (TA)
Low dynamic	d1, d3	10:38, 36:31	Meaconing attack (MA) and Coordinated (CA)
High dynamic	d2	24:50	Meaconing attack (MA) and Coordinated (CA)

Table 3-1: Overview of use cases and datasets. Each dynamic dataset consists of three rounds of the track.

In this study three types of spoofing have been simulated, timing attacks, meaconing attacks and coordinated attacks. Timing attacks are applied to the static use cases, meaconing and coordinated attacks are applied to the dynamic use cases, see Table 3-1. Meaconing represents an uncoordinated attack in which navigation signals are rebroadcast without any synchronization to the authentic signals. In contrast, timing attacks and coordinated attacks require a proper synchronization to the authentic signals to achieve sophisticated attacks that let the user position slowly drift away from the original position for example.

Reference Trajectory Generation

First of all, after data acquisition, a highly precise reference trajectory has been generated, which is visualized in Figure 3-3. The reference trajectory has been used for sensor data simulation on the one hand, and on the other hand for performance evaluation of the anti-spoofing techniques.

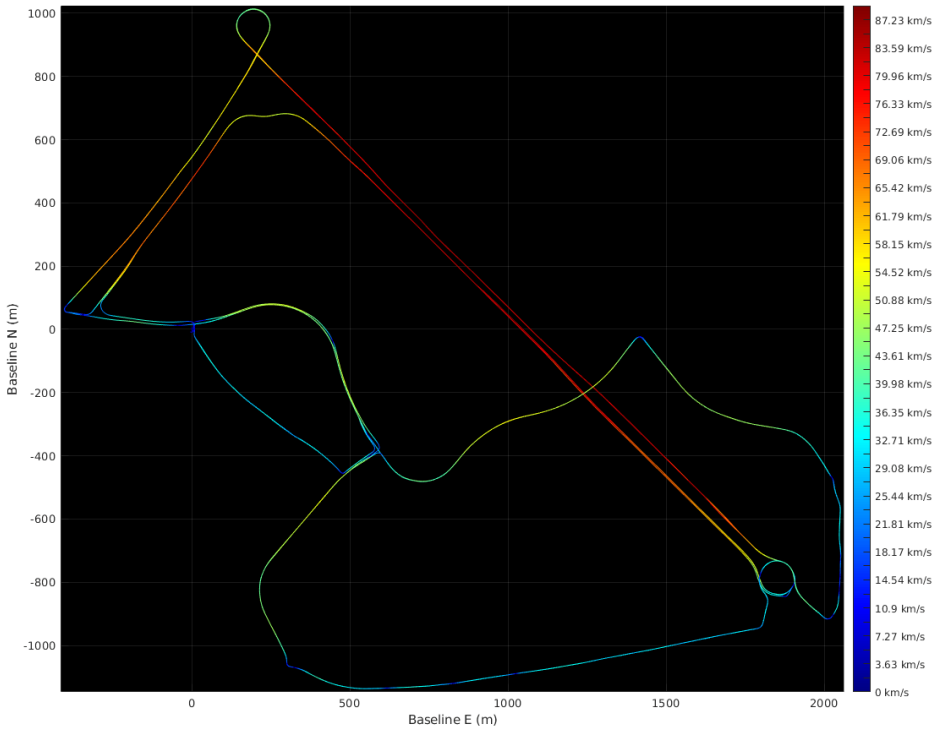

Figure 3-3: Reference trajectory computed using sensor fusion with multi-GNSS RTK (three Ublox receivers), IMU and wheel odometry. The velocity (magnitude) is color-coded. The highway high-dynamic track 2 appears "red", with high velocities. Tracks 1 and 2 appear "blue, green, yellow" with low and medium velocities, cp. Figure 3-2.

Simulated Data and Real Sensor Data

This study combines both sensor data simulation and the use of real sensor data. Table 3-2 summarizes the data simulated in the testbed, which comprises sensor data, network aiding ranges and simulated spoofing attacks represented by spoofing injected GNSS baseband samples. Real sensor data are acquired using the Integrated Sensor Platform (ISP) and the Averna BB-Sample recorder. Since GNSS baseband sample recording yields large raw data files,

in this activity the GNSS signals to be recorded were restricted to GPS and Galileo signals at a samples rate of 12.5 MHz.

Simulated Data	**Real Sensor Data**
Sensor Data Simulation • Barometer • High-end IMU Network Aiding Simulation • Simulated LTE 4G-5G network positioning at ranging level Spoofing Attack Simulation • Timing Attack • Meaconing Attack • Coordinated Attack	Real Sensor Data recorded with the Integrated Sensor Platform (ISP) and the Averna BB-Sample Recorder • GNSS Baseband Samples (GPS and Galileo, samples rate 12.5 MHz) • Low-cost IMU • Wheel odometry • Camera • LiDAR

Table 3-2: Overview of simulated data and recorded real sensor data.

In the following samples of the resulting simulated and real sensor measurements are presented including:

- Real Sensor Low-Cost IMU and Simulated High-End IMU
- Camera Images
- Anchor Selection for Network Aiding Simulation
- Spoofing attack simulation

Figure 3-4 shows samples of the recorded low-cost IMU data compared to the simulated high-end IMU data. The high-end IMU shows less noise, which mainly reflects the fact that vehicle vibrations are present in the real tests but not in simulation. Apart from the noise both IMU data are very similar to each other, which validates the generation of simulated IMU data.

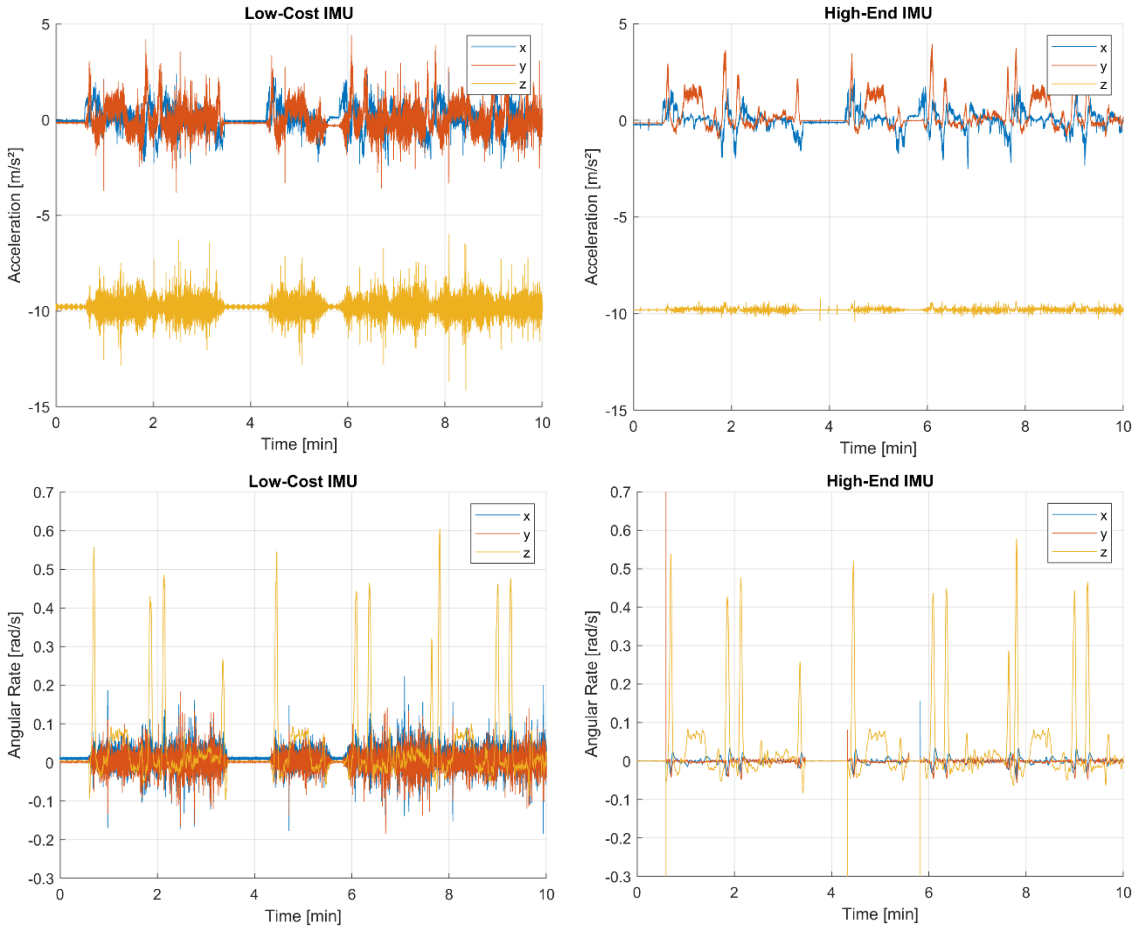

Figure 3-4: Comparison between Low-Cost IMU (left column) and High-End IMU (right column) for acceleration (top) and angular rates (bottom).

In Figure 3-5 sample images of the high-quality RGB camera are shown, which is mounted inside the Integrated Sensor Platform. Exemplarily images from low-dynamic urban scenes and from high-dynamic highway scenes are shown. In this study these camera images have been used to generate a precise road map on the one hand, and to perform camera image to map matching on the other hand. Resulting camera-based pose measurements support the sensor fusion and camera image to map matching has been demonstrated to be applicable for GNSS-pose verification and spoofing detection.

Figure 3-5: Sample camera images recorded in low-dynamic (top row) and high-dynamic scenarios (bottom row).

Network aiding simulation has been performed to simulate LTE 4G-5G network positioning at ranging level. To increase realism anchor positions have been chosen at real radio mast positions, in our study two anchors have been chosen, which is the minimal reasonable configuration, see Figure 3-6.

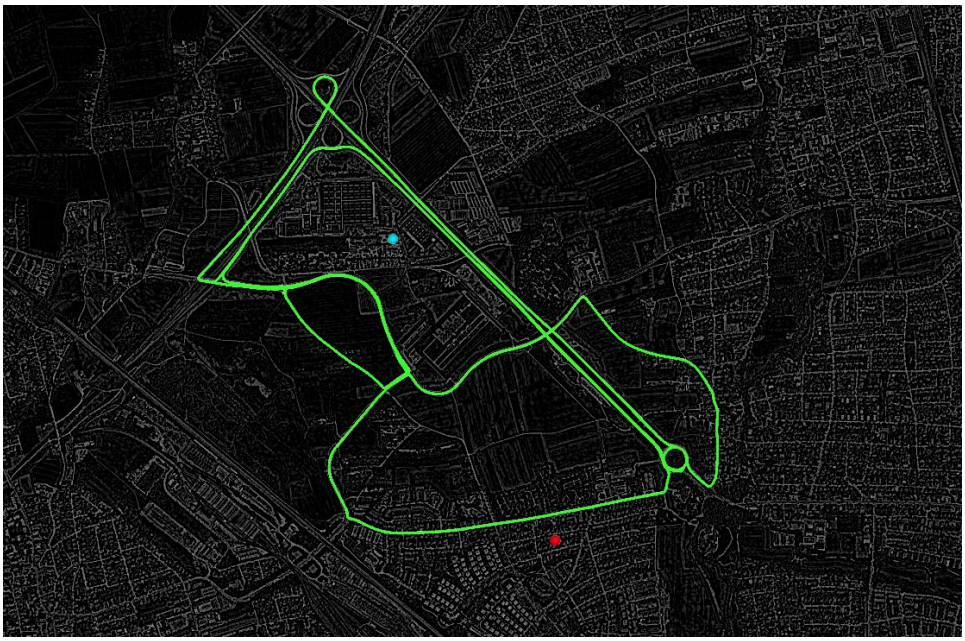

Figure 3-6: Anchors selected for network aiding ranges simulation. Complete trajectory (green) with 1st anchor (blue) and 2nd anchor (red).

The developed spoofing attack simulator uses the recorded baseband samples to generated spoofing injected baseband samples that can be processed by the GNSS SW receiver. Three types

of spoofing have been implemented: Timing attack (spoofs time information), meaconing attack and coordinated attack (spoof position information). Various parameters can be configured to adjust spoofing attacks, such as start and end times, signal strength and target fake position, for example. Figure 3-7 shows example results of a meaconing and a coordinated attack compared to the reference trajectory.

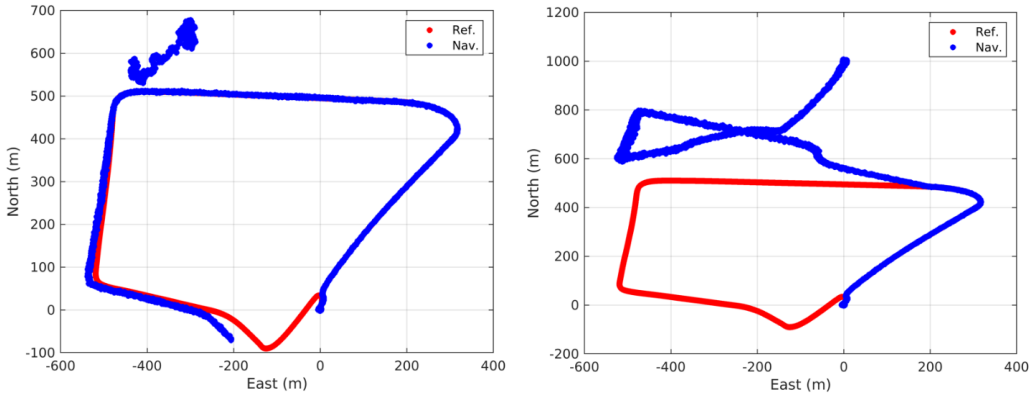

Figure 3-7: Results of the validation of the uncoordinated (left) and the coordinated attack simulation (right). The reference trajectory (red) is plotted against the estimated receiver position (blue). Meaconing leads to a jump to the fake position (left), the coordinated attack leads to a smooth drift from the correct to the fake position (right).

3.2 Tests Overview

Table 3-3 gives an overview of the test categories applied for generating the test plan. The categories include the three use cases, the phases testbed training and performance evaluation, nominal and faulty scenarios and finally the three anti-spoofing techniques, classical anti-spoofing techniques (GNSS SW receiver), machine learning based anti-spoofing sensor fusion and the commercial-off-the-shelf (COTS) receiver.

Use Cases	Phases	Scenarios	Anti-Spoofing Techniques
UC1: static UC2: low dynamics UC3: high dynamics	TT: Testbed Training PE: Performance Evaluation	NS: Nominal FS: Faulty	CL: Classical ML: Machine Learning CO: Commercial RX

Table 3-3: Test categories overview

The following two sections show results obtained during performance evaluation for the static and the dynamic use cases in faulty conditions for all the three anti-spoofing techniques. Faulty scenarios start in nominal conditions without spoofing until the spoofing attack starts. Timing attacks last until the end of the dataset. Meaconing and coordinated attacks are tested within a

single dataset, where the coordinated attack follows the meaconing attack after a 100sec period without spoofing.

3.3 Static Use Case

Figure 3-8 illustrates the performance of the GNSS SW receiver and the COTS receiver anti-spoofing techniques in a static use case when facing a simulated timing attack starting at 150s.

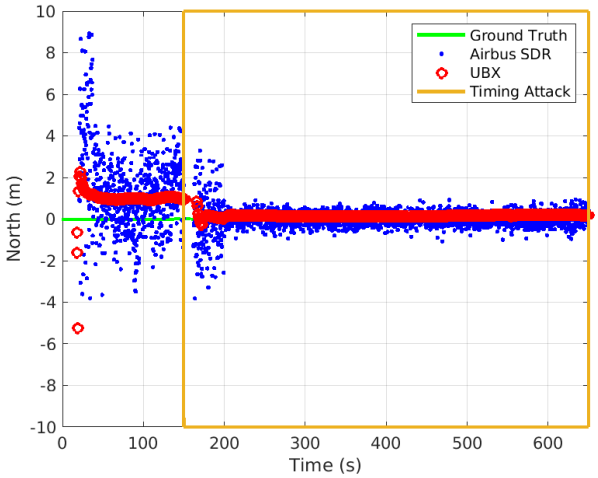

Figure 3-8: GNSS position error in North direction observed in static condition, faulty scenario. The start of the timing attack is indicated (golden line). The position error of the COTS receiver (red) is shown in comparison to the GNSS SW receiver (blue) and the reference (green).

First of all, the COTS receiver applies stronger smoothing and filtering compared to the GNSS SW receiver that does not apply any smoothing. Thus, the GNSS SW receiver PVT solution shows an increased noise level. Furthermore, both the COTS receiver and the GNSS SW receiver were fooled by the timing attack, since a significant change in the noise characteristics takes over as soon as the spoofing attack starts. The fact that the position is driven towards the correct position when spoofing starts is due to the fact that only the timing information has been spoofed but not the positional information (e.g. pseudorange measurements), and prior to the spoofing attack, the position solution shows the residual impact of ionosphere, troposphere and other error sources, which are not present to the same extend in the spoofing signals.

Figure 3-9 and Figure 3-10 show the results of the ML-based spoofing detection and Anti-Spoofing Sensor Fusion for the static use case with timing attack. Figure 3-9 (left) indicates the GPS and Galileo signals that have been detected as being spoofed (red dots) and the start of the spoofing attack. Spoofing was correctly identified for all spoofed signals during the complete spoofing attack. Moreover, the start of the spoofing attack was detected accurately. Very few false positive spoofing detections occurred at the beginning of the dataset. However, when these spoofing detections occurred, no measurements were available and the signals could not be tracked by the software receiver. These false positive detections correctly identify unusually low signal quality and have no effect since no measurements were generated. In Figure 3-10

(right) the performance of the ML-based Anti-Spoofing Sensor Fusion is shown, compared to the performance of the other techniques presented in Figure 3-8. The position estimate converges closely to the reference position before spoofing starts, in contrast to the GNSS only techniques (Figure 3-8), which can be attributed to the additional fusion of simulated network aiding range measurements. No obvious effect of spoofing can be observed from the position error plotted in Figure 3-9 (right). Spoofing mitigation was successful with all detected spoofed samples being discarded. During the spoofing attack, since all GNSS measurements are rejected, only the simulated range measurements contribute to the PVT solution.

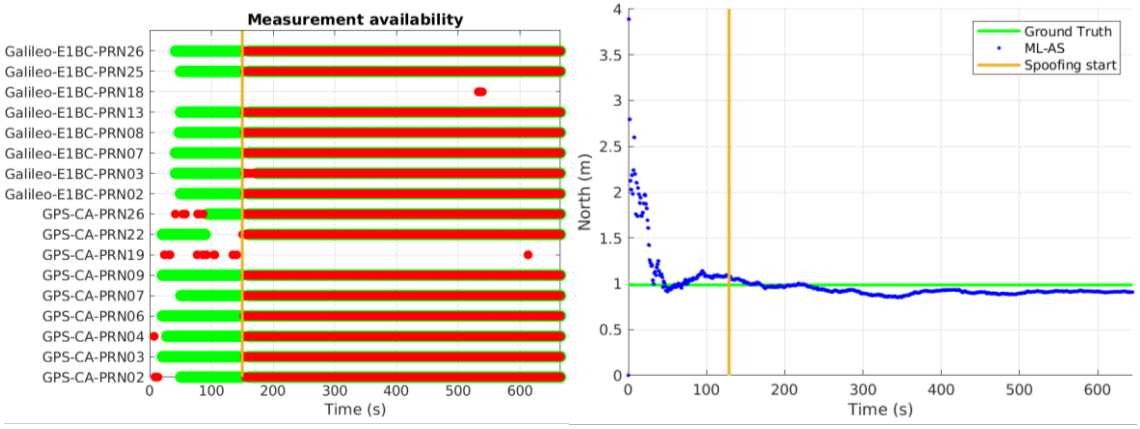

Figure 3-9: ML-based spoofing detection in static condition, faulty scenario. Left: GNSS measurement availability (green) and spoofing detections (red). Right: Anti-Spoofing (AS) Sensor Fusion position in North direction. The ML-AS position estimate (blue) is shown in comparison to the reference (green). The start of the timing attack is indicated (golden line).

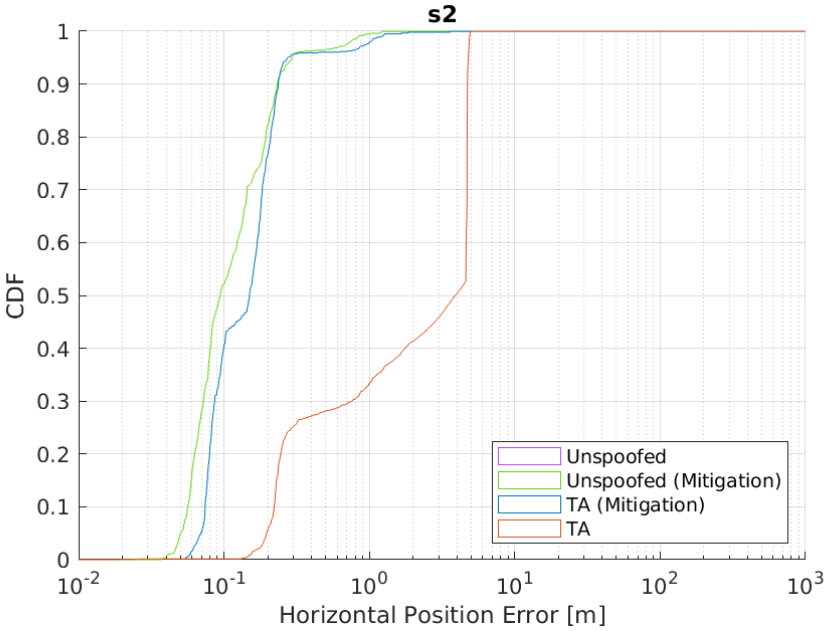

Figure 3-10: ML-based Anti-Spoofing Sensor Fusion: Comparison of spoofing detection and mitigation performance between nominal and faulty scenarios for timing attack with and without mitigation enabled. The horizontal position error is given as a cumulative distribution function (CDF).

Figure 3-10 reports the horizontal position error as a cumulative distribution function comparing timing attack without spoofing mitigation (orange) and with spoofing detection and mitigation enabled (blue). Similar plots are provided for the nominal case with no spoofing injected (unspoofed). A clear performance improvement is achieved when applying the developed spoofing detection and mitigation, and the performance is nearly recovered to the performance achieved in the nominal case. Both result plots for the unspoofed cases (purple and green) fall onto each other, showing that the anti-spoofing technique does not decrease the performance if no spoofing is present.

3.4 Dynamic Use Cases

Figure 3-11 illustrates the performance of the GNSS SW receiver and the COTS receiver anti-spoofing techniques in a low dynamic use case when facing a coordinated attack followed by a meaconing attack. The position estimates in North direction and in North and East direction are plotted against the reference.

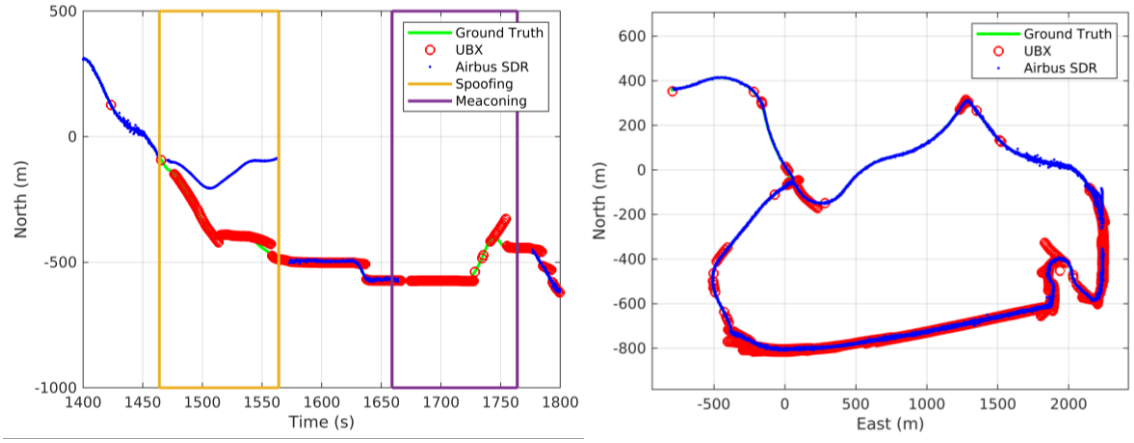

Figure 3-11: Position estimate in North direction in low dynamic conditions, faulty scenario for the COTS receiver (red) in comparison to the GNSS SW receiver (blue) and the reference (green). Left: Position in North direction. Right: Position in North and East direction. The time ranges of spoofing attacks are indicated. A coordinated attack is followed by a meaconing attack.

The GNSS SW receiver was fooled by both spoofing attacks. The coordinated attack results in a position drift, see Figure 3-11 first box, and the meaconing attack results in a jump to the fake position, not visible in Figure 3-11, because the fake position is far off. The COTS receiver was not affected by the spoofing attacks, however the availability of the COTS receiver position fixes was rather low (red markers) and the COTS receiver shows strong position extrapolation (red markers overshooting the reference trajectory) even when valid GNSS measurements are available, as confirmed by the GNSS SW receiver.

Figure 3-12 and Figure 3-13 show the results of the ML-based spoofing detection and Anti-Spoofing Sensor Fusion for the low dynamic use case with a meaconing attack followed by a coordinated attack. Figure 3-12 indicates the GPS and Galileo signals that have been detected as being spoofed (red dots) and the start and end times of the spoofing attacks. Spoofing was

correctly identified for all spoofed signals during both spoofing attacks. Moreover, the start of the spoofing attacks was detected accurately, also the end of spoofing attacks was detected correctly, though some detections extend a bit beyond the spoofing end. An exception represent false alarm detections for the three Galileo signals (PRN08, PRN07 and PRN02) in the time range between both spoofing attacks. Apart from this very few false positive spoofing detections occurred before the first and after the second attack. These are related to cases where no measurements were available and the signals could not be tracked. These false positive detections correctly identify unusually low signal quality and have no effect since no measurements were generated. The false alarm detections between the spoofing attacks are not critical since they do not introduce spoofed samples but discard signals that could be used for navigation. Still, enough GNSS signals remain and together with the additional sensors yield a performant solution as shown in and Figure 3-13 (right).

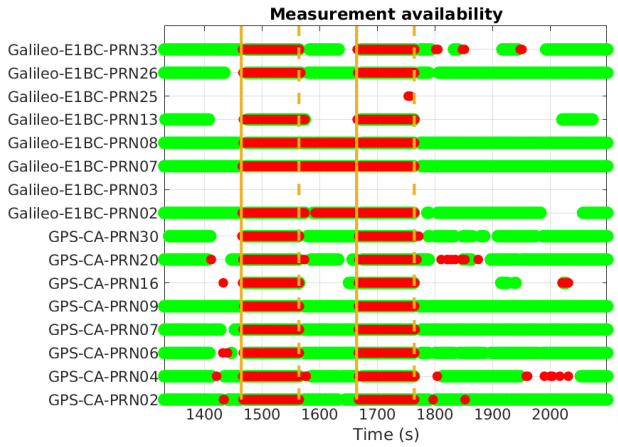

Figure 3-12: GNSS measurement availability (green) and Machine Learning based spoofing detections (red) for low dynamic conditions, faulty scenario. The start and end of spoofing attacks is indicated (golden lines and dashed lines). A coordinated attack is followed by a meaconing attack.

Figure 3-13 shows that spoofing affects the solution if mitigation is disable (left) and spoofed samples are processed. On the other hand, when spoofing mitigation is enabled (right) spoofed signals are detected (Figure 3-12) and discarded which results in a solution close to the reference. During coordinated attack the position is dragged away (top left, first attack window), and during meaconing attack (top left, second attack window) the position is affected as well, but does not jump to the fake position as in the GNSS-only solution from the GNSS SW receiver (Figure 3-11, left). Sensor fusion with absolute camera-based poses and network aiding ranges helps to compensate the meaconing attack when spoofed GNSS samples as used. After the first spoofing attack the position solution is still falsified but recovers to the correct solution after 66 seconds approximately. The sensor fusion filter was affected by the first attack such that it could not recover immediately after the first spoofing attack. In contrast, after the meaconing attack ends, the correct position was achieved much faster. Notably the scenario of multiple spoofing attacks following each other has not been targeted initially and thus was not considered before in the training and testbed implementation phase.

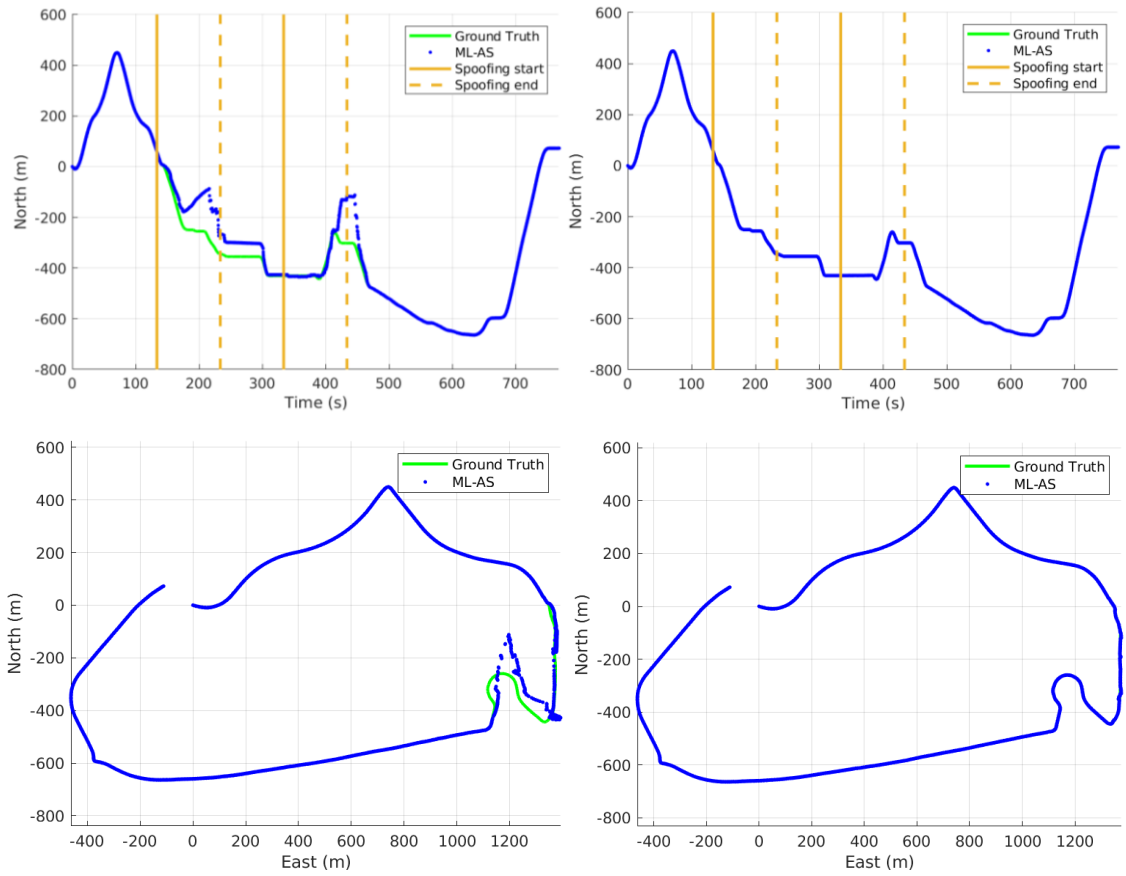

Figure 3-13: ML-based Anti-Spoofing (AS) Sensor Fusion position error in North (top) and in North and East direction (bottom) observed in low dynamic conditions, faulty scenario. Left: Spoofing mitigation disabled. Right: Spoofing mitigation enabled. The start and end of spoofing attacks is indicated (golden lines and dashed lines). A coordinated attack is followed by a meaconing attack. The ML-AS position error (blue) is shown in comparison to the reference (green).

Figure 3-14 reports the horizontal position error as a cumulative distribution function comparing the scenario of a meaconing attack followed by a coordinated attack without spoofing mitigation (orange, red) and with spoofing detection and mitigation enabled (dark and light blue). Two variants have been tested, “ver2” represents a more sophisticated attack in which the spoofer power level is only increased at the start of the spoofing attack (duration of 30s) instead applying a constant high power level during spoofing. Similar plots are provided for the nominal case with no spoofing injected (unspoofed). A clear performance improvement is achieved when applying the developed spoofing detection and mitigation, and the performance is nearly recovered to the performance achieved in the nominal case. The performance achieved for the more sophisticated attack “ver2” is comparable to the original attack variant. Both result plots for the unspoofed cases (purple and green) fall onto each other, showing that the anti-spoofing technique does not decrease the performance if no spoofing is present. Figure 3-15 shows validation results that illustrate the performance of the developed ML-based Anti-Spoofing Sensor Fusion under the tested meaconing and coordinated spoofing attacks plotting the resulting trajectories in the region where spoofing was injected.

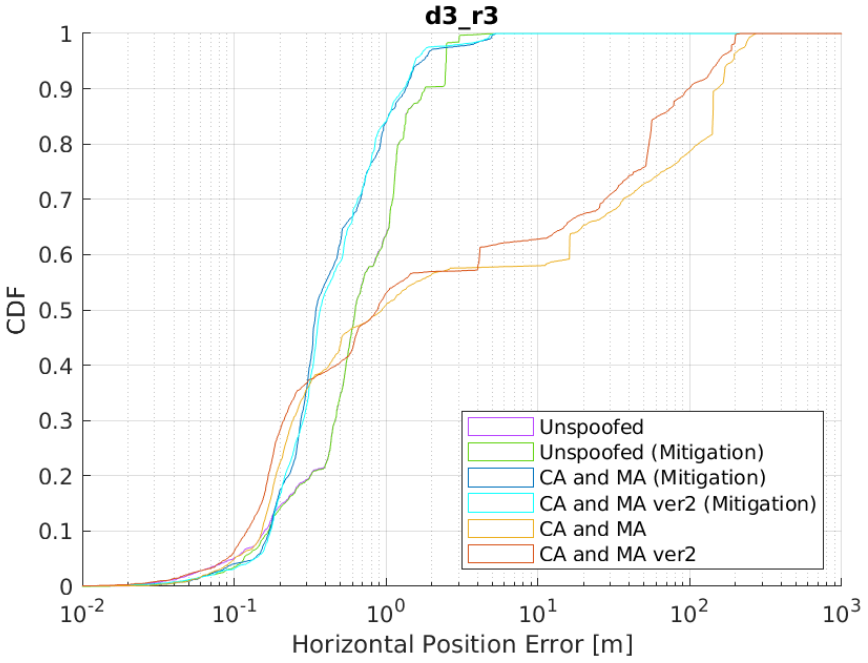

Figure 3-14: ML-based Anti-Spoofing Sensor Fusion: Comparison of spoofing detection and mitigation performance between nominal and faulty scenarios for coordinated followed by meaconing attack with and without mitigation enabled. The horizontal position error is given as a cumulative distribution function (CDF).

Figure 3-15: Validation results comparing the ML-based Anti-Spoofing Sensor Fusion performance with enabled spoofing detection and mitigation (green, dark and bright blue) and without spoofing detection and mitigation (orange, red). Mitigation enabled with no spoofing (green), mitigation enabled coordinated attack (CA) (dark blue), mitigation enabled meaconing attack (MA) (bright blue). Without spoofing detection and mitigation: CA (orange) and MA (red). The region where the spoofing attack was injected is shown.

4 Conclusion

In this study, state-of-the-art and emerging technologies that can be used for PVT assurance at the receiver level have been consolidated and a flexible anti-spoofing testbed has been developed that can integrate different sensors and technologies. The testbed implementation contains modules for sensor data, network aiding and spoofing attack simulation, a GNSS software receiver for processing GNSS baseband samples and a machine learning based anti-spoofing sensor fusion. A performance analysis tool allows to evaluate and compare different solutions. Three anti-spoofing techniques, classical techniques implemented in the GNSS SW receiver, the machine learning based anti-spoofing sensor fusion and the anti-spoofing capabilities of a COTS receiver have been evaluated in static and dynamic use cases using simulated spoofing attacks injected by the spoofing attack simulated. The developed machine learning based anti-spoofing sensor fusion shows very promising results for spoofing detection and mitigation in the evaluation performed in the testbed using simulated spoofing attacks. The main contributing machine learning techniques developed are a neural network approach that predicts spoofing from time series of GNSS multi-correlator data and a camera-image-to-map-matching approach that yields camera-based georeferenced poses for sensor fusion, which has been demonstrated to be effective for spoofing detection as well. The neural network approach has also been demonstrated to be effective for spoofing detection using sensor data. Besides the machine learning techniques, the extended Kalman Filter (EKF) based anti-spoofing sensor fusion engine is a main component. It includes classical consistency checks, spoofing mitigation by rejecting spoofed GNSS signals and fusion of all real and simulated testbed measurements, including GNSS measurements, IMU, wheel odometry and barometer measurements, camera-based poses and network aiding ranges.

In conclusion, the goals of the project have been achieved completely. Next steps can be built on the results obtained using visual sensors and maps that allow to obtain independent global position and heading information, without the need of additional infrastructure, apart from a prebuilt map. In this sense, the development of secure, robust and accurate positioning enabled by camera-, LiDAR and/or RADAR-based localization on custom and publicly available maps is a highly interesting direction. Further detailed discussion on potential system future evolution is provided in the project technical note on “Guidelines for Sensor fusion-based GNSS Anti-Spoofing and Potential System Evolution”.

End of Document