Validation of Beamforming Techniques for Handheld GNSS Receivers in the Presence of Correlated Signals

Lucía Pallarés-Rodríguez[1], Gonzalo Seco-Granados[1], José A. López-Salcedo[1]
Bárbara de Matos[2], Pedro Boto[2], Alexandru Budianu[3]
[1]*Universitat Autònoma de Barcelona,* [2]*GMV Portugal,* [3]*European Space Agency*

BIOGRAPHIES

Bárbara de Matos received her M.Sc. degree in physics engineering from Universidade de Coimbra, Portugal in 2022. Since joining GMV in 2024, she has been involved in several projects related with GNSS receiver design, including acquisition, tracking and beamforming.

Pedro Boto received his M.Sc. degree in aerospace engineering from Instituto Superior Técnico of Lisbon, Portugal in 2014. Since joining GMV in 2014, he has been involved in several projects related with GNSS receiver design, specifically on signal processing techniques for acquisition, tracking and also on monitoring the quality of GNSS signals.

Lucía Pallarés-Rodríguez received her M.Sc. degree in Telecommunication Engineering from Universitat Autònoma de Barcelona in 2024. She is currently a Ph.D. student at the SPCOMNAV research group from Universitat Autònoma de Barcelona, and her research is centered on array signal processing for positioning systems.

Gonzalo Seco-Granados received the Ph.D. degree in Telecommunication Engineering from Universitat Politècnica de Catalunya in 2000, and the M.B.A. degree from IESE Business School in 2002. He is currently a Professor in the Department of Telecommunications at Universitat Autònoma de Barcelona. His research interests include GNSS, and beyond 5G integrated communications, localization, and sensing.

José A. López-Salcedo received the M.Sc. and Ph.D. degrees in Telecommunication Engineering from Universitat Politècnica de Catalunya, in 2001 and 2007, respectively. He currently serves as a Professor in the Department of Telecommunications at Universitat Autònoma de Barcelona. His research interests lie in the field of signal processing for communications and navigation.

Alexandru Budianu received his Engineer's Degree in Telecommunication Engineering from "Gheorghe Asachi" Technical University, Romania, in 2009, and his Ph.D. degree from the University of Twente, Netherlands, in 2015. He currently serves as a Radio Navigation System Engineer at the European Space Agency.

ABSTRACT

With the proliferation of the Internet of Things (IoT), autonomous driving, and the smart city paradigm, the number of applications relying on accurate positioning and timing information has grown significantly. In this context, Global Navigation Satellite Systems (GNSS) represent a key technology, providing global coverage and high accuracy. However, GNSS performance is known to degrade in urban and suburban environments, due to signal blockages and propagation impairments. In particular, two phenomena can severely impact receiver performance in these scenarios: multipath and spoofing. Their high level of correlation with the desired line-of-sight (LoS) signal distorts the correlation output, inducing large errors in code delay and phase estimation.

A promising approach to detect these undesired contributions and counteract their effects is the deployment of multiple antennas. In this work, multi-antenna detection and mitigation techniques effective against multipath and spoofing are analyzed, with a focus on mass-market GNSS receivers equipped with a limited number of antennas. The first part of this work presents the theoretical formulation of the proposed techniques, supported by numerical simulations under representative scenarios. To verify the performance observed in the simulations, the second part presents the results of a thorough evaluation of the algorithms in a controlled laboratory environment, serving as an initial validation bridging simulations and real-world implementation, and providing a comparison with commercially available mass-market receivers. Finally, to complete the assessment of beamforming in small GNSS

receivers, the techniques are tested using real signals captured under open-sky conditions, addressing the antenna calibration challenge and enabling future evaluation under real multipath and spoofing environments.

1 INTRODUCTION

Global Navigation Satellite Systems (GNSS) play a crucial role in modern society, enabling a wide range of applications across numerous sectors. Their ubiquity and global coverage make them a cornerstone technology for navigation, positioning, and timing applications. However, a significant portion of these applications target urban environments, where GNSS performance tends to be degraded due to obstructed visibility and other propagation impairments. As a result, mass-market receivers operating in these scenarios are highly vulnerable.

In this context, a major source of performance degradation is multipath reflections. Multipath occurs when replicas of the line-of-sight (LoS) signal arrive at the receiver after being reflected by nearby structured. These replicas are delayed and attenuated versions of the LoS signal coming directly from the satellite, and thus remain strongly correlated with it. This high level of correlation has severe impact on the despreading process, distorting the output and biasing both code and carrier tracking. A similar effect occurs during a spoofing attack. Spoofed signals are counterfeit transmissions that aim to deceive the receiver, so it computes a false positioning solution. As in the presence of multipath reflections, the high correlation between authentic and spoofed signals causes a distortion at the output of the correlation function, leading to large errors in subsequent stages of the GNSS receiver. Despite the difference between these phenomena, in both cases the receiver encounters correlated replicas that degrade its measurements and compromise its navigation solution.

Several approaches have been proposed in the literature to mitigate these effects. Conventional methods include narrow correlators, multipath estimating delay lock loops, and other refined code-tracking schemes. However, the rise of multi-antenna solutions in the wireless communications domain over the past decades has motivated the adoption of this technology in GNSS. Deploying multiple antennas enables the receiver to exploit the spatial domain, allowing it to filter out undesired contributions based on their spatial characteristics. The array's capacity to detect and suppress these contributions is proportional to the number of antennas, which is often limited in small GNSS receivers. Furthermore, these smaller devices, such as smartphones and tablets, are typically the most vulnerable to the propagation impairments in urban and suburban areas. To overcome this limitation, advanced spatial diversity or beamforming techniques are required, offering promising detection and mitigation capabilities even with small antenna arrays.

The DIVAL activity (Multipath and Spoofing Mitigation in Handheld Devices with Diversity Algorithms), funded by the European Space Agency (ESA), explored this trade-off by designing, implementing, and validating advanced spatial diversity algorithms tailored to small arrays, consisting of two antennas (size of a smartphone) or four antennas (size of a tablet). After an extensive study of the existing techniques in the literature, and a trade-off analysis between accuracy and complexity, the selected detection techniques included Single-Correlation Multi-Antenna (SC-MA) analysis (Egea-Roca, López-Salcedo, & G. Seco-Granados, 2014) and Carrier Phase Differences (CPD) methods (Domínguez, et al., 2015), while mitigation relied on Hybrid (HYB) (Seco-Granados, Fernández-Rubio, & Fernández-Prades, 2005) and Power-Based Capon (PBC) (Mañosas-Caballú, Swindlehurst, & Seco-Granados, 2021) beamformers.

The first part of this paper covers the theoretical formulation of the proposed techniques and presents numerical simulations under representative scenarios. To validate the performance observed in these simulations, the second part addresses laboratory validation using synthetic signals, providing a controlled environment to assess the techniques under a variety of conditions, including multipath and spoofing scenarios with both single and multiple sources, and allowing comparison with mass-market receivers. Lastly, to complete the evaluation of the techniques, the paper extends the analysis to real-world experiments with signals captured under open-sky conditions, providing end-to-end validation of the techniques implementation and laying the groundwork for future tests under real multipath and spoofing conditions. This progression from theory to simulation, laboratory, and field testing ensures a comprehensive assessment of the proposed methods, and their suitability for practical integration into handheld GNSS receivers.

2 TECHNICAL DESCRIPTION

Consider an antenna array equipped with L elements receiving the superposition of a GNSS signal coming directly from the satellite, i.e., LoS signal, along with a correlated replica product of either a multipath reflection or a spoofed signal. After the despreading process, these contributions rise above the noise floor and become visible. At this stage, the overall received signal at the different antennas, referred to as the post-correlation signal, can be expressed as a $L \times 1$ vector following,

$$\boldsymbol{x}[n] = \alpha_0 \boldsymbol{a}(\theta_0) s[n;\tau_0] + \alpha_m \boldsymbol{a}(\theta_m) s[n;\tau_m] + \boldsymbol{n}[n], \quad (1)$$

where $s[n;\tau_i]$ is the code correlation output of the i-th contribution with propagation delay τ_i, α_i denotes the amplitude of the signal, and $\boldsymbol{a}(\theta_i) \in \mathbb{C}^{L\times 1}$ is the spatial signature or steering vector of the array in direction θ_i. The subscript 0 in Equation (1) refers to the LoS signal, whereas m denotes for the undesired source. Finally, the term $\boldsymbol{n}[n] \in \mathbb{C}^{L\times 1}$ models white Gaussian noise with zero mean and variance σ_n^2; $\boldsymbol{n}[n]\sim\mathcal{CN}(0,\sigma_n^2\boldsymbol{I})$.

2.1 Detection Techniques

The spatial information of the received signals, captured by the steering vectors, can be used by the receiver to detect the presence of undesired contributions, such as multipath reflections or spoofed signals. Although different approaches can be followed, this work focuses on two detection techniques: SC-MA, which targets multipath, and CPD, which targets spoofing.

2.1.1 Single-Correlation Multi-Antenna (SC-MA)

The SC-MA multipath detection method is a sequential detection algorithm based on the cumulative sum (CUSUM) approach, which tracks the cumulative deviation of a predefined metric from its reference value to detect changes (Egea-Roca, López-Salcedo, & G. Seco-Granados, 2014). In the absence of any undesired component, the received signal contains only the LoS contribution, meaning that the output of the prompt correlator is expected to take the form $\boldsymbol{x}^0[P;n] = \alpha_0\boldsymbol{a}(\theta_0) + \boldsymbol{n}[n]$, where P indicates the sample corresponding to the prompt. However, in the presence of a multipath replica, the prompt output becomes $\boldsymbol{x}^1[P;n] = \alpha_0\boldsymbol{a}(\theta_0) + \alpha_m\boldsymbol{a}(\theta_m)s[n;\tau_m] + \boldsymbol{n}[n]$. Assuming the DoA of the LoS signal is available, obtained through the GNSS ephemeris, the receiver can perform the following computation

$$\boldsymbol{u}[P;n] = \boldsymbol{x}[P;n] \odot \boldsymbol{a}^*(\theta_0), \tag{2}$$

where $\odot$ is the elementwise or Schur-Hadamard product. Under the scenario in which the only signal present is the LoS, $\boldsymbol{u}[P;k]$ becomes set of constant amplitudes across all the antennas affected by the noise contribution, i.e., $\boldsymbol{u}^0[P;n] = \alpha_0\boldsymbol{1}_L + \boldsymbol{n}[n]$. On the contrary, with additional contributions present in the received signal, the output of Equation (2) yields $\boldsymbol{u}^1[P;n] = \alpha_0\boldsymbol{1}_L + \alpha_m[\boldsymbol{a}(\theta_m) \odot \boldsymbol{a}^*(\theta_0)]s[n;\tau_m] + \boldsymbol{n}[n]$, where the spatial contribution of the reflection remains. Accordingly, the signal model after removing the LoS spatial contribution is given by the following hypotheses:

$$\begin{aligned}\mathcal{H}^0&: \boldsymbol{u}[P;n] = \alpha_0\boldsymbol{1}_L + \boldsymbol{n}[n],\\ \mathcal{H}^1&: \boldsymbol{u}[P;n] = \alpha_0\boldsymbol{1}_L + \alpha_m[\boldsymbol{a}(\theta_m) \odot \boldsymbol{a}^*(\theta_0)]s[n;\tau_m] + \boldsymbol{n}[n].\end{aligned} \tag{3}$$

This structure is exploited to define the observations for implementing the CUSUM algorithm. Since, under $\mathcal{H}^0$, the amplitude across all antennas is approximately α_0, this scenario can be distinguished from $\mathcal{H}^1$ using the following metric,

$$\rho_{SC-MA}(k) \doteq \frac{1}{L-1}\sum_{\boldsymbol{l=2}}^{\boldsymbol{L}}\left(\left[\boldsymbol{u}[P;k]\right]_{\boldsymbol{l}} - \left[\boldsymbol{u}[P;k]\right]_{\boldsymbol{1}}\right), \tag{4}$$

where $\left[\boldsymbol{u}[P;k]\right]_{\boldsymbol{l}}$ denotes the l-th entry of vector $\boldsymbol{u}[P;k]$. Note that under hypothesis $\mathcal{H}^0$, $\rho_{SC-MA}(k) \approx 0$, whereas under $\mathcal{H}^1$, the different contributions from each antenna will result in a non-zero detection metric. To implement the CUSUM using the metric in Equation(5) (4), Gaussian distributions are assumed for the hypotheses in Equation (3):

$$\begin{aligned}\mathcal{H}^0&: \rho_{SC-MA}(k)\sim\mathcal{N}\left(0,\frac{2\sigma_n^2}{L-1}\right),\\ \mathcal{H}^1&: \rho_{SC-MA}(k)\sim\mathcal{N}\left(\mathbb{E}_{\mathcal{H}^1}\{\rho_{SC-MA}\},\frac{2\sigma_n^2}{L-1}\right).\end{aligned} \tag{5}$$

Note that while σ_n^2 can be obtained in practice, $\mathbb{E}_{\mathcal{H}^1}\{\rho_{SC-MA}\}$ depends on both the amplitude α_m and DoA θ_m of the multipath reflection, which are unknown. For this reason, the maximum likelihood estimate of these parameters is used, resulting in the implementation of the generalized likelihood ratio (GLR) algorithm as the decision rule for the CUSUM, given by

$$g_{GLR}(k) = \frac{1}{2\sigma_{SC-MA}^2}\max_{1<n<k}\frac{1}{k-n+1}\left[\sum_{i=n}^{k}(\rho_{SC-MA}(i) - \mathbb{E}_{\mathcal{H}^0}\{\rho_{SC-MA}\})\right]^2, \tag{6}$$

where in the case under analysis $\mathbb{E}_{\mathcal{H}^0}\{\rho_{SC-MA}\} = 0$ and $\sigma^2_{SC-MA} = {2\sigma_n^2}/{L-1}$. Finally, a change is detected when $g_{GLR}(k)$ surpasses a certain threshold h, defined to guarantee a given false alarm probability.

2.1.2 Carrier Phase Differences (CPD)

The CPD is a multi-antenna method for spoofing detection that exploits the spatial diversity offered by the antenna array by comparing the carrier phases at a pair of elements (Domínguez, et al., 2015). In a real GNSS scenario, the different satellites and their movements lead to nonidentical phases. However, in the presence of a spoofer at a fixed position, all the spoofed satellites will exhibit the same carrier phase.

Based on this, the algorithm computes the carrier phase difference between two antennas for all the available satellites as,

$$\Delta CP[n]_{(s,v)} = |\Delta\theta^s[n] - \Delta\theta^v[n]| \qquad \forall v \neq s \tag{7}$$

where the phase difference between antennas for satellite s, i.e., $\Delta\boldsymbol{\theta}^s[n]$, is computed as

$$\Delta\theta^s[n] = \left[\boldsymbol{\theta}^s[n]\right]_2 - \left[\boldsymbol{\theta}^s[n]\right]_1. \tag{8}$$

If the value $\Delta CP[n]_{(s,v)}$ is close to zero for any pair of satellites (s, v), a spoofer is detected. The main advantage of this technique is that it comes at no extra cost, since carrier phases are already available within the PLL tracking loop, making it particularly suitable for handheld receivers that require low-complexity algorithms.

2.2 Mitigation Techniques

The spatial diversity introduced by the antenna array can also be exploited for beamforming. The latter is the linear combination of signal samples received at each of the antennas with the aim of concentrating the array beam towards a certain DoA, and filtering out all contributions arriving from different directions. This entails applying a set of coefficients or weights to the signal in (1), as

$$y[n] = \boldsymbol{w}^H \boldsymbol{x}[n] \tag{9}$$

where $\boldsymbol{w} \in \mathbb{C}^{L\times 1}$ is the weight vector and $y[n]$ is the signal at the beamformer output. The computation of these weights is performed according to a given criterion, which defines the beamforming technique applied. In this work, the focus is on two algorithms that surpass the performance of traditional beamforming techniques in the presence of correlated replicas: the HYB and the PBC.

2.2.1 Hybrid (HYB) Beamformer

The HYB beamformer is a weighted linear combination of the MVDR or Capon beamformer (Capon, 1969) and the temporal or LMSE (Konovaltsev, Antreich, & Hornbostel, 2007), therefore exploiting both spatial and temporal references (Seco-Granados, Fernández-Rubio, & Fernández-Prades, 2005). The use of both references, the DoA of the LoS contribution and a replica of the desired signal, makes it a particularly robust technique, capable of offering promising mitigation levels even for correlated replicas, a case where the MVDR alone fails (Widrow, Duvall, Gooch, & Newman, 1982; Zoltowski, 1988) The weight vector for the HYB beamformer is computed according to

$$\boldsymbol{w}_{\mathrm{HYB}}[n] = \alpha_0^* \boldsymbol{w}_{\mathrm{LMSE}}[n] + (1 - \alpha_0^* \boldsymbol{a}^{\boldsymbol{H}}(\theta_0)\boldsymbol{w}_{\mathrm{LMSE}}[n])\boldsymbol{w}_{\mathrm{MVDR}}[n], \tag{10}$$

where the weights for the LMSE beamformer $\boldsymbol{w}_{\mathrm{LMSE}}$ are computed as

$$\boldsymbol{w}_{\mathrm{LMSE}}[n] = \boldsymbol{R}_{\boldsymbol{x}}^{-1}[n]\, \boldsymbol{r}_{xs}[n], \tag{11}$$

where $\boldsymbol{R}_{\boldsymbol{x}}[n] = \boldsymbol{x}[n]\boldsymbol{x}^H[n]$ is the spatial covariance matrix and $\boldsymbol{r}_{\boldsymbol{xs}}[n] = \boldsymbol{x}[n]s^*[n]$ is the cross-correlation between the received signal and the local replica $s[n]$. The weights for the MVDR are given by

$$\boldsymbol{w}_{\mathrm{MVDR}}[\mathrm{n}] = \frac{\boldsymbol{R}_{\boldsymbol{x}}^{-1}[n]\boldsymbol{a}(\theta_0)}{\boldsymbol{a}^{\boldsymbol{H}}(\theta_0)\boldsymbol{R}_{\boldsymbol{x}}^{-1}[n]\boldsymbol{a}(\theta_0)}. \tag{12}$$

As shown in Equation (10), the weighting factor controlling the linear combination of the LMSE and MVDR techniques is the amplitude of the LoS signal. Although this value is unknown, the HYB provides the following closed-form expression for its estimation, allowing the direct implementation of the technique:

$$\hat{\alpha}_0[n] = \boldsymbol{w}_{\mathrm{HYB}}^H[n]\boldsymbol{r}_{\boldsymbol{xs}}[n] \tag{13}$$

By replacing Equation (13) into Equation (10), the final solution of the HYB beamformer weights is obtained. Note that to estimate the amplitude in Equation (13), the weights vector is required, showing the iterative nature of the algorithm. The weights computed from the n-th sample are used to estimate the amplitude $\hat{\alpha}_0[n]$, which in turn is used to compute the next weights at sample $n+1$.

2.2.2 Power-Based Capon (PBC) Beamformer

As briefly mentioned in the previous section, the MVDR beamformer fails to effectively mitigate contributions that are highly correlated with the desired signal, resulting in signal cancellation (Mañosas-Caballú, Swindlehurst, & Seco-Granados, 2021). To circumvent this effect, the PBC removes the terms responsible for it from the spatial covariance matrix and subsequently applies the same principle as the MVDR, but over a modified version of the covariance matrix.

Assuming the received signal follows the model in Equation (1), the spatial covariance matrix is given by

$$\begin{aligned}\boldsymbol{R}_{\boldsymbol{x}}[n] = {\alpha_0}^2 P_s \boldsymbol{a}(\theta_0)\boldsymbol{a}^H(\theta_0) + {\alpha_m}^2 P_m \boldsymbol{a}(\theta_m)\boldsymbol{a}^H(\theta_m) + \rho_{0m}[n]\boldsymbol{a}(\theta_0)\boldsymbol{a}^H(\theta_m)\alpha_0\alpha_m + \rho_{m0}[n]\boldsymbol{a}(\theta_m)\boldsymbol{a}^H(\theta_0)\alpha_0\alpha_m \\ + \boldsymbol{R}_{\boldsymbol{n}}[n],\end{aligned} \tag{14}$$

where ${\alpha_0}^2 P_s \boldsymbol{a}(\theta_0)\boldsymbol{a}^H(\theta_0)$, ${\alpha_m}^2 P_m \boldsymbol{a}(\theta_m)\boldsymbol{a}^H(\theta_m)$, and $\boldsymbol{R}_{\boldsymbol{n}}[n]$ denote the LoS, multipath and noise contributions, respectively, and $\rho_{0m}[n]\boldsymbol{a}(\theta_0)\boldsymbol{a}^H(\theta_m)\alpha_0\alpha_m + \rho_{m0}[n]\boldsymbol{a}(\theta_m)\boldsymbol{a}^H(\theta_0)\alpha_0\alpha_m$ is the cross-correlation between the LoS signal and the multipath component.

Since the objective of the MVDR in Equation (12) is to minimize the power at the output of the beamformer subject to a distortionless constraint in the DoA of the LoS signal, i.e., $\boldsymbol{w}^H\boldsymbol{a}(\theta_0) = 1$, this cross-correlation term can be combined with the direct signal contribution to reduce the output power, causing LoS cancellation.

To avoid this, the PBC estimates the cross-correlation term as

$$\begin{aligned}\boldsymbol{r}_{\boldsymbol{xs}}[n] = P_s\alpha_0^2\boldsymbol{a}(\theta_0) + \rho_{m0}[n]\boldsymbol{a}(\theta_m)\alpha_m\alpha_0 \implies \\ \rho_{m0}[n]\boldsymbol{a}(\theta_m)\boldsymbol{a}^H(\theta_0)\alpha_m\alpha_0 = \big(\boldsymbol{r}_{\boldsymbol{xs}}[n] - P_s\alpha_0^2\boldsymbol{a}(\theta_0)\big)\boldsymbol{a}^H(\theta_0),\end{aligned} \tag{15}$$

where both spatial (θ_0) and temporal ($s[n]$) references are used. Furthermore, the power of the LoS signal also needs to be known by the receiver to implement the technique. With the expression in Equation (15) available, it can be used to remove the cross-correlation contributions from the spatial covariance matrix, yielding the following modified matrix used by the PBC to compute the weights:

$$\boldsymbol{R}_{\mathrm{PBC}}[n] = \boldsymbol{R}_{\boldsymbol{x}}[n] - \big(\boldsymbol{r}_{\boldsymbol{xs}}[n] - P_s\alpha_0^2\boldsymbol{a}(\theta_0)\big)\boldsymbol{a}^H(\theta_0) - \boldsymbol{a}(\theta_0)\big(\boldsymbol{r}_{\boldsymbol{xs}}[n] - P_s\alpha_0^2\boldsymbol{a}(\theta_0)\big)^H. \tag{16}$$

The weight vector is therefore obtained as

$$\boldsymbol{w}_{\mathrm{PBC}}[n] = \frac{\boldsymbol{R}_{\mathrm{PBC}}^{-1}[n]\boldsymbol{a}(\theta_0)}{\boldsymbol{a}^H(\theta_0)\boldsymbol{R}_{\mathrm{PBC}}^{-1}[n]\boldsymbol{a}(\theta_0)} \tag{17}$$

which corresponds to the MVDR formulation using the modified spatial covariance matrix, $\boldsymbol{R}_{\mathrm{PBC}}[n]$.

2.3 Simulation Results

An initial performance analysis of the proposed techniques was conducted using simulations. The evaluation considered urban and suburban scenarios, with the number of antennas limited to two and four, representative of typical handheld receivers. Particularly, in the two-antenna case a uniform linear array (ULA) is used, whereas in the four-antenna configuration a uniform rectangular array (URA) is assumed. The results presented in this work are those obtained in an urban setting, which represents the most challenging

case, as undesired contributions are expected to arrive from directions close to that of the LoS signal. Analyzing this scenario provides valuable insight into the overall behavior of the techniques, allowing the performance in the suburban context to be extrapolated.

The simulation parameters used to test the urban scenario are shown in Table 2-1. Note that the angular difference between LoS and the replica is 20°, which is small given the use of only two and four antennas. In contrast, the difference in carrier-to-noise density ratio (C/N_0) is considerably large, 10 dB, reflecting the attenuation expected under urban propagation conditions.

Table 2-1: Configuration for urban scenario

	LoS Signal	Multipath/Spoofing
C/N_0	45 dB-Hz	35 dB-Hz
θ	80°	60°

2.3.1 Performance of Detection Algorithms

In order to characterize the performance of the SC-MA technique, the detection time is analyzed. This metric reflects the time required by the algorithm to identify the presence of the undesired contribution. Accordingly, the goal is to minimize the time of detection so that mitigation techniques can be triggered as quickly as possible.

In Figure 2-1, the detection time of the SC-MA as a function of the multipath delay is shown for the two- and four-antenna configurations assuming the parameters in Table 2-1. For the sake of comparison, the performance of the Multi-Correlation Multi-Antenna (MC-MA) technique is also presented, whose main difference with respect to the SC-MA is that it uses multiple correlation points instead of just the prompt. Despite the robustness provided by using additional points, the performance in the presence of closely spaced replicas, i.e, those with small relative delays, is better for the SC-MA than for the MC-MA. However, as the multipath delay increases beyond, approximately, 0.5 chips, this behavior is reversed, with the MC-MA outperforming the SC-MA.

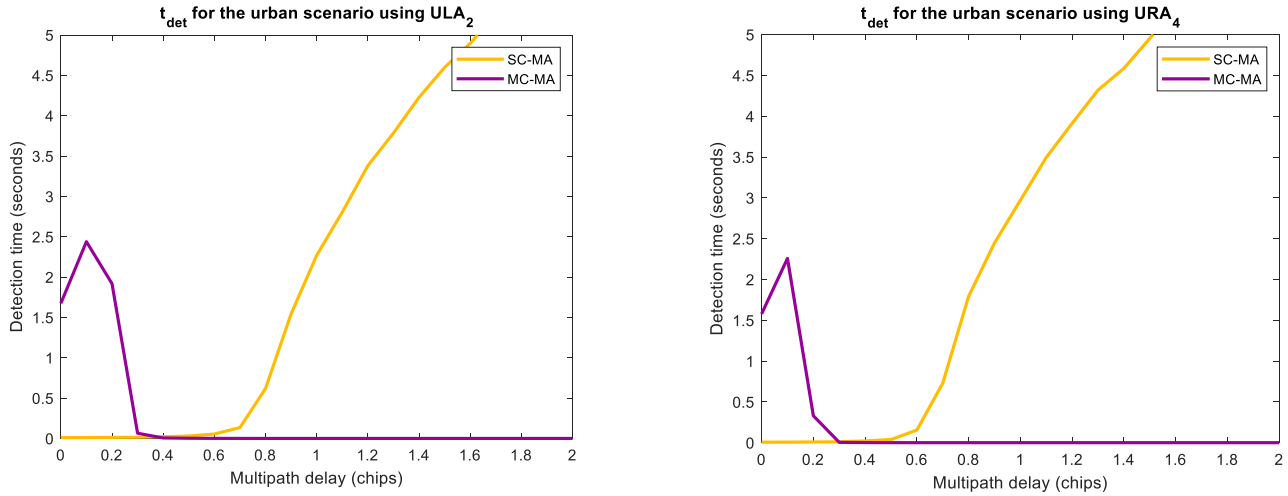

(a) ULA with two antennas. **(b) URA with four antennas.**

Figure 2-1: Detection time of SC-MA and MC-MA in urban scenario for the case of (a) two antennas and (b) four antennas.

Since highly correlated replicas have a greater impact on the performance of GNSS receivers, it is preferable to ensure a quick detection for small delays. For larger multipath delays, the level of correlation between the LoS signal and the reflection decreases, leading to smaller errors in the positioning solution. Based on this, and according to the results presented, the SC-MA is selected over the MC-MA for further analysis.

Regarding the CPD technique for spoofing detection, it achieves minimal detection time, providing nearly instantaneous identification. This promising performance results from the large changes in the metric when spoofed signals are present. In Figure 2-2, the carrier phase difference between satellites is shown for both spoofed and non-spoofed signals. As expected, the phase

differences for the real signals are nonzero, with values ranging from 1.8 to 3.8 radians. In contrast, when the satellites are spoofed, the phase difference is very close to 0 radians, as shown in Figure 2-2, where the differences for all the spoofed satellites are around 0.05 radians, exhibiting minimal variance.

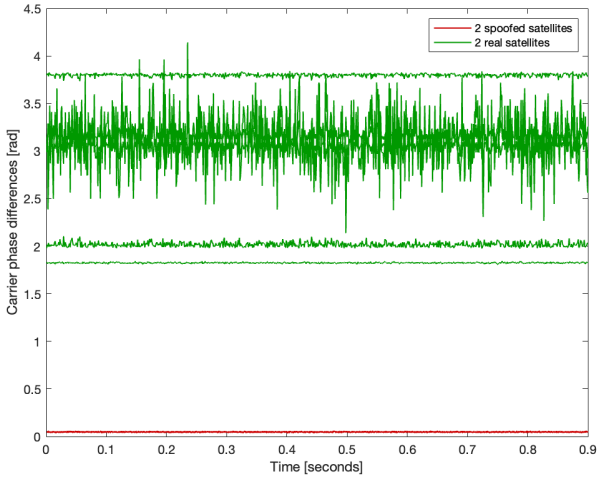

Figure 2-2: Comparative of carrier phase differences of pairs of spoofed and non-spoofed satellites.

2.3.2 Performance of Mitigation Algorithms

The performance of the mitigation algorithms is measured according to three metrics. The first one is the mitigation capacity, defined as the ratio between the array response at the DoA of the undesired contribution and the array response at the DoA of the LoS signal:

$$\text{Mitigation Capacity [dB]} = 10\, log_{10}\left(\frac{|\boldsymbol{w}^H \boldsymbol{a}(\theta_m)|^2}{|\boldsymbol{w}^H \boldsymbol{a}(\theta_0)|^2}\right). \tag{18}$$

Moreover, to fully characterize the impact of the techniques on the GNSS receiver performance, the root mean square error (RMSE) of the Delay-Locked Loop (DLL) and Phase-Locked Loop (PLL) is analyzed. These metrics, obtained from the output of the DLL and PLL, reflect the influence of the remaining contribution after mitigation on the tracking loops and, ultimately, on the final positioning solution. For these simulations, a normalized dot product DLL discriminator with an early-late spacing of 1 chip was implemented, along with a four-quadrant arctangent discriminator for the PLL.

For comparison purposes, in addition to the proposed techniques, several common beamforming techniques were also implemented and analyzed. Specifically, the LMSE/temporal (TEM) beamformer in Equation (11), the MVDR/Capon (CAP) algorithm in Equation (12), and the phased array (PHA) were evaluated together with the HYB and PBC. The results for these techniques in terms of mitigation capacity as a function of the replica delay are shown in Figure 2-3 for the case of two and four antennas.

Note that, in both plots, the CAP beamformer performs poorly when the delay of the undesired contribution is small. Particularly, for relative delays below 0.4 chips, the receiver is not capable of tracking the LoS signal due to the cancellation effect, leading to loss-of-lock. For slightly larger delays, between 0.4 and 0.6 chips, the CAP algorithm still amplifies the replica, showing mitigation values greater than 0 dB, even though it does not fully cancel the LoS contribution. This clearly demonstrates the deficiencies of the CAP beamformer in the presence of correlated replicas. As expected, the PHA beamformer exhibits constant behavior across all replica delays, since its weights do not depend on the received samples and are fixed to steer the beam toward the DoA of the LoS. The TEM, HYB and PBC algorithms show similar mitigation capabilities, attenuating the undesired component even for small relative delays. The performance difference among these three are revealed when the DLL and PLL RMSE are examined.

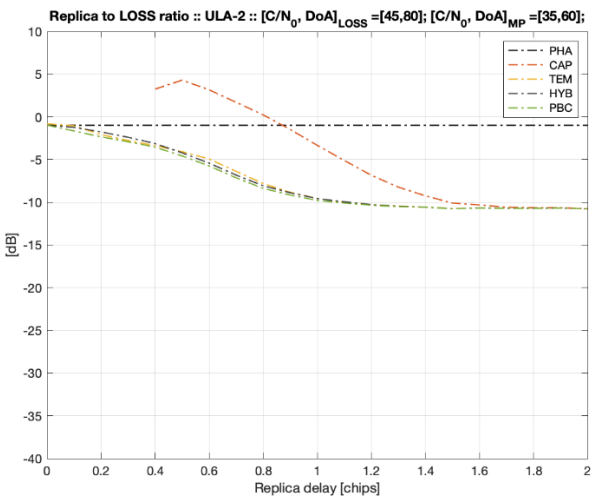

(a) ULA with two antennas

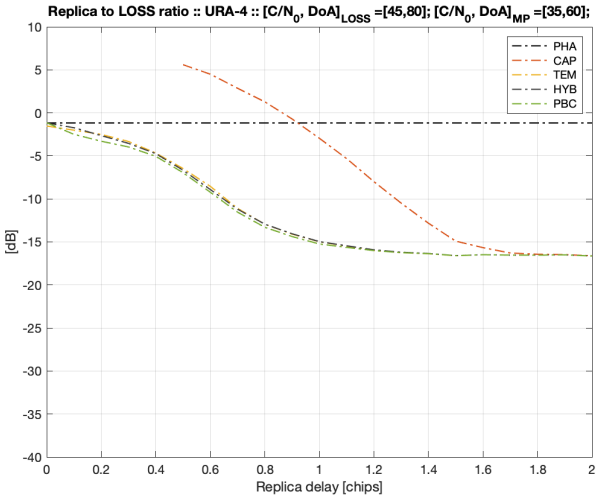

(b) URA with four Antennas

Figure 2-3: Mitigation capacity in urban scenario for the case of (a) two antennas and (b) four antennas.

The DLL and PLL RMSE are presented in Figure 2-1Figure 2-4 for both antenna configurations. The DLL performance of the three techniques, TEM, HYB and PBC, is quite promising, with the TEM slightly outperforming the other two for small relative delays, below 0.4 chips. However, the PLL RMSE of the latter is considerably larger than that of the other techniques and, although not explicitly shown in the results, such a large PLL error causes loss-of-lock, preventing position calculation. The reason for the large PLL RMSE in the TEM beamformer is that this technique does not use a spatial reference in the computation of the weights, only a temporal one. Consequently, there is no constraint in the DoA of the LoS signal, unlike the HYB and PBC, and thus the weights for the TEM introduce a random phase in the LoS direction each time they are updated. This random phase offset from the weights cannot be tracked by the PLL, resulting in loss-of-lock.

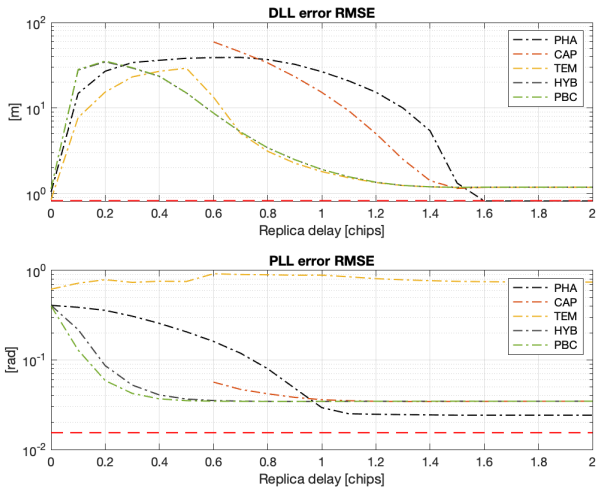

(a) ULA with two antennas

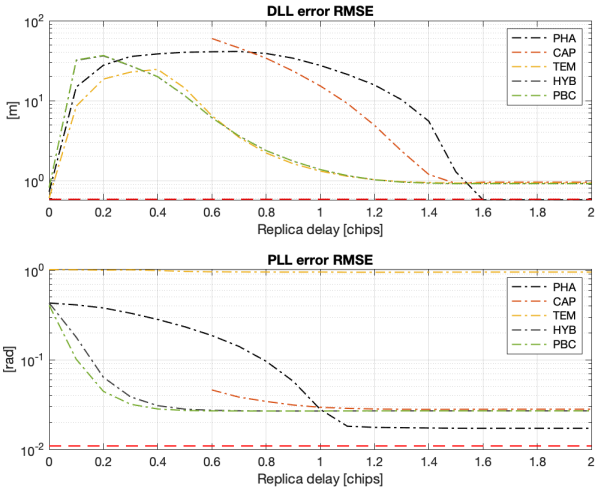

(b) URA with four antennas

Figure 2-4: DLL and PLL error in urban scenario for the case of (a) two antennas and (b) four antennas.

Overall, the analysis shows that the HYB and PBC beamformers exceed the performance of more conventional approaches, such as the CAP and TEM. Furthermore, it should be noted that, although using four antennas clearly increases the mitigation capacity of the techniques, the use of two antennas already provides partial attenuation of the undesired signals, demonstrating the suitability of the proposed techniques even in small antenna arrays.

3 LABORATORY SETUP

To validate the performance observed in the simulations, the proposed techniques were analyzed in laboratory settings using synthetic signals, providing a first validation of the algorithms and offering a comparison with mass-market GNSS receivers. To assess the performance, different scenarios were processed considering the two previously described impairments, multipath and spoofing, and different scenario conditions such as urban, semi-urban and open sky, along with single and multiple spoofers. In this section, processing of the synthetic scenarios is detailed.

3.1 Baseline Scenario

A main static scenario was employed to generate synthetic signals in the radio frequency constellation simulator (RFCS), which served as the reference scenario to be subjected to multipath and spoofing. The total number of antennas available was set to four, of which either two or four were used depending on the test specification.

Figure 3-1 shows the skyplot of GPS and Galileo constellations used for signal generation, which were the L1CA and E1C signals, respectively. In Figure 3-2 the array configuration is presented, where the antennas are arranged in a URA with a minimum inter-element spacing of 13 cm, half the L5 wavelength. These two figures define the baseline configuration for the test cases described in the following sections. After the injection of the external disturbances (multipath time series) or the addition of spoofed signals in the baseline component generated by the RFCS, the overall signal is processed by the receiver, and the performance is analyzed.

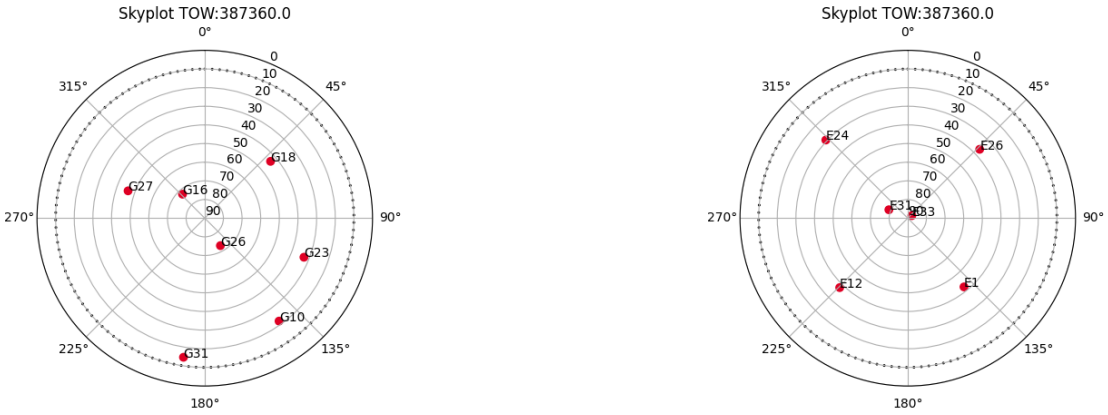

(a) GPS constellation **(b) Galileo constellation**

Start-time: 2022-05-26 11:37:00

Figure 3-1: Baseline Scenario Skyplot (a) GPS and (b) Galileo.

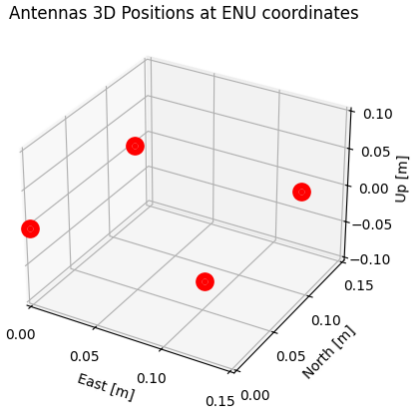

Figure 3-2: Antennas arrangement of baseline scenario.

3.1.1 Multipath Scenarios Description

For the laboratory tests, several multipath scenarios were generated using the Orolia Skydel-SDX RFCS, along with an ESA multipath plugin, previously developed by GMV in the framework of the SDXPAND project. The multipath time series was injected into the baseline scenario signal described in the previous section, and the generated IQ files were directly processed in the software receiver XRC.

The time series used in SDXPAND for multipath injection were generated using the Land Mobile Multipath Channel Model V3.0 tool, standardized in ITU-R P.681-7 (10/09) "Propagation data required for the design of Earth-space land mobile telecommunication systems". In this tool, the user can configure either an urban or suburban environment, automatically setting the desired parameters for each type of scenery, being these the height and size of the trees, the size of the road, the number and height of the buildings, etc. Based on this, three pedestrian scenarios were created: urban, sub-urban and dense foliage environments, and for each one, the detection and mitigation techniques were applied using two and four antennas, leading to a total of six multipath test cases.

In Table 3-1, the general configuration for the multipath test cases is presented, and Table 3-2 presents the parameters used for the detection and mitigation techniques analyzed.

Table 3-1: Multipath configurations for Test Cases

Simulation duration [s]	Multipath Start time [s]	Multipath duration [s]
600	120	360

Table 3-2: Detection and mitigation configurations for multipath Test Cases

Detection Technique	Detection Parameters	Mitigation Technique	Mitigation Parameters	Output Metrics
SC-MA	$\mu_0 = 12$ $\mu_1 = 36$ $h = 70:\ 2\ antennas$ $h = 100:\ 4\ antennas$	HYB	K=0.01	Tracking Accuracy PVT Accuracy Multipath Detection Capability

3.1.2 Spoofing Scenarios Description

In order to generate the Spoofing scenarios, the Orolia Skydel-SDX RFCS was also used. In this scenario, the spoofed signals are added to the baseline signal, producing a single IQ file containing both real and spoofed contributions. The test cases considered both single and multiple spoofing sources. In the case of multiple spoofers, the transmissions of the visible satellites were distributed among four spoofers. The spoofing sources were generated with geometrical dispositions relative to the simulated receiver using the interference plugin of the RFCS. In Figure 3-3 the location of the spoofers and the real and spoofed receiver position is shown.

Figure 3-3: Spoofing Locations.

To characterize the spoofing detection and mitigation techniques in the laboratory setting, a total of four test cases were generated, evaluating the algorithms with two and four antennas, for single and multiple (four) spoofing transmitters. In Table 3-3, the configuration for the spoofing simulation is shown, in Table 3-4, the parameters for the detection and mitigation algorithms are presented.

Table 3-3: Spoofing simulation configurations.

Simulation Start Time	Simulation duration [s]	Spoofing Start Time [s]	Spoofing duration [s]	Spoofing Signal Power [dB]
2022-05-26 11:37:00	600	60	240	Receiver Signal Power [dB] + 4 dB

Table 3-4. Detection and mitigation configurations for spoofing Test Cases

Detection Technique	Detection Parameters	Mitigation Technique	Mitigation Parameters	Output Metrics
CPD	$h = 0.015$ $N_{corr} = 400$ $S = 5$	PBC	K=0.04	Tracking Accuracy PVT Accuracy Spoofing Detection Capability

3.2 Multipath Results

The multipath detection results using synthetic signals are shown in the Table 3-5, for all the environments:

Table 3-5: Average of multipath detection metrics across all satellites.

Multipath	Number of antennas	Detection [%]	False Detection [%]
Urban	2	93.9	9.5
	4	96.5	12.4
Suburban	2	90.6	12.9
	4	96.5	14.0
Dense Foliage	2	94.7	10.1
	4	98.0	12.0

With two antennas, the detection technique (SC-MA) reached accuracies between 90 and 98%, with a maximum false detection of 14%. The main source of detection error occurred when signal tracking was intermittently lost and reacquired, resembling conditions such as periodic blockage in multipath environments. These fluctuations prevented the accumulative detection metric from reaching the threshold within its reset time, leading to missed detections. The same effect was reflected in the C/N₀, which oscillated between strong and weak values, highlighting the difficulty of reliable detection under rapidly varying signal conditions. Another factor that contributed to the detection faults was due to GPS satellites, where the highest rate of false detections occurred before multipath was present, suggesting that the multipath detection threshold may be too low for GPS signals. However, considering the overall performance with both two and four antennas and both constellations, the chosen threshold can be regarded as well tuned and capable of providing reliable results.

In the laboratory campaigns, multipath mitigation was also evaluated in urban, suburban, and dense foliage scenarios using two- and four-antenna configurations. At the tracking level, the main benefit of the beamforming technique, the HYB beamformer, was seen in the recovery of C/N₀. With two antennas, modest improvements were observed, typically around +1 dB compared to single-antenna processing, while four antennas consistently achieved increases of 2–3 dB, effectively doubling the available signal power. This gain confirmed that beamforming was correctly reinforcing the LoS component while attenuating replicas.

While the gains were scenario-dependent, the techniques consistently enhanced signal strength and stability, laying the foundation for more accurate positioning solutions. In Table 3-6, Table 3-7, and Table 3-8, the code delay jitter (third order difference) results for urban, suburban and dense foliage environments are shown, respectively. Furthermore, the results obtained in each case with the reference mass-market receiver used to assess the added value of the techniques are also presented.

Table 3-6: Code Jitter for Urban scenario signals.

Scenario	Number of Points for Analysis	95th Percentile [m]
Clean Signal - GPS	2392	1.10
Clean Signal - GAL	1794	0.84
Technique Off – GPS	1046	3.46
Technique Off – GAL	1333	2.52
2 Antenna Mitigation - GPS	942	3.37
2 Antenna Mitigation – GAL	1348	2.46
4 Antenna Mitigation - GPS	991	3.27
4 Antenna Mitigation - GAL	1362	2.36
U-blox – GPS (L1 C/A)	940	1.34
U-blox – GAL (E1B+E1C)	1133	0.97

Table 3-7. Code Jitter for suburban scenario signals.

Scenario	Number of Points for Analysis	Standard Deviation [m]	95th Percentile [m]
Clean Signal - GPS	2392	0.4	1.1
Clean Signal - GAL	1794	0.5	0.8
Technique Off – GPS	1358	1.2	2.8
Technique Off – GAL	1358	1.1	2.2
2 Antenna Mitigation - GPS	1337	3.6	3.4
2 Antenna Mitigation – GAL	1359	1.5	1.6
4 Antenna Mitigation - GPS	1411	3.8	3.4
4 Antenna Mitigation - GAL	1366	1.5	1.6
U-blox – GPS (L1 C/A)	1412	0.6	1.3
U-blox – GAL (E1B+E1C)	1286	0.3	0.7

Table 3-8 Code Jitter for dense-foliage scenario signals.

Scenario	Number of Points for Analysis	Standard Deviation [m]	95th Percentile [m]
Clean Signal - GPS	2392	0.4	1.1
Clean Signal - GAL	1794	0.5	0.8
Technique Off – GPS	1519	1.5	1.7
Technique Off – GAL	1028	1.6	2.9
2 Antenna Mitigation - GPS	1513	1.5	1.6
2 Antenna Mitigation – GAL	1078	1.5	3.4
4 Antenna Mitigation - GPS	1523	1.3	1.5
4 Antenna Mitigation - GAL	1026	1.6	3.0
U-blox – GPS (L1 C/A)	1875	0.3	0.5
U-blox – GAL (E1B+E1C)	864	0.7	1.3

First, it can be seen from the “Technique Off” performance that sub-urban scenarios were the less harmful, and dense foliage the most harmful regarding multipath robustness. For the mitigation performance against multipath using two antennas, slight improvements are seen in the code delay estimations compared to the technique off scenario in urban and dense foliage. On the other hand, with four antennas, further improvements are observed, especially in dense foliage scenario. For the suburban scenario the

gains of the technique were seen in the increase on the number of tracking points with four antennas, that are reflected in the positioning results presented in Table 3-9.

Overall, despite large improvements in C/N$_0$, code jitter statistics showed limited improvements when averaged across all satellites, from two antennas to four antennas, as many signals were either fully blocked or minimally affected by multipath, reducing the observable impact of mitigation. However, when analyzing individual satellites, clear gains in code tracking accuracy were identified, particularly in cases with moderate multipath.

Regarding the reference receiver, U-blox tracking accuracy remained high and comparable to clean signal levels, even though multipath still affected its performance. This accuracy came with slightly lower robustness (number of points), and did not fully translate into reliable positioning, as will be seen next.

At positioning level, the mitigation delivered a more accurate and stable position, close to clean-signal performance and outperforming the reference receiver, as can be seen in Table 3-9. While overall accuracy per satellite set may appear unchanged, improvements in individual satellite tracking can significantly enhance the final positioning solution. It was also found that further gains could be achieved by incorporating mitigation at the acquisition stage or avoiding signal acquisition when multipath is detected, preventing unreliable measurements from being used in the position, velocity and time (PVT) computation.

Table 3-9: Multipath positioning results.

Urban	Suburban	Dense Foliage
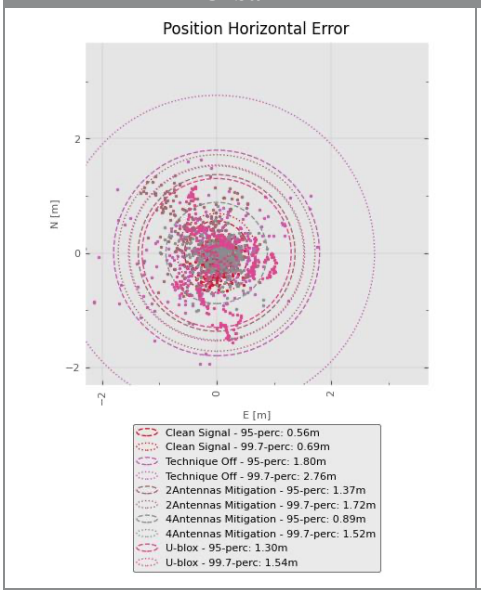	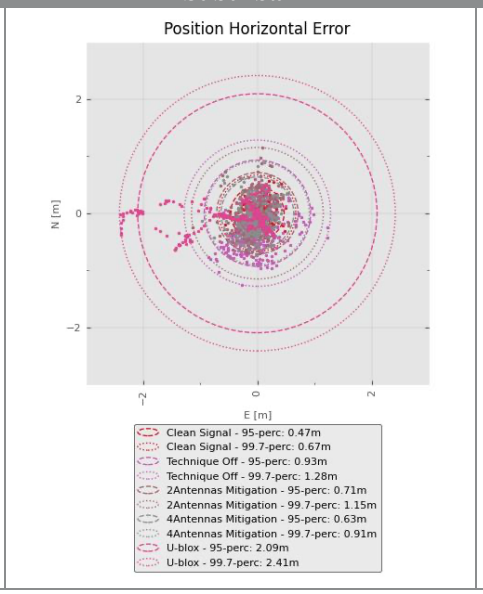	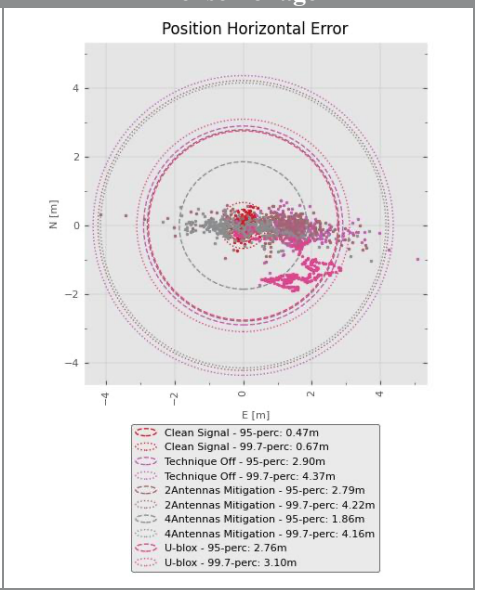

3.3 Spoofing Results

The metric used to detect spoofing is unique for all the satellites, since all spoofed satellites arrive from the same direction. In Figure 3-4 the spoofing metric (S) during the simulation for a single spoofer is shown. A threshold value of S =5 was used as the minimum number of satellite pairs required to consider a spoofing event, for both single and multiple spoofers.

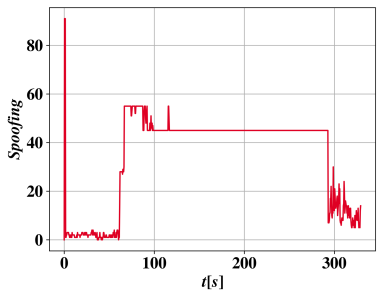

(a) Detection with two antennas

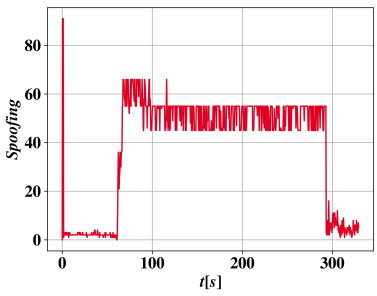

(b) Detection with four antennas

Figure 3-4: Pairs of spoofed satellites for single spoofer with (a) two and (b) four antennas.

Once the spoofed signals disappeared, some satellites could not resume tracking of the real signals because their estimated parameters were too far from the true values. As a result, the receiver lost lock, contaminating the detection metric and preventing reliable results at the end of the simulation. This is the main reason why the S parameter does not return to the values before spoofing appeared.

In Table 3-10, the detection results for the different spoofing scenarios processed are presented:

Table 3-10. Detection Statistics for two and four antennas.

Conditions	Detection [%]	Miss Detection [%]	False Detection Full Test [%]	False Detection Excluding post-Spoofing [%]
2 antennas – Single Spoofer	99.6	0.4	36.5	2.0
4 Antennas – Single Spoofer	99.6	0.4	14.3	
2 antennas – Multiple Spoofer	97.5	2.5	4.4	
4 antennas – Multiple Spoofer	98.8	1.2	2.6	

The spoofing detection technique, the CPD, showed great performance, with a probability of detection superior to 98% on single and multiple sources test cases, for both two and four antenna scenarios. Since the CPD algorithm uses only two antennas, the extra spatial diversity achieved with four antennas did not improve the detection performance.

Regarding spoofing mitigation, the PBC algorithm was used. When spoofing is active, and the tracking is locked on the spoofed peak, the mitigation is only observed through a decrease in the C/N_0. This occurs because, since the receiver is tracking the spoofed signal, once it is mitigated, a lower power is observed in the tracking estimate. Spoofing mitigation was only possible with four antennas, at tracking level, as seen in Table 3-11 with the decrease of the C/N_0.

Table 3-11: Average C/N_0 across satellites, at spoofing scenarios.

C/N_0 Average	Mitigation Off – 1 Antenna [dB-Hz]	Mitigation On – 2 Antennas [dB-Hz]	Mitigation On – 4 Antennas [dB-Hz]
Single Spoofer	47.00	46.29	35.46
Multiple Spoofer	45.07	44.60	39.10

While spoofed signals can be attenuated, switching back to the real GNSS signal requires mitigation to be active during both acquisition and tracking. This can be seen in the positioning results presented in Table 3-12, where techniques are unable to return their solution to the correct receiver location. In multi-spoofer setups, position estimation collapses due to incoherent measurements between different spoofed signals. In order to spoof the receiver with a multiple spoofer transmission setup, it is necessary to properly synchronize the different delays considering the different distances between spoofer transmitter and receiver location.

Table 3-12: Spoofing positioning results.

Single Spoofer	Multiple Spoofers
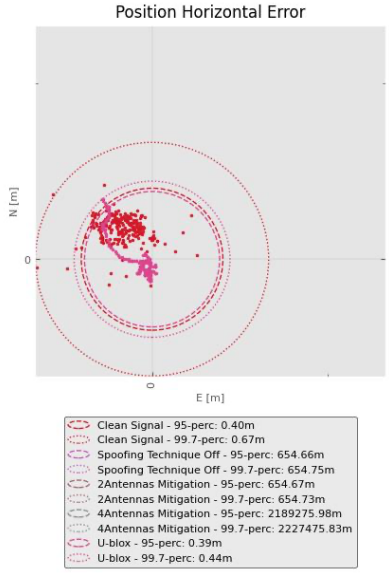	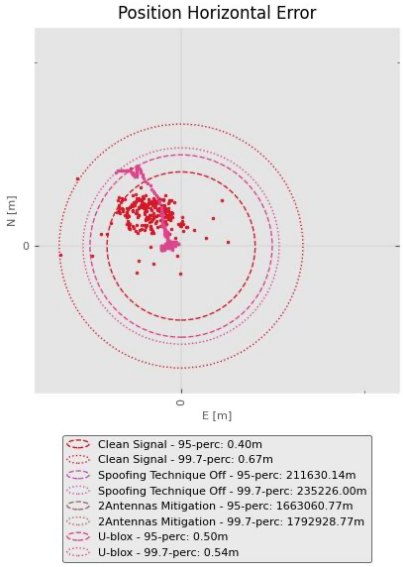

4 FIELD TESTS

To finalize the validation of the multipath mitigation technique, the HYB beamformer, real signals were processed to demonstrate the full end-to-end process, validating the technique under a real-life scenario. For this test, the HYB beamformer was applied to GNSS signals captured by ESA in the Netherlands under open-sky conditions. In this case, no undesired contributions were present; instead, the technique was used to validate the proposed antenna calibration and evaluate its performance with only LoS signals.

4.1 Setup Description

The recordings were conducted using four antennas arranged in a URA configuration, with an inter-element spacing of half the L1 wavelength, and the array was connected to 4 - NI VST SDR for the data recording. This setup configuration is depicted in Figure 4-1, and information about the data capture, such as the reference antenna location and recording date, is presented in Table 4-1.

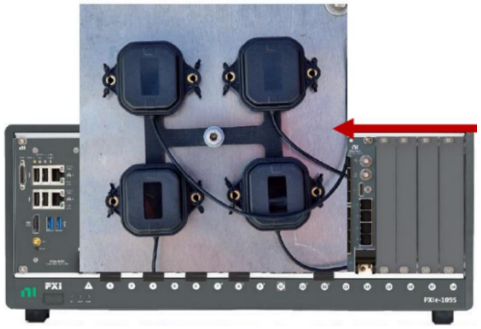

Figure 4-1: ESA Capture Set up (U-Blox antennas and NI VST SDR).

Table 4-1: Location and Date of open-sky recordings by ESA.

Location	Date
• Latitude: 52.2195N • Longitude: 4.4231E • Height: 52m	2024-08-01T14:37:29

In Table 4-2, the front-end configuration for the signal recording is shown. In the software receiver, the acquisition and tracking/technique configurations were the same as in the laboratory tests already presented in Section 3.

Table 4-2: ESA open-sky recording main configurations.

Parameter	Value
Signal duration	300 s
Sampling frequency	40 MSPS
Number of bits	16
Band recorded	L1
Signals Processed	E1C, L1C/A

4.2 Antennas Calibration

In the field tests, antenna calibration was required to compensate hardware delays and ensure phase coherence across the array. Calibration was carried out in open-sky conditions by computing the assisted phase error of each antenna relative to a reference. Ideally, this error should be constant across satellites, but large variations were observed, confirming non-negligible hardware-induced delays.

To correct the average phase between antennas, a phase offset was applied to antennas 2, 3, and 4, in order to compensate different hardware delays previously measured. The given phase offsets were 1º, -4.5º and -4.5º for antenna 2, 3, and 4, respectively. In Figure 4-2, the phase error of antenna 2, 3, and 4 when using the estimations of the beamforming signal, with all antennas after calibration, is shown.

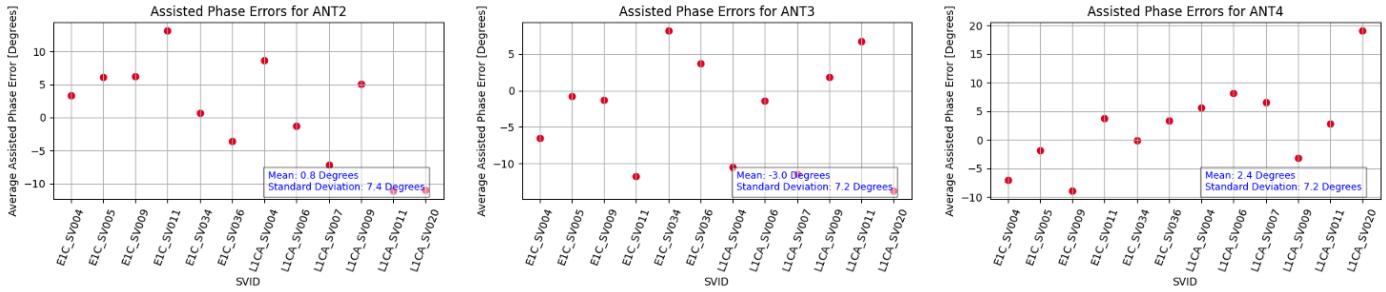

(a) Antenna 2 **(b) Antenna 3** **(c) Antenna 4**

Figure 4-2: Phase error of antenna 2, 3, and 4 after calibration.

The previous figure highlights the low mean, and low standard deviation of the phase error between antennas after calibration, which will allow the technique to successfully combine the different antennas measurements and provide a robust and precise tracking.

4.3 Results

After the calibration of each antenna, the HYB beamformer was applied using four antennas. In the tracking domain, the effect of the technique is reflected in the code delay jitter, as shown in Table 4-3.

Table 4-3: Code Jitter of ESA capture signals.

Scenario	Number of points for Analysis	95th Percentile [m]
Technique Off – GPS	1428	2.1
Technique Off – GAL	1452	1.5

Scenario	Number of points for Analysis	95th Percentile [m]
2 Antenna Mitigation - GPS	1416	1.2
2 Antenna Mitigation – GAL	1476	0.9
4 Antenna Mitigation - GPS	1428	1.1
4 Antenna Mitigation - GAL	1494	0.7

In terms of tracking performance enhancement, the application of the technique with two antennas results in a reduction of the 95th percentile code jitter to nearly half, with a similar improvement observed in the standard deviation. This indicates a significant gain achieved simply by transitioning from a single antenna to a dual-antenna configuration.

Four antennas added more tracking points but brought minimal further improvement due to open-sky environment, where the conditions were already near-optimal with two antennas. It is worth noting that the code jitter statistics of field tests using two antennas is aligned with the clean signal results of the laboratory tests, free of hardware defects.

Finally, the results at positioning level are presented in Figure 4-3.

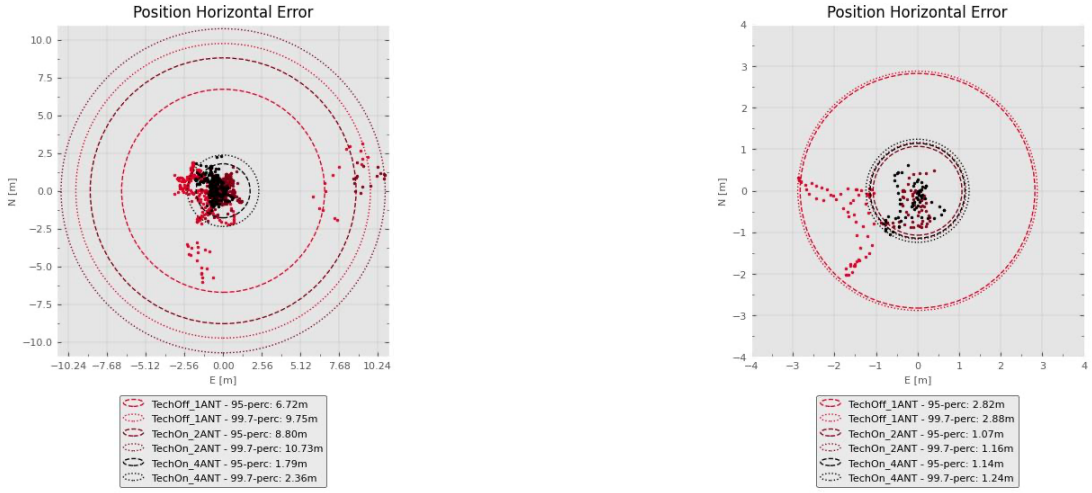

(a) Total simulation **(b) Last minute**

Figure 4-3: Positioning solution for ESA Capture for (a) total simulation and (b) last minute.

When the total simulation is considered, with two antennas, the average position error was reduced compared to the single-antenna case, from 2.56 m to 1.3 m. However, some statistics, such as the 95th percentiles, were slightly higher due to slower convergence in the early part of the test. The four-antenna setup provided the best results, with the mean for horizontal error reduced by roughly a factor of three relative to one antenna (0.87 m), and faster convergence to a stable solution thanks to improved satellite availability due to higher signal power. When considering only the final 60 s of data, both two- and four-antenna setups showed clear improvements over single-antenna processing in all metrics, with the four-antenna case achieving sub-meter accuracy and 100% availability. Overall, the results confirmed that multi-antenna configurations, especially with four elements, improve positioning accuracy and stability under open-sky conditions. This paves the way for future testing of the technique under multipath environments, to quantify the added value of beamforming and corroborate the results observed in the laboratory tests.

5 CONCLUSIONS

This paper presented the design, integration, and validation of multipath and spoofing detection and mitigation techniques within GMV's software receiver XRC. The work focused on algorithms suitable for handheld GNSS devices equipped with compact two- or four-antenna arrays, aiming to balance accuracy, computational complexity, and robustness. The results from laboratory and field tests confirmed the feasibility of the techniques an validated the performance observed in the numerical simulations, while also highlighting the practical requirements for successful deployment.

For multipath, the SC-MA detector achieved high detection rates in all simulated scenarios, with miss detections typically below 10%. The use of four antennas offered modest improvements compared to two, although false detections—particularly for GPS in noisy conditions—emphasized the need for adaptive thresholds tuned per modulation. Multipath mitigation using the HYB beamformer was clearly more effective with four antennas, consistently outperforming both the two-antenna setup and the commercial U-blox receiver. While overall tracking accuracy averaged across satellites showed limited gains, significant improvements were observed at the individual satellite level, where even a small number of enhanced measurements could meaningfully improve the final positioning solution. These results suggest further gains may be possible if mitigation is extended to the acquisition stage, or if acquisition of signals already identified as multipath is avoided.

Spoofing detection was addressed with the CPD algorithm, which proved highly effective, achieving detection probabilities above 98% for both single- and multi-spoofer scenarios. Importantly, two antennas provided the same performance as four, since the technique only uses two elements to compute the phase differences. Mitigation, in contrast, required a larger array: the PBC was only effective with four antennas, where it was able to substantially attenuate the spoofed signals. Despite this, the mitigation did not enable recovery of the authentic signal once the spoofed signal disappeared, as tracking loops could not reacquire the correct delays. To ensure seamless recovery, mitigation must be active during acquisition as well as tracking.

Antenna calibration emerged as a critical prerequisite for both multipath and spoofing mitigation. Calibration must accurately estimate and correct phase offsets among antennas, as residual errors directly limit performance, particularly in four-antenna arrays. Hardware setup also proved decisive: ESA's recordings, made with the NI VST SDR, which ensures channel coherency, produced stable results aligned with theoretical expectations.

Future work should prioritize fine-tuning detection thresholds per modulation, extending multipath mitigation to the acquisition stage, and ensuring spoofing mitigation operates during both acquisition and tracking to enable recovery of real signals. In addition, since the techniques have proven effective with current settings, it is important to explore the minimum viable configurations, specifically, the lowest usable number of bits and sampling rate. This will help optimize the solution for integration into low-SWAP (Size, Weight, and Power) platforms, such as handheld or embedded systems.

To conclude, this activity has demonstrated that handheld GNSS receivers with compact antenna arrays can effectively detect and mitigate multipath and spoofing threats. With careful calibration and appropriate hardware support, these techniques pave the way for more reliable and trustworthy GNSS positioning in mass-market devices operating in challenging environments.

6 ACKNOWLEDGMENTS

The work presented in this paper is fully funded by ESA in the DIVAL project. This paper contains the personal comments and beliefs of the authors and does not necessarily represent the opinion of ESA.

7 REFERENCES

Capon, J. (1969). High-resolution frequency-wavenumber spectrum analysis. *Proceedings of the IEEE*, 1408-1418.

Domínguez, E., J.M. López-Almansa, Seco-Granados, G., López-Salcedo, J., Egea-Roca, D., Aguado, E., . . . Boyero, J. (2015). Multi-antenna techniques for NLOS and spoofing detection using vehicular real signal captures in urban and road environments. *Proc. ION GNSS+*.

Egea-Roca, D., López-Salcedo, J. A., & G. Seco-Granados. (2014). Interference and multipath sequential tests for signal integrity in multi-antenna GNSS receivers. *Proc. IEEE SAM*.

Konovaltsev, A., Antreich, F., & Hornbostel, A. (2007). Performance assessment of antenna array algorithms for multipath and interferers mitigation. *2nd Workshop on GNSS Signals & Signal Processing, GNSS SIGNALS2007.* Noordwijk.

Mañosas-Caballú, M., Swindlehurst, A., & Seco-Granados, G. (2021). Power-based Capon beamforming: avoiding the cancellation effects of GNSS multipath. *Signal Processing*.

Seco-Granados, G., Fernández-Rubio, J., & Fernández-Prades, C. (2005). "ML estimator and hybrid beamformer for multipath and interference mitigation in GNSS receivers". *IEEE Trans. on Signal Proc.*

Widrow, B., Duvall, K., Gooch, R., & Newman, W. (1982). Signal cancellation phenomena in adaptive antennas: Causes and cures. *IEEE Transactions on Antennas and Propagation*, 469-478.

Zoltowski, M. (1988). On the performance analysis of the MVDR beamformer in the presence of correlated interference. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 945-947.