

Robust Interference Mitigation in GNSS Snapshot Receivers

Helena Calatrava¹ | Adrià Gusi-Amigó² | Floor Melman³ | Pau Closas¹

¹ Department of Electrical and Computer Engineering, Northeastern University, Boston, MA, USA

² Albora Technologies, London, UK

³ European Space Agency (ESA/ESTEC), Noordwijk, Netherlands

Correspondence

Helena Calatrava
Department of Electrical and Computer Engineering at Northeastern University, Boston, MA, USA.
Email: calatrava.h@northeastern.edu

Funding Information

This research was supported by the European Space Agency NAVISP Program under activity NAVISP-EL2-096 and by the National Science Foundation under awards ECCS-1845833 and CCF-2326559.

Abstract

The robust interference mitigation (RIM) framework offers a promising solution to jamming attacks on global navigation satellite system (GNSS) receivers. By identifying interfered samples as outliers in the selected transform domain, RIM operates without relying on jamming waveform assumptions. This paper adapts RIM for GNSS snapshot architectures, assessing the impact of low-bit quantization on receiver performance under continuous wave (CW) and chirp interference. Using simulated data, snapshot RIM demonstrates significant improvements, achieving gains of 10, 20, and 35 dB in detected satellites for 2-, 4-, and 8-bit quantization in the presence of CW jamming. We also analyze the effect of quantization on the effective jammer-to-noise ratio, waveform distortion, and robust variance estimation. An experiment with realistic recordings shows that snapshot RIM achieves a 20-dB gain in the carrier-to-noise ratio over a professional receiver. Finally, a 24-h specifications test supports the feasibility of RIM integration in snapshot receivers with a maximum time-to-first-fix increase of 0.31 s.

Keywords

Anti-jamming, median absolute deviation, robust statistics, signal quantization, snapshot positioning

1 | INTRODUCTION

Jamming interferences pose a threat to services relying on global navigation satellite system (GNSS) technology by causing reduced accuracy, increased time to first fix (TTFF), and loss of positioning information. Thus, the escalating proliferation of jamming sources in the past years has attracted special attention in the literature, with a substantial focus on jamming cancellation and mitigation techniques (Bamberg et al., 2023; Meng et al., 2022; Zidan et al., 2020). The online marketplace provides a diverse selection of jammers that are easily accessible, owing to their affordable prices and compact sizes (Garcia-Molina & Crisci, 2016). Whereas some individuals may use jamming devices with malicious intent against other GNSS users, the primary motivation for human-made jamming revolves around safeguarding personal privacy (Medina, Lass, et al., 2019). Additionally, GNSS receivers can be disrupted by legitimate waveforms, including signals emitted by distance-measuring equipment (DME) (Zhu et al., 2023) and tactical air navigation (TACAN) aiding systems (Pullen & Gao, 2012). Another example of unintentional jamming is the continuous wave (CW) interferences produced by damaged

electronics. The overcrowded radio frequency spectrum also enters into the picture, with signal harmonics from other wireless systems being present in the GNSS frequency bands (Gao et al., 2016). Although out-of-band interferences can also jeopardize GNSS receivers because of, for instance, saturation of the receiver circuitry and reciprocal mixing effects (Hegarty et al., 2020), this paper focuses on in-band interferences.

Adding to the complexity, jamming signals are received with remarkably higher power levels than useful GNSS signals, given the significant propagation losses caused by the high altitude of satellite orbits (Morales Ferre et al., 2019). As a consequence, service disruption areas can cover up to several square kilometers (Borio et al., 2016; Mitch et al., 2011). Although numerous anti-jamming algorithms have been developed to enhance GNSS performance, our understanding of how jamming signals impact the stability of conventional receivers remains limited. A study on the performance of baseline GNSS receivers with second-order tracking loops under different jamming interferences was conducted by Ding et al. (2023), who investigated the minimum jammer-to-signal ratio (J/S) causing a loss of lock. Moreover, Wu et al. (2023) showed that anti-jamming filters affect the performance of multipath suppression algorithms.

1.1 | Snapshot Positioning

Conventional GNSS receivers are designed to provide robust performance in the presence of interference and are commonly used in safety-critical applications. Alongside these, low-power GNSS receivers, including snapshot receivers, are increasingly essential for applications in which energy constraints are critical, such as internet-of-things (IoT) sensors, wearable devices, and automated vehicles operating in challenging urban environments with limited power resources and intermittent connectivity (Grenier et al., 2023). As the operating principles and capabilities of GNSS snapshot receivers differ from those of conventional architectures, there is a need to study the suitability of these receivers for interference mitigation. A snapshot is defined as a digitized GNSS signal captured over only a few milliseconds (Fernández-Hernández & Borre, 2016). Snapshot processing, which may also be referred to as cloud-GNSS or server-based processing (Van der Merwe et al., 2019), provides an instantaneous position fix from a single snapshot. With this small amount of data, signal tracking, symbol decoding, and navigation message extraction are not possible, and the snapshot position, velocity, and time (PVT) solution relies solely on the acquisition results with external ephemeris information. Snapshot positioning strategies are known for their low power consumption because all of the processing steps, except for snapshot recording, can be outsourced from the user equipment to the cloud through software-defined radio (Janssen et al., 2023). This results in the cloud-based snapshot positioning framework outlined in Figure 1. This approach reduces the TTFF, as complex tasks such as signal acquisition can be performed more quickly using the high computational power of cloud infrastructure. Additionally, snapshot receivers can implement a duty cycle mechanism to further reduce power consumption by switching the receiver off or putting it into sleep mode when positioning is not needed. This feature is particularly advantageous for applications such as asset tracking, where location updates are required only a few times a day (Liu et al., 2021).

The open-loop architecture of these energy-saving devices requires external assistance and modification of the navigation filter for computing the PVT solution (Linty, 2015; Liu et al., 2020). In terms of external assistance, as seen in assisted GNSS (Van Diggelen, 2009), the receiver obtains supplementary data from external

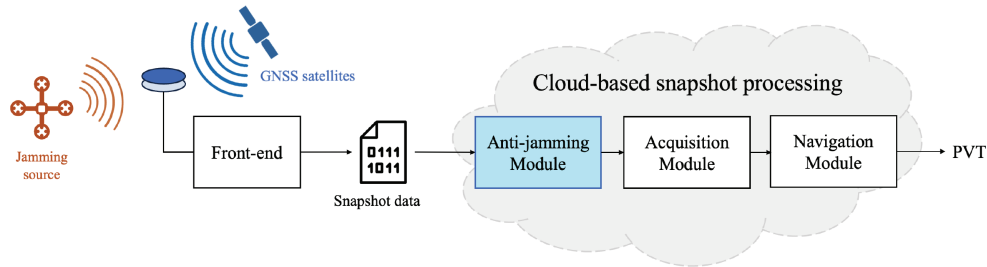


FIGURE 1 Overview of the system under study

An anti-jamming module is integrated into the GNSS baseline receiver with a snapshot architecture. In this setup, the acquisition and navigation modules are outsourced to the cloud, whereas the snapshot recording remains within the user equipment. Our study aims to assess the impact of low-bit quantization on RIM and variance estimation in the presence of three common types of jamming interference found in mass-market jammers.

sources to aid in positioning. For cloud-based snapshot positioning, the cloud serves as the external source. The assistance data may include a combination of approximate user position, ephemerides, almanac, time, and frequency information, which are crucial to help the filter compute the position with minimal signal data. Regarding time ambiguity, because snapshot receivers do not decode the full navigation message, they rely on a rough estimate of the GNSS signal transmission time. The coarse-time navigation filter allows the receiver to estimate the time of week (TOW) as an additional unknown in the navigation solution (Muthuraman et al., 2012). This approach offers advantages such as a shorter TTFF and energy savings, as the typical time spent tracking satellites (6 s) to decode the TOW can instead be used to obtain a position fix. However, this method requires a valid set of ephemerides, an approximate user position (within kilometers), and an approximate TOW (within seconds) (Van Diggelen & Abraham, 2007). The coarse-time navigation filter is also beneficial in weak signal conditions, where the receiver may be unable to reliably decode the TOW because of a high bit error rate.

These changes in the navigation filter typically reduce the position accuracy to tens of meters (Dumitraschkewitz, 2020). As a result, there has been an increasing focus in the literature on enhancing accuracy in snapshot positioning. Recent advancements have shown that snapshot receivers can achieve high-accuracy real-time kinematic (RTK) solutions. Using simulated data, Medina et al. (2020) demonstrated that snapshot RTK (SRTK) is feasible under certain signal quality conditions. Moreover, Liu et al. (2020) validated SRTK performance with real GNSS signals, although this analysis was limited to zero-baseline configurations (i.e., where the receiver and base station are connected to the same antenna). Finally, Liu et al. (2021) extended these findings to baselines up to 50 km. As these architectures are limited to the processing of narrow data chunks, which, if too large, may cause bottlenecks in the system, the minimum snapshot duration required to achieve a suitable performance must be investigated. This length may vary depending on the application requirements. For instance, Nardin et al. (2023) identified the minimum chunk length to be imposed as a requirement for lunar missions. The impact of different snapshot durations in the context of SRTK positioning has also been studied. Liu et al. (2021) tested durations of 40, 60, 80, and 100 ms to explore their effect on RTK fix rates and positioning accuracy. Better fix rates and positioning accuracy were achieved with longer snapshot durations, shorter baselines (i.e., shorter distance between the reference station and receiver), and higher signal bandwidths. Nevertheless, even with a short integration time of 40 ms, a solid performance was achieved, such as a 90% RTK fix rate at a 10-km baseline.

Despite the challenges associated with cloud-based snapshot positioning, such as the need for a communication module and potential delays from data transmission and cloud processing times (Liu et al., 2021), this approach remains a relevant strategy for low-power applications with flexible accuracy requirements. Snapshot architectures are of particular interest in areas such as IoT and location-based services, where manufacturers can integrate snapshot modules into mobile platforms and low-size, weight, and power (SWaP) devices (Van der Merwe et al., 2019). An interesting application of server-based processing is user position verification through encrypted GNSS signals (Vroom et al., 2023).

1.2 | GNSS Anti-Jamming

Classical jamming mitigation techniques are formulated as an estimation problem in which the jamming signal is detected and estimated to enable interference cancellation (IC), commonly through a parametric model. This principle consists in reconstructing the radio frequency interference (RFI) and subtracting it from the received signal. In this way, it is possible to work with a clean signal that is similar to the signal containing useful GNSS data. Adaptive notch filtering and pulse blanking are popular mitigation techniques that make use of the IC principle. Kalman filters can also be used to track and reconstruct jamming signals (Mitch et al., 2013). A major drawback of these IC techniques is that they require detection and estimation of the jammer waveform. When waveform-dependent methods are applied, jamming signal classification is required (Morales Ferre et al., 2019; Wu et al., 2025), given the need to make assumptions on the jamming signal waveform. To improve the traditional notch filter, an adaptive notch filter has been proposed (Borio et al., 2008, 2014). This method incorporates a transfer function with a filter zero coefficient that is adapted via a stochastic gradient approach to minimize the energy at the filter output. Although this method can track changes in the instantaneous frequency, it is only effective if the adaptation of the filter zero is fast enough to follow the frequency variations. This requires the adaptation parameters to be properly tuned before filtering, and the notch width must also be adjusted to mitigate jamming without significantly degrading the useful GNSS signal. Numerous multi-antenna techniques designed for GNSS applications (Fernández-Prades et al., 2016) have been shown to offer considerable benefits in addressing multipath and jamming interferences (Parkinson & Bartone, 2023; S. Li et al., 2023; Sun et al., 2022). Specifically, sparse array architectures may be leveraged to reduce hardware and computational costs in the task of estimating the direction of arrival of jamming sources and mitigating their impact (Amin et al., 2016). These methods can complement the previously mentioned techniques as well as the robust anti-jamming approach described hereafter.

A different jamming mitigation strategy emerged with the use of robust signal processing (Borio, 2017b), which makes it possible to discard disrupted samples by identifying them as outliers in certain domains (H. Li et al., 2023). Narrowband interferences assume a sparse representation in the frequency domain, whereas DME-like signals or pulsed interferences assume a sparse representation in the time domain. To understand why robust methods are so powerful, it is necessary to refer to the *breakdown point* concept. This concept corresponds to the proportion of disrupted samples that a method can handle before computing an incorrect result. Non-robust methods typically have a breakdown point of 0, whereas robust methods can achieve a breakdown point of 0.5. Robust anti-jamming techniques are built upon zero-memory nonlinear (ZMNL) functions and nonlinear correlators

based on robust methods. Some works in the literature use myriad nonlinearity (Borio, 2017a), complex signum nonlinearity (Borio & Closas, 2018), and Huber's nonlinearity (Borio et al., 2018; Li, 2023). While the focus of this paper is on interference mitigation, robust methods may also prove helpful in other tasks within the GNSS field, such as PVT estimation (Medina, Li, et al., 2019).

In this paper, we focus on the robust interference mitigation (RIM) framework proposed by Borio (2017a), which is a robust pre-correlation technique that removes outliers from the received signal without the need for interference detection, estimation, or waveform assumptions, unlike classical IC methods. RIM can be applied in several domains (e.g., frequency and time), showing successful results when compared with other state-of-the-art techniques such as the aforementioned notch filtering and pulse blanking (Borio et al., 2018). H. Li et al. (2019) proposed dual-domain RIM (DD-RIM), in which RIM is sequentially applied in the time and frequency domains. This approach exploits the sparsity of the jamming signals in both domains, with the cost of a duplicated loss of efficiency (LoE). Additionally, Gioia and Borio (2021) suggested that the RIM approach may be combined with other anti-jamming techniques in multilayer multi-constellation architectures.

The disparities between conventional and snapshot architectures must be considered in the implementation of anti-jamming techniques. For instance, Van der Merwe et al. (2019) emphasized that many spoofing detection techniques monitor PVT changes over time using the tracking module, which is unavailable in snapshot architectures. To address this limitation, those authors proposed a method to enhance spoofing detection for multi-antenna snapshot receivers through machine learning techniques. In a conventional GNSS architecture, RIM is applied to the coherent integration segment, which can be shorter than the snapshot itself. Therefore, the primary challenge of implementing RIM in a snapshot receiver lies not in the absence of the tracking loop, but in the requirement for low-bit quantization in low-power applications, where snapshot receivers offer particular advantages. To the best of the authors' knowledge, this is the first work to evaluate the performance of RIM methods in a snapshot architecture and with low-bit quantization, addressing a gap in the literature that typically assumes that a sufficient number of quantization bits is available to fully represent the jamming signal.

1.3 | Contributions

In this paper, we expand upon the work of Calatrava et al. (2023) by providing an extensive analysis of the snapshot RIM and presenting a benchmark study against a professional GNSS receiver. The main contributions of our research are as follows:

- The RIM framework is adapted for GNSS snapshot architectures, particularly to provide resilience against interferences with time-varying instantaneous frequency when applied in the frequency domain. The snapshot samples are divided into M batches, and Huber's ZMNL is applied to each batch independently. Additionally, a robust estimation of the noise variance is recursively calculated by considering all batches.
- Whereas other studies on RIM methods assume a signal resolution that is sufficiently high to fully represent the jamming signal, our study assesses the impact of low-bit quantization on the RIM framework and robust variance estimation. This feature is particularly relevant in snapshot architectures that prioritize cost-effective transmission.

- The performance of RIM is evaluated through two experiments. The first experiment uses simulated RFI data, and the second experiment uses realistic recordings provided by the European Space Agency (ESA). The latter includes a benchmark study against a commercial off-the-shelf (COTS) receiver.

The remainder of this paper is organized as follows. Section 2 formulates the GNSS and jamming signal models and offers insights into GNSS signal quantization. This section includes expressions for the amplitude of the quantized signal, the quantization thresholds, and the optimal automatic gain control (AGC) gain factor, which is obtained by minimizing the quantization loss. Section 3 outlines the processing steps of RIM, including the use of Huber’s nonlinearity for outlier detection, a robust estimation of the standard deviation of the snapshot noise, and the LoE expression. In Section 4, we describe how the RIM framework is adapted for the GNSS snapshot receiver architecture as well as the setup of the two conducted experiments. Section 5 presents our results and a discussion, starting with a comparative analysis between frequency- and time-domain snapshot RIM, followed by a study on the impact of signal quantization on the interference waveform, robust variance estimation, and RIM framework. Subsequently, we conduct a benchmark experiment comparing the performance of the baseline receiver, with the option to enable RIM, to that of a COTS GNSS receiver. Finally, Section 6 concludes the paper.

2 | SIGNAL MODEL

Given the quasi-orthogonality of the spreading codes, GNSS receivers are able to process L useful satellite signals independently. A signal received in a one-path additive Gaussian channel in the presence of a jamming attack can be modeled as follows (Kaplan & Hegarty, 2005; Morton et al., 2021):

$$y(t) = \sqrt{2C}d(t - \tau_0)c(t - \tau_0)\cos(2\pi(f_{\text{RF}} + f_0)t + \phi_0) + \eta(t) + i(t) \quad (1)$$

where C is the useful signal power, $\eta(t)$ is zero-mean additive white Gaussian noise (AWGN), and $i(t)$ corresponds to the jamming interference. The navigation message and pseudorandom code are represented by $d(\cdot)$ and $c(\cdot)$, respectively. Owing to the relative dynamic between the satellite and receiver, there is a Doppler shift f_0 with respect to the signal radio frequency f_{RF} . In addition, the communications channel introduces a delay and phase shift denoted as τ_0 and ϕ_0 , respectively. After the receiver front-end amplifies, filters, and down-converts $y(t)$, the signal is digitized by the analog-to-digital converter (ADC) in two steps: sampling and amplitude quantization. Assuming that Equation (1) is sampled without introducing significant distortions (Borio et al., 2012), the down-converted in-phase/quadrature baseband samples are obtained after front-end filtering as follows:

$$y[n] = \sqrt{C}d(nT_s - \tau_0)c(nT_s - \tau_0)e^{j2\pi f_0 nT_s + j\phi_0} + \eta[n] + i[n] \quad (2)$$

where n is the time index and $[\cdot]$ denotes a discrete time sequence sampled at the receiver sampling frequency $f_s = \frac{1}{T_s}$. The signal $y(t)$ has been down-converted to the baseband frequency, and consequently, the noise component $\eta[n]$ and the interference $i[n]$ are baseband signals. We model $\eta[n]$ as complex AWGN with independent and identically distributed real and imaginary parts, each with

variance σ^2 . This variance can be expressed as $\sigma^2 = \frac{1}{2}N_0B_{R_x}$, where B_{R_x} is the receiver front-end two-sided bandwidth and N_0 is the power spectral density (PSD) of the input noise $\eta(t)$.

2.1 | Jamming Signals

The received jamming interference can assume different waveforms. According to Morales Ferre et al. (2019), jamming signals can be classified into five categories: Class I jammers are CW-modulated signals with a bandwidth of up to 100 kHz; Class II and III jammers include single-chirp and multi-chirp signals; Class IV are chirp signals with frequency bursts that aim to expand the band of disrupted frequencies; and Class V jammers include pulsed and DME-like signals. Class I jammers, also known as amplitude-modulated jammers, generate single- or multi-tone sinusoidal signals. These jammers represent the simplest and most widely studied waveform type. Chirp jammers from Classes II–IV are characterized by an instantaneous frequency $f_j(t)$ that changes over time, often showing a sawtooth-like waveform with rapidly changing frequency or phase characteristics. In this paper, we consider a single-chirp jammer with frequencies periodically sweeping from f_{\min} to f_{\max} every T_{sweep} seconds. Sweep periods are commonly around 10 μs , and sweep ranges ($\Delta f = f_{\max} - f_{\min}$) can vary within the range of 10–40 MHz (Mitch et al., 2011). Regarding Class V, signals transmitted by DME and TACAN systems are examples of legitimate pulsed jamming waveforms that can potentially disrupt GNSS performance. Aircraft interrogators in DME systems transmit paired high-power pulses, with each pulse potentially modeled by a Gaussian function. TACAN combines pulse pairs with amplitude modulation, where the amplitude of the pulses is modulated by a rotating antenna to provide both distance and azimuth information. Gao (2007) described the DME/TACAN pulse structure, including DME pulse pairs as well as TACAN pulse amplitude modulation and reference pulse group patterns.

Several metrics have been adopted in the literature to characterize the relationship between jamming and noise signals. The jammer-to-noise ratio (J/N) is defined as the ratio between the received jamming power J and the noise power N , whereas the jammer-to-noise PSD ratio (J/N_0) is defined as the ratio between J and N_0 . These two ratios are related as follows:

$$\frac{J}{N} = \frac{J}{N_0} \frac{1}{B_{R_x}} = \frac{JT}{N_0} \frac{1}{N_{\text{snap}}} = \frac{A_j^2}{2\sigma^2} \quad (3)$$

where T is the snapshot duration in seconds, A_j is the amplitude of the jamming signal, and N_{snap} corresponds to the number of snapshot samples or chunk length. One of the main signal quality indicators in GNSS is the carrier-to-noise PSD ratio (C/N_0), defined as the ratio between the carrier power C and C/N_0 depends solely on the received signal and is continuously estimated by the receiver and determined post-correlation, usually in logarithmic units, i.e., dB-Hz. GNSS receivers first estimate the signal-to-noise ratio (SNR) and then use a model that relates the SNR to C/N_0 (Borio et al., 2012). The C/N_0 estimated by the receiver may be referred to as an effective C/N_0 , given that SNR estimation does not usually account for the degradation resulting from, for instance, jamming interferences (Betz, 2000, 2001). Considering this, it is possible to quantify the effect of a jammer as the degradation in the effective C/N_0 .

2.2 | Signal Quantization

RFI detection and mitigation techniques usually assume that input data are quantized with a significant number of bits B that is sufficient to fully represent the jamming signal, i.e., at least 8 (Díez-García & Camps, 2019). Nevertheless, this is not the case with most GNSS receivers, as 2-bit ADCs are very common in the market. Moreover, in the absence of jamming interferences and under suitable conditions, GNSS receivers provide reasonable results for $B = 1$. To maximize the effectiveness of state-of-the-art jamming detection and mitigation techniques, it is beneficial to have enough quantization levels available so that the full jamming signal is represented. A sufficient number of quantization bits is required to ensure that the receiver is sensitive to the power differences between non-interfered and interfered samples.

The number of quantization bits and the quantization dynamic range can be dynamically altered according to the characteristics of the received signal. The quantization dynamic range is defined as the difference between the largest and smallest values that a quantized signal can assume. A fixed dynamic range guarantees a proper representation of the expected noise signal. However, samples affected by jamming suffer from clipping because, as the name implies, the range is fixed and is not susceptible to changes in the received power. This clipping effect occurs when the number of quantization bits is not sufficient to represent the signal, in cases such as when the signal is overpowered by a jammer. Díez-García and Camps (2019) proposed an adaptive dynamic range system based on the AGC, where the quantization thresholds and the dynamic range are determined by the variance of the input signal.

The quantized signal can assume the amplitude values in the set $B = \{-(2^B-1), \dots, -3, -1, 1, 3, \dots, 2^B-1\}$, which corresponds to odd numbers $\{2i+1\}_{i=-2^{B-1}, -2^{B-1}+1, \dots, 2^{B-1}-1}$ (Borio, 2008, §6.2.3). The quantization thresholds can be calculated as $Q_0 = \{-(2^B-1) + 2, \dots, -3, -1, 0, 1, 3, \dots, (2^B-1) - 2\}$. Additionally, the input signal must be multiplied by an optimal AGC gain factor. This factor depends on the standard deviation of the snapshot noise before the quantization block, denoted as σ . This term may be found by minimizing the quantization loss expression proposed by Borio (2008, §6.2.4), which is expressed as follows:

$$L(A_g) = \frac{2}{\pi} \frac{\left(1 + 2 \sum_{i=1}^{2^{B-1}-1} e^{-\alpha_i^2}\right)^2}{1 + 8 \sum_{i=1}^{2^{B-1}-1} i \cdot \operatorname{erfc}(\alpha_i)} \quad (4)$$

where $\alpha_i = -i / (\sqrt{2} A_g \sigma)$ and consequently $A_{\text{opt}} = \min_{A_g} L(A_g)$. Figure 2 shows quantization loss values as a function of the normalized AGC gain, denoted as $A_g \sigma$, for the values of B considered in the experimental analysis. The figure suggests that the requirement for an optimal AGC gain can be relaxed for higher values of B because these cases exhibit a larger range of normalized AGC gain values at which the quantization loss remains almost constant. The optimal AGC gain factor values and their corresponding quantization losses are shown in Table 1. The loss introduced by signal quantization for $B = 8$ is almost negligible for an increasing noise standard deviation up to a normalized AGC gain of 70. In contrast, the minimum loss for $B = \{2, 4\}$ is higher, with a very limited range of normalized AGC gain. The loss in 1-bit quantization systems cannot be improved by adjusting the AGC gain, as it remains constant at 1.96 dB.

GNSS positioning can still be reliable with 1-bit quantization owing to the robustness of pseudorandom noise codes. This feature makes low-bit




FIGURE 2 Quantization loss as a function of the normalized AGC gain for the values of B considered in the experimental analysis

The quantization loss expression in Equation (4) was proposed by Borio (2008, §6.2.4).

TABLE 1

Optimal AGC Gain Factor A_{opt} for the Number of Quantization Bits B Considered in the Experimental Analysis

These values have been obtained by minimizing the quantization loss expression in Equation (4) proposed by Borio (2008, §6.2.4). Minimum quantization loss values are also provided.

B	1	2	4	8
A_{opt}	-	$1.0043/\sigma$	$2.9833/\sigma$	$32.5072/\sigma$
Minimum Loss (dB)	-1.96	-0.55	-0.05	-4×10^{-4}

quantization receivers especially attractive for low-power applications requiring simple hardware. However, higher quantization levels are often necessary in scenarios involving interference, multipath, or high-precision needs. As a result, most studies on RIM assume a sufficient bit depth to fully represent jamming signals, typically $B > 8$. Recognizing that such high bit depths may not suit low-power applications in which snapshot architectures are more practical, our work evaluates the performance of snapshot RIM at lower ADC quantization levels. This approach allows us to investigate the limits of the well-established RIM strategy.

3 | ROBUST INTERFERENCE MITIGATION

RIM can be implemented in three processing steps, namely, a first linear transformation \mathbf{T}_1 , a ZMNL, and a second linear transformation such that $\mathbf{T}_1 \circ \mathbf{T}_2 = \mathbf{I}$, where \mathbf{I} is the identity operator. The first linear transformation projects the jamming component into the transform domain (TD). This domain must be selected such that the jamming signal assumes a sparse representation, affecting only a limited number of samples, which can then be identified as outliers. As RIM is a pre-correlation technique, \mathbf{T}_1 is applied directly to the samples $y[n]$ before correlation, yielding $\mathbf{Y}[k] = \mathbf{T}_1(y[n])$, where k is the discrete frequency index. \mathbf{T}_2 inverts the effects of \mathbf{T}_1 , bringing the TD samples back to the time domain as $\tilde{y}[n] = \mathbf{T}_2(\mathbf{Y}_v[k])$. When RIM is applied in the time domain, \mathbf{T}_1 is a low pass-filter, and \mathbf{T}_2 is the identity operator. The low-pass filter is applied so that the jamming component appears as a sequence of time pulses (Borio, 2017a). When applied in the frequency domain, \mathbf{T}_1 and \mathbf{T}_2 are the fast Fourier transform (FFT) and inverse FFT (IFFT), respectively. By applying these transformations, the impact of the jammer $i[n]$ on $\tilde{y}[n]$ is reduced with respect to the initial time samples $y[n]$.

After \mathbf{T}_2 has been applied, a standard receiver architecture is adopted with the filtered samples $\tilde{y}[n]$. The resulting bidimensional robust cross ambiguity function (CAF) is computed as follows (Borio, 2017a):

$$C_\psi(\tau, f_d) = \sum_{n=0}^{N-1} \tilde{y}[n] c(nT_s - \tau) e^{-j2\pi f_d n T_s} \quad (5)$$

where $T_c = NT_s$ corresponds to the coherent integration time and $N = N_{\text{snap}}$. The CAF depends on the locally tested code delay τ and Doppler frequency f_d , which are obtained by maximizing the function in the acquisition and tracking stages.

3.1 | Huber's ZMNL

Following the work by Borio et al. (2018), we use Huber's nonlinearity to reduce the impact of interfered samples in the TD. This ZMNL is the derivative of Huber's loss function and allows us to identify outliers by comparing the amplitude of the TD samples with a decision threshold. Huber's ZMNL is defined as follows (Wang & Poor, 1999):

$$\psi_H(\Upsilon[k]) = \begin{cases} \Upsilon[k] & \text{for } |\Upsilon[k]| \leq T_h \\ T_h \text{ csign}(\Upsilon[k]) & \text{for } |\Upsilon[k]| > T_h \end{cases} \quad (6)$$

where $|\cdot|$ is the absolute value operator and T_h is the nonlinearity decision threshold, with an optimal value of $T_{\text{opt}} = 1.345\sigma$ for the real-valued signal case (Fox, 2002). The complex signum operator preserves the phase of the samples $\Upsilon[k]$ and can be defined as follows:

$$\text{csign}(\Upsilon[k]) = \begin{cases} \frac{\Upsilon[k]}{|\Upsilon[k]|} & \text{for } \Upsilon[k] \neq 0 \\ 0 & \text{for } \Upsilon[k] = 0 \end{cases} \quad (7)$$

The standard deviation of the snapshot noise must be known to calculate the Huber nonlinearity decision threshold T_h . In the presence of interference, we need a robust measure of statistical dispersion such as the median absolute deviation (MAD) to estimate σ . The MAD can be expressed as follows:

$$\text{MAD}(\mathbf{x}) = \beta \text{Med}(|\mathbf{x} - \text{Med}(\mathbf{x})|) \quad (8)$$

with $\text{Med}(\mathbf{x})$ denoting the median of \mathbf{x} and β being a constant scale factor. This method is widely used for obtaining an asymptotically consistent estimator of the standard deviation σ due to its simplicity and its breakdown point of 0.5 (Rousseeuw & Croux, 1993). The value of the scale factor β is distribution-dependent, being set to $\beta = 1.4815$ under a Gaussian assumption. The variance must be estimated after the first linear transformation has been applied, i.e., with the samples in $\Upsilon[k]$.

Borio et al. (2018) interpreted Huber's nonlinearity as a switch between the Gaussian and Laplace regimes, based on the fact that the nonlinearity in Equation (6) when a TD sample is recognized as an outlier, i.e., when $|\Upsilon[k]| > T_h$, is equivalent to the nonlinearity that results from the Laplace assumption (Borio & Closas, 2018). The Gaussian regime is applied when interference is weak, effectively modeling noise under nominal conditions but remaining sensitive to outliers. In contrast,

the Laplace regime is more robust to outliers owing to its heavy-tailed distribution, although it sacrifices some precision under nominal conditions. The impact of this trade-off in performance is quantified through an LoE analysis, as detailed in the next subsection. The combination of these two regimes, as depicted in Figure 3, allows the framework to dynamically adapt to varying interference environments without requiring specific detection or estimation of the interference waveform.

3.2 | Loss of Efficiency

The LoE caused by Huber's nonlinearity has been evaluated by Borio et al. (2018). The LoE is defined as the performance degradation caused by the ZMNL in the absence of interference and is given by the following ratio:

$$L_0(T_h) = \frac{\text{SNR}_{\text{out}}^\psi}{\text{SNR}_{\text{out}}} \quad (9)$$

where SNR_{out} corresponds to the post-correlation SNR that measures the quality of $C(\tau, f_d)$. This ratio may be defined as follows (Betz, 2000, 2001; Borio et al., 2018):

$$\text{SNR}_{\text{out}} = \max_{\tau, f_d} \frac{|\text{E}\{C(\tau, f_d)\}|^2}{\frac{1}{2} \text{Var}\{C(\tau, f_d)\}} \quad (10)$$

$\text{SNR}_{\text{out}}^\psi$ corresponds to the post-correlation SNR that measures the quality of $C_\psi(\tau, f_d)$. Borio et al. (2018) demonstrated that the LoEs obtained under time- and frequency-domain processing are comparable, with a maximum value of 1.05 dB. As mentioned in Section 1, the DD-RIM framework used by Li et al. (2019) exploits the sparsity of the jamming signals in the time and frequency domains simultaneously, with the cost of a duplicated LoE. This approach results in a maximum LoE of 2.1 dB.

4 | METHODS

RIM has been implemented in a cloud-based GNSS snapshot receiver, which we consider our baseline receiver. Once the snapshot samples are recorded, they are sent to the cloud, where the impact of interference can be mitigated by following the steps shown in Figure 3. In Section 4.1, we elaborate on how the outlier detection block is adapted for snapshot processing. In this study, two experiments have been conducted: the first experiment aims to assess the impact of signal quantization on receiver performance in the presence of interference and when RIM is available, whereas the second experiment is designed for comparative purposes with a benchmark receiver. The setup of the two experiments is described in Sections 5.1 and 5.2.

4.1 | Snapshot RIM

In this section, we describe how frequency-domain RIM has been adapted for snapshot processing in the presence of interferences with time-varying instantaneous frequency. Chirp interferences commonly sweep a significant portion of

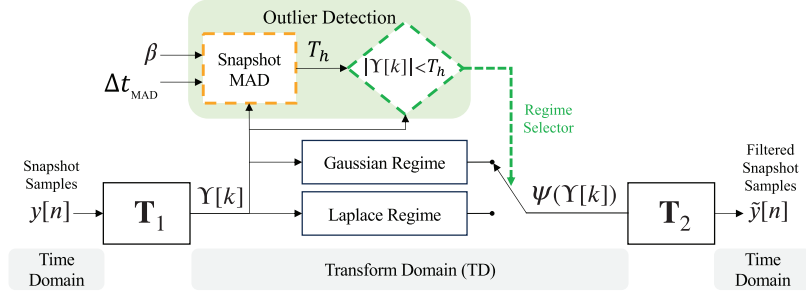


FIGURE 3 This overview of the RIM framework includes the interpretation of Huber’s nonlinearity, denoted as $\psi_H(\cdot)$, as a switch between Gaussian and Laplace regimes, as proposed by Borio et al. (2018). The RIM algorithm is implemented in three processing steps: a first linear transformation to the TD denoted as \mathbf{T}_1 , a ZMNL such as Huber’s nonlinearity, and a second linear transformation subject to $\mathbf{T}_1 \circ \mathbf{T}_2 = \mathbf{I}$. To compute the Huber ZMNL threshold T_h , we estimate σ^2 with the MAD measure of statistical dispersion. When applied in the frequency domain, the MAD is recursively calculated by dividing the snapshot into batches, as shown in Figure 4.

the receiver bandwidth over time periods that are significantly shorter than the snapshot duration. Under these conditions, the snapshot samples no longer admit a sparse representation of the jamming component in the frequency domain. To address this issue, we recursively update the snapshot MAD every Δt_{MAD} seconds. Figure 4 provides an overview of how the snapshot MAD is robustly estimated by using recursion and how this connects to the regime selector from the RIM framework in Figure 3. The number of snapshot samples in the time domain is denoted as $N_{\text{snap}} = M \times N_M$, where $N_M = f_s \Delta t_{\text{MAD}}$ represents the number of samples per batch and M is the number of batches into which the snapshot is divided. Here, $m = [1, \dots, M]$ represents the batch index, and ψ_m denotes the application of the ZMNL in Equation (6) applied to the samples of batch m . The snapshot MAD is calculated recursively as follows:

$$\bar{\sigma}_m^2 = \frac{\sigma_m^2}{m} + \frac{m-1}{m} \bar{\sigma}_{m-1}^2 \quad (11)$$

where $\bar{\cdot}$ is defined as the recursive operator. Considering this, $\bar{\sigma}_m^2$ corresponds to the recursive estimate at update m and is dependent on the recursive estimation at update $m-1$, denoted as $\bar{\sigma}_{m-1}^2$. The non-recursive estimate σ_m^2 is computed as the MAD of the N_M samples of batch m in the frequency domain, i.e., after the first linear transformation \mathbf{T}_1 has been applied as an FFT. Once we have obtained the robust estimate $\bar{\sigma}_M^2$, we calculate the corresponding optimal value of the decision threshold for this snapshot, i.e., $T_h = 1.345 \bar{\sigma}_M^2$. The ZMNL is also applied to every N_M samples, with the difference that no recursion is used for this step and, consequently, each batch is treated independently. Moreover, batches from the same snapshot share the same decision threshold value after its estimation via the recursive MAD. In the conducted experiments, frequency-domain RIM is applied every $\Delta t_{\text{MAD}} = 1$ ms, and the robust MAD measure is computed recursively, considering snapshot batches of the same duration. In contrast, time-domain RIM is applied to the whole snapshot, and the MAD measure is computed by considering all snapshot samples N_{snap} .

Several aspects are worth discussing regarding the implementation of RIM. In conventional receivers, an FFT can be used in fast parallel algorithms to accelerate CAF computation during acquisition. These same FFTs can serve as the outputs of the first linear transformation \mathbf{T}_1 when RIM is applied in the frequency domain.

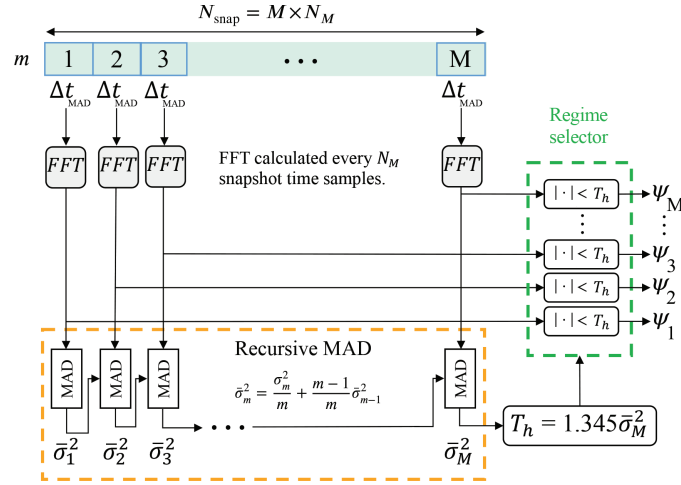


FIGURE 4 Overview of the snapshot RIM framework, where the outlier detection module is adapted to preserve the sparse nature of chirp jammers in the frequency domain. The snapshot samples are divided into M batches of Δt_{MAD} seconds, and Huber's nonlinearity, denoted as ψ_m , is applied to the N_M samples of each batch independently. The MAD measure of statistical dispersion is calculated recursively.

Based on findings in the literature and our results presented in Section 5.1.1, frequency-domain RIM is typically sufficient to mitigate most interferences, owing to the periodic nature of common jamming signals (Borio & Closas, 2019). Consequently, no additional computational load is required for FFT calculation with RIM when FFT-based algorithms for rapid acquisition are used. Furthermore, because RIM is applied pre-correlation, this approach only requires that the ZMNL be applied once, rather than repeating the processing for each correlator or satellite (Borio & Closas, 2019). In a conventional receiver, RIM is applied to the samples within the coherent integration segment, which is usually shorter than a snapshot. This shorter duration does not allow partitioning into multiple batches, as is done in snapshot RIM. In contrast, the practical implementation of RIM in snapshot receivers benefits from the availability of multiple batches within a snapshot, allowing for recursive estimation of signal variance across these batches. As previously mentioned, batches are selected to have a duration of 1 ms, matching the duration of a coherent integration segment. This choice ensures consistency with the implementation of RIM in a conventional architecture. Ultimately, if snapshots are shortened, fewer batches are available for refining variance estimates; however, RIM still mitigates interference with 1-ms chunks, as demonstrated in the literature, when applied to conventional receivers (Borio et al., 2018). Overall, the main challenge for the implementation of RIM in snapshot receivers does not arise from the lack of tracking module or snapshot duration constraints, as coherent integration segments in conventional receivers are typically shorter than snapshots. Rather, the primary challenge arises from the need for low-bit quantization to meet communication and power constraints in low-power applications.

4.2 | RFI Simulation

In the first experiment, we use Python to generate 120-ms snapshots with GNSS and jamming data under various configurations. Global Positioning System (GPS) L1, Galileo L1, and BeiDou B1 signals are generated with the same sampling frequency as the two-sided bandwidth of the baseline receiver, i.e., $f_s = B_{R_x} = 4.5 \text{ MHz}$.

TABLE 2

Classification of Baseband Signal Models for the Interferences Implemented in the First Experiment

 This classification follows the taxonomy proposed by Morales Ferre et al. (2019) (Class I: CW, Class II: single-chirp, Class V: pulsed/DME-like). The value of A_J is not specified in this table due to its variability with J/N . Values of J/N under study vary from 0 to 95 dB.

ID	Class	Type	Baseband model	Design parameters	Configuration
1	I	CW	$i[n] = A_J \exp(j2\pi f_J nT_s)$	A_J, f_J	$f_J = -1$ kHz
2	II	Chirp	$i[n] = A_J \exp(j2\pi T_s \sum_{j=0}^n f_J[j])$	$A_J, T_{\text{sweep}}, \Delta f$	$T_{\text{sweep}} = 10$ μ s; $\Delta f = 9$ MHz
3	II, V	Pulsed chirp	Chirp with $\Delta f > B_{R_x}$	$A_J, T_{\text{sweep}}, \Delta f$	$T_{\text{sweep}} = 10$ μ s; $\Delta f = 20$ MHz

The baseband signal models used for interference simulation are presented in Table 2, which considers the taxonomy proposed by Morales Ferre et al. (2019). Jammer 1 is a single-tone CW with f_J set to -1 kHz with respect to the baseband frequency, Jammer 2 is a chirp interference sweeping the whole receiver bandwidth ($f_{\min} = -\frac{f_s}{2}$ and $f_{\max} = \frac{f_s}{2}$), and Jammer 3 is a wideband jammer with a sweep range of 20 MHz (from -10 to 10 MHz). Both Jammers 2 and 3 have a sweep period of duration $T_{\text{sweep}} = 10$ μ s. The behavior of CW, chirp, and pulsed chirp jammers in the frequency domain over time can be understood from their spectrogram plots (Morales Ferre et al., 2019, Figure 1).

Wideband interferences with a sweep range greater than the receiver bandwidth can be interpreted as pulsed interferences, as the receiver is only affected by the jamming attack when the instantaneous frequency is within the receiver bandwidth. This is the case for Jammer 3, whose bandwidth (20 MHz) exceeds the receiver bandwidth (4.5 MHz), causing the interference to appear as pulsed interference in the time domain. Because of this pulsed effect, Jammer 3 does not contaminate all samples within the snapshot, as there are periods during which its instantaneous frequency lies outside the receiver bandwidth. In contrast, Jammer 2 consistently targets frequencies within the receiver bandwidth over the entire snapshot duration, $t \in [0, 120]$ ms. The three jammers show a sparse nature in the frequency domain, whereas only Jammer 3 shows a sparse nature in the time domain, given its pulsed effect. Jammer 3 is categorized as both a Class II single-chirp jammer and a Class V DME-like jammer; thus, we are evaluating the snapshot RIM performance against signals similar to those from DME/TACAN systems.

As a preliminary step to proceed with the experimental section, we determined the appropriate ranges of J/N ratio values to consider for evaluation. Characterizations of mass-market jammers in previous literature (Mitch et al., 2011) have shown that the maximum jamming transmission peak power typically varies from -10 dBm to 30 dBm. This power corresponds to a J/N ratio of approximately 95–130 dB, respectively, in the case of our receiver. To convert from J/N to J/N_0 (see Equation (3)), we consider the two-sided receiver bandwidth B_{R_x} to be equal to the sampling frequency $f_s = 4.5$ MHz, as we are performing complex sampling. To consider the worst-case scenario, in which the jammer suffers from the least possible attenuation, the free-space path loss (FSPL) factor is computed as $\text{FSPL} = \frac{4\pi d^2}{\lambda}$, where d corresponds to the jammer–receiver distance and λ corresponds to the signal wavelength, which depends on the frequency of transmission. As L-band wavelengths range between 15 and 30 cm, the worst-case scenario (i.e., least path-loss attenuation for the jammer) is given by $\lambda = 30$ cm, resulting in 60 dB for $d = 100$ m. Considering that the maximum jammer transmitted peak power can reach 130 dB, the maximum J/N ratio that a receiver will experience is 70 dB for $d = 100$ m. Consequently, the J/N ratio values that we considered for assessing the performance of the RIM algorithm in the experimental section vary from 0 to 95 dB, with an additional 25 dB considered as a safety margin.

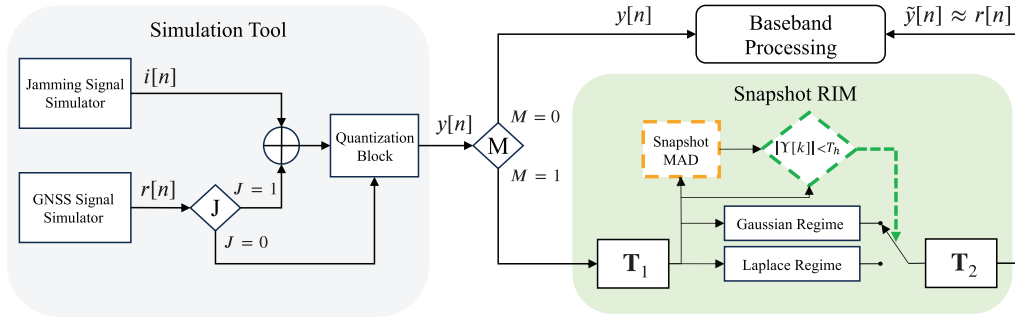


FIGURE 5 Diagram of the system setup in the first experiment with simulated RFI data. The parameter J controls the simulation and addition of a jamming signal to the clean GNSS signal, whereas the parameter M determines whether the RIM algorithm is enabled.

A diagram of the system model for this experiment is presented in Figure 5. In this diagram, $r[n]$ corresponds to the clean snapshot samples generated by the GNSS signal simulator, and $i[n]$ represents the jamming interference generated by the jamming signal simulator. The anti-jamming module is placed between the simulation tool, which is equivalent to the system front-end, and the baseband processing block. The parameter J controls the option to simulate and add a jamming signal to the clean GNSS snapshot, and the parameter M controls the option to enable RIM. If $J = 1$, the output of the simulation tool is polluted by the jammer as $y[n] = r[n] + i[n]$, whereas $y[n] = r[n]$ if $J = 0$. When $M = 1$, RIM is applied to the snapshot samples $y[n]$. In this case, the input of the baseband processing block is $\tilde{y}[n]$, where $\tilde{\cdot}$ represents the impact of RIM. If the anti-jamming module is successful, $\tilde{y}[n] \approx r[n]$. If $M = 0$, the input signal of the baseband processing block corresponds to the simulated snapshot samples $y[n]$. The LoE analysis is conducted under the conditions of $J = 0$ and $M = 1$.

The quantization block in Figure 5 reduces the resolution of the snapshot samples $y[n]$ from $B = 16$ to the desired value of B . The performance of RIM is assessed for $B = \{2, 4, 8\}$ and an adaptive quantization dynamic range that is proportional to the snapshot variance, which increases with the presence of jamming. For large dynamic ranges, a higher resolution is needed to perceive the contribution of useful GNSS signal, which can be achieved by increasing the number of quantization bits. The optimal AGC gain factor is calculated according to the expressions provided in Table 1. It is relevant to highlight that the variance used to calculate the AGC gain factor is not calculated robustly and, consequently, contains the effect of the jamming component. Each case of study within the first experiment is defined by a combination of values for J and M and a specific jammer type. For each case of study, ten snapshots with a duration of 120 ms (i.e., 1200 ms of recording) are processed. For the LoE assessment, two 120-ms snapshots are processed.

4.3 | ESA Test-Bed

In the second experiment, the performance of a COTS GNSS receiver is compared with that of the baseline receiver in the processing of realistic GNSS jammed signal recordings provided by ESA. A diagram of the setup used to generate and record realistic GPS L1 and Galileo L1 signals under three different scenarios is shown in Figure 6. A different jamming interference is simulated for each scenario, including a single-tone CW centered at -1 kHz with respect to the baseband frequency, a chirp with $T_{\text{sweep}} = 10$ ms and $\Delta f = 9$ MHz, and another chirp with $T_{\text{sweep}} = 10$ ms and $\Delta f = 20$ MHz. These interferences resemble those simulated in the first study

TABLE 3

RFI Features for the Three RFCS Input Scenarios in the Second Experiment
 These interferences resemble those used in the first experiment (see Table 2). Values of I/S under study vary from 0 to 100 dB.

ID	Class	Type	Configuration
1	I	CW	$f_j = -1$ kHz
2	II	Chirp	$T_{\text{sweep}} = 10$ ms; $\Delta f = 9$ MHz
3	II, V	Pulsed Chirp	$T_{\text{sweep}} = 10$ ms; $\Delta f = 20$ MHz

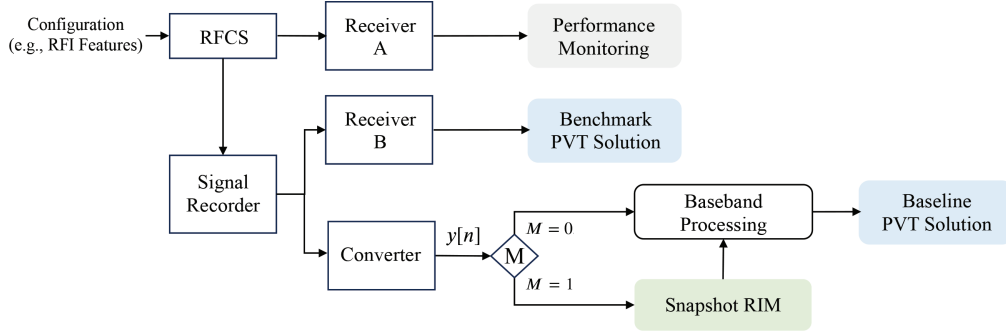


FIGURE 6 Setup used by ESA to generate and record realistic GPS L1 and Galileo L1 signals in the presence of three RFI signals with I/S values from 0 to 100 dB
 Benchmark PVT solutions are derived based on RINEX observation files recorded by receiver B and are compared with the baseline positioning results. Receiver A is included only to confirm results during testing.

for consistency, but with experimental data. As in Experiment 1, Jammer 3 resembles the pulsed signals emitted by DME/TACAN systems. This feature is a result of the pulsed effect described in Section 4.2, which arises when the interference bandwidth is significantly greater than the receiver bandwidth.

To generate RFI data, three configurations are defined as the input of the radio frequency constellation simulator (RFCS), each consisting of several files describing the characteristics of the signals to be generated, including the RFI features (e.g., sweep period, transmit power). Two GNSS benchmark receivers are used. The first one (receiver A) is connected to the RFCS, and the second one (receiver B) is connected to the signal recorder. The output of receiver A is not used for further evaluation; its purpose is to provide a confirmation of the results while the test-bed is being used. The signal recorder provides a binary file that can be processed by the snapshot baseline receiver, providing a baseline PVT solution. Benchmark PVT solutions are derived based on the receiver-independent exchange (RINEX) observation file recorded by receiver B. In Experiment 1, the J/N ratio was considered in the definition of the interference features; in contrast, in this experiment, the interference-to-signal ratio (I/S) is used instead. These two ratios are proportional to the interference power, with I/S ratio values typically being lower because they correspond to the ratio between the interference power and the useful signal power, not just the noise. For this experiment, I/S values range from 0 to 100 dB, with power transitions of 1 dB every 2 s of recording (i.e., approximately 16 120-ms snapshots).

5 | RESULTS AND DISCUSSION

5.1 | Experiment 1

In this section, we assess the performance of the snapshot RIM framework with the GNSS and jamming data described in Section 4.2. We first compare results

obtained with the frequency- and time-domain RIM algorithms in Section 5.1.1, investigating which is the most suitable option to ensure that the baseline receiver is resilient against typical interferences while avoiding an unnecessary increase in LoE. In Section 5.1.2, we discuss the effect of low-bit quantization on the baseline receiver in the presence of interference and in terms of the effective J/N , waveform distortion, and quantization gain. The impact of quantization on robust variance estimation is also studied. We provide a general performance assessment and LoE results in Section 5.1.3.

5.1.1 | *Frequency- vs. Time-Domain Snapshot RIM*

It is of interest to compare the results obtained with the snapshot RIM approach in the time and frequency domains. The three implemented jamming signals assume a sparse representation in the frequency domain, with Jammer 3 also assuming a sparse representation in the time domain owing to its wide sweep range, which causes a narrowband pulsed effect. Figures 7(c) and 7(d) show the performance of the frequency- and time-domain snapshot RIM for $B = 8$ in the presence of the three jammers implemented in the first experiment (see Table 2). The number of satellite vehicles (SVs) detected by the receiver is plotted as a function of the simulated interference J/N . When the receiver is jeopardized by a jammer and RIM is not enabled, this metric drops dramatically. Nevertheless, when RIM is enabled in a domain where the interfered samples can be identified as outliers, we observe an increase in the range of J/N values at which the receiver can detect a considerable number of SVs.

The results show that Jammer 3 can be successfully mitigated by both the time- and frequency-domain snapshot RIM, with the performance of the frequency-domain RIM being slightly better. Frequency-domain RIM can reduce the impact of Jammers 1 and 2, whereas time-domain RIM is not successful, as the interfered samples from these two jammers cannot be identified as outliers in the time domain. For instance, in the presence of Jammer 1, the gain offered by RIM is 35 and 5 dB in the frequency and time domains, respectively. Morales Ferre et al. (2020) suggested that most interferences found in typical environments assume a sparse representation at least in the frequency domain. This trend suggests that applying the frequency-domain RIM is sufficient to mitigate jamming interferences in typical environments. A solution would be to use DD-RIM, in which the RIM algorithm is sequentially applied in the time and frequency domains. However, this approach would lead to a duplicated LoE. Considering this, the RIM algorithm is only applied in the frequency domain throughout the remainder of our study.

5.1.2 | *Effect of Signal Quantization*

Effective J/N

Snapshot samples affected by jamming suffer from a clipping effect when the number of quantization bits is low. This phenomenon occurs because the number of available quantization values is not sufficient to allow for full signal representation. Consequently, the J/N observed at the receiver, which we refer to as the effective J/N , differs from its actual value, similar to how the effective C/N_0 differs from its actual value, as explained in Section 2.1. In Figure 7(a), the value of effective J/N observed at the receiver is studied as a function of the actual value of J/N used to simulate the jamming signal, which is a known input parameter of the interference

simulator. This figure shows that the quantization dynamic range increases proportionally with B and that a saturation point occurs when the whole dynamic range is used. Once this saturation point is reached, an increase in jamming power cannot be further perceived by the receiver. The saturation point occurs at J/N values of 0, 12, 17, and 35 dB with effective J/N values of approximately 0, 5, 12, and 33 dB for 1, 2, 4, and 8 quantization bits. Thus, for instance, in the case of $B = 4$, the value of J/N measured at the receiver cannot be greater than 12 dB, and this value is reached when the J/N of the simulated interference is 17 dB.

For a low B , the magnitude of J/N that the receiver can perceive is also low. For $B = 2$, as quantization values are limited and signal power is more spread owing to waveform distortion, the effective J/N at the receiver is very low. Although one might think that this low value would diminish the effect of powerful interference without the need for mitigation, this is not the case. The results in Section 5.1.3 suggest that when RIM is not enabled, the receiver observes a protection gain in terms of detected SVs and PVT availability only by lowering the number of quantization bits B from 8 to 2. However, the C/N_0 plots show a strong performance decrease for $B = 2$ when RIM is both enabled and disabled. The pseudo-resilience offered by low resolution does not translate in terms of actual signal quality, i.e., the number of detected SVs or PVT availability can be high while the actual C/N_0 at the receiver is very low. This finding suggests that, in the presence of interference, it is indeed necessary to use high signal resolution and to enable RIM to achieve true resilience.

Waveform Distortion

Figure 7(a) shows how, for $B = 8$ and before saturation, effective J/N values coincide with the actual J/N used for interference simulation. However, this is not the case for lower values of B . We believe that this phenomenon is due to the waveform distortion introduced by a low number of quantization bits, which causes artifacts in the form of multiple frequency harmonics in the signal spectrum. This phenomenon is evidenced by the results presented in Figure 8, where the PSD and spectrogram of Jammer 1 are shown for $B = \{2, 8\}$. The interference waveform strongly differs from the baseband signal model, as shown in Table 2, when there is a low number of quantization bits. Additionally, in the case of $B = 2$, the power is spread across multiple frequencies, differing from the single tone obtained for $B = 8$.

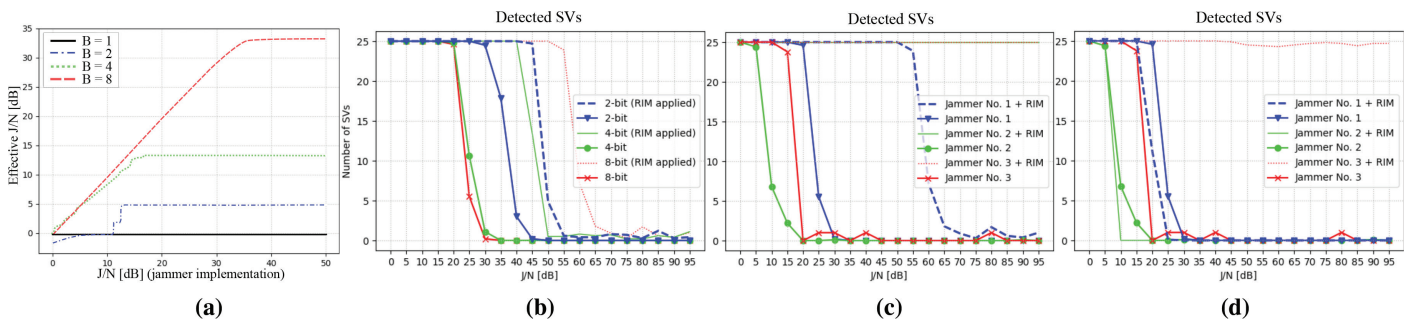


FIGURE 7 Results on the effect of signal quantization and snapshot RIM performance: (a) evolution of the effective J/N value at the receiver for $B = \{1, 2, 4, 8\}$ as a function of the actual J/N , (b) performance of snapshot RIM under Jammer 1 attack for $B = \{2, 4, 8\}$, with the number of detected SVs used as a metric, (c)–(d) comparison of frequency- and time-domain RIM frameworks for 8-bit quantization under three types of jamming, with detected SVs versus J/N used as a performance metric




FIGURE 8 PSD and spectrograms of 120-ms simulated GPS L1, Galileo E1, and BeiDou B1 signals with interference from Jammer 1 (see Table 2) for $B = \{2, 8\}$, $C/N_0 = 45$ dB, and $J/N = 30$ dB (a) PSD for $B = 2$ (b) PSD for $B = 8$ (c) Spectrogram for $B = 2$ (d) Spectrogram for $B = 8$

TABLE 4

Snapshot Variance Estimation Results with the MAD as a Robust Measure of Statistical Dispersion for $y[n] = r[n]$ (Clean Signal, Left Subtable) and $y[n] = r[n] + i[n]$ (Jammed Signal, Right Subtable), with $C/N_0 = 45$ dB and $J/N = 30$ dB. Signal quantization is performed with different numbers of bits $B = \{2, 4, 8\}$ over three different dynamic ranges. DR_1 , DR_2 , and DR_3 correspond to the dynamic ranges calculated from the variance of a snapshot with interference from Jammers 1, 2, and 3, respectively (see Table 2). After the ADC, both the signal and its variance are represented in unitless digital counts.

	Clean Signal						Interfered Signal							
	B	2		4		8		B	2		4		8	
Domain	Time	Freq.	Time	Freq.	Time	Freq.	Time	Freq.	Time	Freq.	Time	Freq.	Time	Freq.
DR_1	0	0.4906	0	0.0557	0.0021	0.0011	2.1731	0.0820	0.9860	0.0117	1.0987	0.0011		
DR_2	0	0.4906	0	0.0557	0.0021	0.0012	0	0.1020	0.9860	0.0199	1.0052	0.0012		
DR_3	0	0.4906	0	0.0557	0.0021	0.0052	2.1731	0.4245	0.2465	0.0482	0.0021	0.0052		

Quantization Gain

In Figure 7(b), the performance of snapshot RIM is assessed in the presence of Jammer 1 as a function of the J/N ratio for $B = \{2, 4, 8\}$. In this figure, the performance metric is the number of detected SVs. Results are presented for when RIM is both enabled and disabled. A gain of 10, 20, and 35 dB is observed when the RIM algorithm is applied for 2, 4, and 8 quantization bits, suggesting that the gain obtained when the RIM algorithm is applied is proportional to the number of quantization bits.

When RIM is not applied, the baseline performance is better for $B = 2$ in terms of detected SVs. Nevertheless, this improvement is due to the pseudo-resilience discussed earlier.

Impact on Robust Variance Estimation

In Table 4, snapshot variance estimation results with the MAD measure of statistical dispersion are shown with both clean and interfered signal for $B = \{2, 4, 8\}$. These results have been normalized based on the quantization gain values from Table 1. The simulated snapshot contains GPS L1, Galileo E1, and BeiDou B1 signals with a C/N_0 of 45 dB. Three different values of quantization dynamic range (DR) are considered, namely, DR_1 , DR_2 and DR_3 . These dynamic ranges correspond to the variances of snapshots interfered with Jammers 1, 2, and 3, respectively, for $J/N = 30$ dB (i.e., DR_l is proportional to the variance of a snapshot interfered by Jammer l). In each row of the *Interfered Signal* subtable, we present MAD results obtained when estimating the variance of the GNSS signal subject to interference from the jammer used to compute the dynamic range for that particular row. For

example, the row corresponding to DR₁ contains results for the signal with interference from Jammer 1. Given the robustness of the MAD estimation method, it is expected that when the variance of an interfered snapshot is estimated, only the contribution of useful GNSS signal has an impact on the result. In other words, we expect to obtain the same variance estimation results for both a clean and an interfered snapshot because the MAD measure should be able to cancel the contribution of the jamming component.

The results suggest that MAD offers robust results if the following two conditions apply: (1) the number of quantization bits is high enough to allow for full signal representation (i.e., $B \geq 8$) and (2) the MAD is applied in the domain where the interference exhibits a sparse representation. Considering this, results for $B = 8$ in the frequency domain match for all dynamic ranges, as the number of quantization bits is sufficiently high and the three jammers assume sparse representation in the frequency domain. Meanwhile, in the time domain and for $B = 8$, only results with Jammer 3 match between the clean and interfered signals, as this is the only jammer assuming a sparse representation in the time domain.

For $B = \{2, 4\}$ and in the time domain, all results in the *Clean Signal* subtable are equal to 0, indicating that the MAD measure discards all of the useful signal contribution as a consequence of the waveform distortion caused by low-bit quantization. For $B = \{2, 4\}$ and in the frequency domain, the MAD method intercepts the contribution of useful signal when applied to the clean snapshots. Better results are obtained in the frequency domain because samples are mapped to a wider range of values, allowing for more quantization levels than those defined by B . In the *Interfered Signal* column for the same configuration, the variance estimation values are lower, suggesting that the interference causes the MAD measure to mistake clean samples for interfered samples. Overall, high resolutions are necessary to discard the jamming component when estimating the snapshot variance using the robust MAD measure.

5.1.3 | Overview of Snapshot RIM Performance

An assessment of the snapshot RIM performance in the frequency domain for $B = \{2, 8\}$ is presented in Figure 9. The key performance indicators (KPIs) under study are the average C/N_0 , number of detected SVs, and PVT availability. C/N_0 values are plotted only when the receiver passes the integrity test. Overall, the results suggest that in the presence of interference, it is necessary to use high signal resolution to experience the resilience offered by the snapshot RIM framework, with the gain being proportional to B . In the analysis with Jammer 3, we observe that reducing the number of quantization bits can be slightly beneficial when no anti-jamming technique is available, but only in the presence of interferences that are sparse in the time domain. Nevertheless, the best performance is achieved with the highest signal resolution and with snapshot RIM enabled.

In the case of Jammer 1 without RIM, the number of detected SVs and PVT availability fall to 0 at $J/N = 30$ dB for $B = 8$ and at $J/N = 40$ dB for $B = 2$. The 10-dB gain provided by reducing the signal resolution is attributed to the pseudo-resilience experienced by the receiver, as discussed earlier. The fact that this resilience is not genuine becomes apparent when observing the average C/N_0 plots, which show equally poor performance when RIM is not enabled for both signal resolutions. When snapshot RIM is enabled, a gain of 30 dB in terms of average C/N_0 is observed for $B = 8$ with respect to when RIM is disabled, while the improvement for $B = 2$ is only 5 dB.

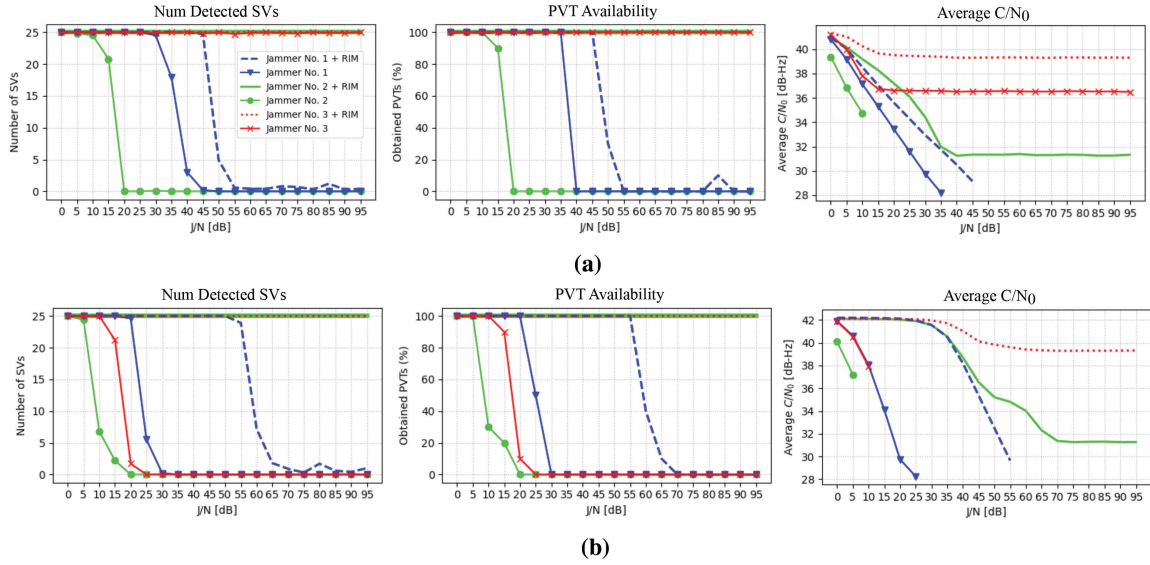


FIGURE 9 Assessment of the frequency-domain snapshot RIM algorithm for $B = \{2, 8\}$ (a) $B = 2$ (b) $B = 8$

The performance metrics under study are the number of detected SVs, PVT availability, and average C/N_0 . C/N_0 values are plotted only when the receiver passes the integrity test.

In the case of Jammer 2 without RIM, the number of detected SVs and PVT availability fall to 0 at $J/N = 20$ dB under both signal resolutions, although the metrics start degrading earlier for $B = 8$. Average C/N_0 values are similar under these two configurations for disabled RIM. When snapshot RIM is enabled, the receiver acquires all of the satellites with 100% PVT availability for $B = \{2, 8\}$ for the full range of tested jamming powers. This result is due to an actual improvement in signal quality only for the case of high signal resolution, i.e., $B = 8$, for which we observe a gain of 30 dB in terms of average C/N_0 . This gain is only 5 dB for $B = 2$. For $J/N > 40$ dB and $B = 8$, the high values of detected SVs and PVT metrics are not meaningful, as the average C/N_0 drops dramatically.

The clipping effect caused by a low number of quantization bits, as described in Section 2.2, slightly mitigates the impact of the interference component in the presence of the pulsed chirp. Jammer 3 assumes a sparse representation not only in the frequency domain but also in the time domain. Consequently, the clipping resembles the use of a ZMNL by setting a limit for the magnitude of the samples overpowering the receiver in the domain in which the interference shows a sparse nature, i.e., the time domain. Consequently, for $B = 2$ and disabled RIM, all satellites are acquired with 100% PVT availability and stable average C/N_0 values. For $B = 8$, the receiver performance begins to decrease for $J/N > 10$ dB when RIM is disabled. Nevertheless, in the presence of Jammer 3, the best overall performance is obtained for $B = 8$ when snapshot RIM is enabled, resulting in the acquisition of all satellites with 100% PVT availability and the maximum average C/N_0 for $J/N > 35$ dB.

Loss of Efficiency

Results have been obtained by averaging each satellite LoE and by setting Huber's nonlinearity threshold to $T_h = 1.345\sigma$ as in the remainder of the experiments. The measured LoE introduced by the frequency-domain snapshot RIM with MAD applied every 1 ms is 0.45, 0.48, and 0.49 dB for 2, 4, and 8 quantization

bits, respectively. These values differ from the theoretical loss of 0.29 dB. The difference of approximately 0.15 dB between the theoretical value and the empirical loss is considered to be small and is not a potential issue for satellite acquisition.

5.2 | Experiment 2

In this section, we report on a benchmark study between the baseline receiver and a COTS GNSS receiver in the presence of three different interference types (see Table 3). The setup used in this experiment is detailed in Section 4.3. The KPIs under study are the average C/N_0 and the number of detected SVs as a function of I/S values from 0 to 100 dB. The results show a significant performance improvement for the baseline receiver when snapshot RIM is used to process realistic RFI recordings, even when compared against the benchmark GNSS receiver.

5.2.1 | Scenario 1: CW

In Figure 10(a), we assess the performance of the baseline and benchmark receivers when processing the recording from scenario 1. An adaptation curve in terms of detected SVs is observed for the benchmark because it needs to obtain the navigation message, unlike the baseline snapshot receiver. The KPIs begin to decline at $I/S = 55$ dB for the benchmark. When the anti-jamming module is disabled, the baseline shows a modest improvement of approximately 2 dB compared with the benchmark results in terms of detected SVs. However, with the anti-jamming module activated, this improvement increases to approximately 20 dB for the baseline. Overall, the results indicate that the performance of the snapshot-based receiver surpasses that of the benchmark receiver in terms of average C/N_0 and the number of detected SVs.

5.2.2 | Scenario 2: Chirp

In Figure 10(b), we assess the performance of the baseline and benchmark receivers when processing the recording from scenario 2. In this scenario, the adaptation curve observed in the bottom subplot for the benchmark is more pronounced. For the benchmark receiver, C/N_0 values begin to decline at $I/S = 30$ dB; in contrast, this KPI remains stable until I/S values of 40 and 60 dB for the baseline with snapshot RIM disabled and enabled, respectively. Consequently, the gain provided by the baseline with snapshot RIM enabled is approximately 30 dB. Although the benchmark outperforms the baseline in terms of detected SVs when RIM is disabled under I/S values of 45–70 dB, the baseline is still able to acquire satellites for I/S values of 70–80 dB. In contrast, the benchmark performance drops significantly at an I/S ratio of 70 dB.

5.2.3 | Scenario 3: Pulsed Chirp

In Figure 10(c), we assess the performance of the baseline and benchmark receivers when processing the recording from scenario 3. Average C/N_0 values start decreasing at approximately $I/S = 60$ dB in the case of both receivers. The number of detected SVs starts decreasing at $I/S > 60$ dB for the benchmark and for the baseline with disabled RIM. However, for enabled RIM, the baseline is resilient to

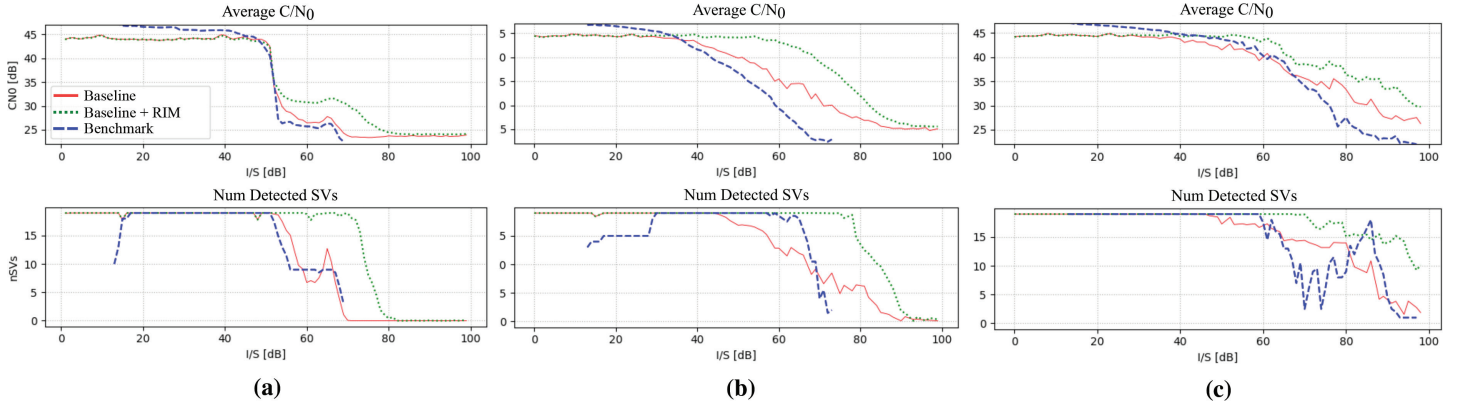


FIGURE 10 Results obtained in Experiment 2 when processing 200 ms of realistic recordings provided by ESA containing GNSS data with interference from the jammers described in Table 3 (a) Jammer 1 (b) Jammer 2 (c) Jammer 3

The baseline receiver performance is assessed with disabled and enabled snapshot RIM and compared with the performance of the benchmark receiver for different values of I/S . The metrics under study are the number of detected SVs and average C/N_0 .

I/S values up to 80 dB in terms of detected SVs. For $I/S > 80$ dB, the average C/N_0 provided by the benchmark is less than 25 dB, while the baseline is still able to provide average C/N_0 values greater than 30 dB when RIM is enabled and greater than 25 dB when RIM is disabled. For I/S values of 70–90 dB, there is a sudden improvement in performance in the benchmark results. This event is unexpected, and considering that the benchmark is a closed-loop receiver, it may indicate that the receiver can track some signals for a short period.

5.3 | 24-h Specifications

The performance of the proposed mitigation technique was validated by processing a 24-h RFI-free recording of L1/G1 and L5 bands in open-sky conditions, with the goal of computing KPIs such as the horizontal position accuracy, velocity accuracy, and TTFF. These KPIs are defined in Table 5. The 24-h test was conducted under three different configurations: one in which the anti-jamming module is disabled and two in which the anti-jamming module is enabled, with the application of frequency-domain RIM and DD-RIM, respectively. Results are presented in Table 6. In the case of the L1/G1 band, the horizontal position accuracy is 2.298 m both without RIM and with frequency-domain snapshot RIM. This value increases to 2.377 m with DD-RIM. In the L5 band, the worst accuracy is also obtained with DD-RIM, whereas the best accuracy is obtained with frequency-domain RIM. We do not consider the error introduced by DD-RIM to be a concern because this error is small and our previous experiments show that applying RIM in the frequency domain is sufficient to counter the jamming signal component. The velocity accuracy obtained in the L1/G1 band is 0.290 m/s with disabled RIM, 0.299 m/s with frequency-domain RIM, and 0.305 m/s with DD-RIM. These errors are considered to be sufficiently small. Moreover, no error is introduced by RIM in terms of velocity accuracy in the L5 band.

The collected results show that the use of the FFT and IFFT functions as the first and second linear transformations in the frequency-domain RIM accounts for the most important increase in TTFF. When frequency-domain RIM is applied in the L1 band, the TTFF increases by 0.25 s with respect to the case in which RIM is

TABLE 5
Definitions of Metrics Used for Receiver Performance Assessment in the 24-h Specification Experiment

Parameter	Definition
Horizontal Position Accuracy	Measures the positioning error in meters, represented by the circular error probable (CEP) at the 50% confidence level. The estimation is conducted under the following conditions: <ul style="list-style-type: none"> • Static receiver, 24 h, strong signal conditions, refresh rate of 5 min. More than 5 satellites per constellation within the period. • Oscillator: temperature-compensated crystal oscillator, stability 0.5 ppm.
Velocity Accuracy	Determines the velocity error, measured in meters per second and represented as the three-dimensional (3D) root mean square error (RMSE). The estimation is performed under the same conditions employed for the horizontal position accuracy.
TTFF	Processing time of a snapshot of specific length, measured in seconds. Computation time was obtained using the AWS Fargate engine (Amazon Web Services' serverless compute engine) configured with 4 GB of memory and 1 virtual CPU (vCPU).

TABLE 6
Impact of Frequency-Domain Snapshot RIM and DD-RIM on the Baseline Receiver in Terms of Horizontal Position Accuracy, Velocity Accuracy, and TTFF

	L1			L5		
	Disabled	Frequency	Dual	Disabled	Frequency	Dual
Horizontal Position Accuracy – CEP [m]	2.298	2.298	2.377	1.746	1.738	1.752
Velocity Accuracy - 3D RMSE [m/s]	0.290	0.299	0.305	0.268	0.267	0.268
TTFF [s]	9.75	10	10.01	15.59	15.9	15.93

disabled. However, only 0.01 s is added to this value when the time-domain RIM algorithm is also enabled with the DD-RIM. The same trend is observed in the L5 band, where the TTFF increases by 0.31 s when frequency-domain RIM is applied and only 0.03 s is added for DD-RIM. Although this could be a motivation to use time-domain RIM, frequency-domain RIM offers a stronger protection against interferences and better position and velocity accuracies. Furthermore, we consider these TTFF increases of 0.25 and 0.31 s to be sufficiently small. Together, the present findings suggest that the implemented anti-jamming module does not negatively affect the product performance.

6 | CONCLUSION

In this paper, the RIM framework is adapted for GNSS snapshot architectures. For low-power applications in areas such as IoT, including asset tracking, lower quantization levels (e.g., 2-bit ADC) are preferable, especially for manufacturers integrating snapshot modules into mobile platforms and SWaP devices. Consequently, we examined the impact of low-bit quantization on RIM performance, contrasting with other studies that typically assume sufficient resolution to fully represent the jamming signal. Snapshot RIM divides the snapshot samples into batches and applies Huber's ZMNL independently, with a robust measure of statistical dispersion, MAD, calculated recursively to set the ZMNL decision threshold. Two experiments were conducted to evaluate the snapshot RIM performance against three interference types: CW, chirp, and pulsed chirp. In the first experiment, we compared frequency- and time-domain RIM, finding that frequency-domain RIM is more effective against the three tested interference types, achieving a 35-dB gain in the presence of CW interference, compared with 5 dB for time-domain RIM. We

also showed that the use of fewer quantization bits introduces signal distortion and limits the effective J/N measured at the receiver. Snapshot RIM provides a gain proportional to the number of quantization bits, achieving 10-, 20-, and 35-dB gains in detected satellites for 2-, 4-, and 8-bit quantization in the presence of a CW jammer. With 8-bit quantization, RIM ensures 100% PVT availability and a 30-dB gain in C/N_0 . However, performance begins to degrade for $J/N > 40$ dB in the presence of CW interference. Overall, a higher number of quantization bits is preferred for optimal interference mitigation, as this enhances interference reconstruction, in line with the literature. Interestingly, reducing the number of quantization bits can be slightly beneficial when no anti-jamming technique is available, particularly for time-sparse interferences. In the second experiment, snapshot RIM demonstrated significant performance improvement, achieving up to a 20-dB gain in C/N_0 compared with a professional COTS GNSS receiver. A 24-h test confirmed that RIM can be implemented without compromising position accuracy or TTFF. Frequency-domain RIM increased the TTFF by 0.25 s in the L1 band and 0.31 s in the L5 band, whereas time-domain RIM added less than 0.05 s in both bands.

ACKNOWLEDGMENTS

This work has been partially supported by the ESA NAVISP Programme under activity NAVISP-EL2-096 and by the National Science Foundation under awards ECCS-1845833 and CCF-2326559.

REFERENCES

- Amin, M. G., Wang, X., Zhang, Y. D., Ahmad, F., & Aboutanios, E. (2016). Sparse arrays and sampling for interference mitigation and DOA estimation in GNSS. *Proceedings of the IEEE*, 104(6), 1302–1317. <https://doi.org/10.1109/JPROC.2016.2531582>
- Bamberg, T., Konovaltsev, A., & Meurer, M. (2023). Enabling RTK positioning under jamming: Mitigation of carrier-phase distortions induced by blind spatial filtering. *NAVIGATION*, 70(1). <https://doi.org/10.33012/navi.556>
- Betz, J. W. (2000). Effect of narrowband interference on GPS code tracking accuracy. *Proc. of the 2000 National Technical Meeting of the Institute of Navigation*, Anaheim, CA, 16–27. <https://www.ion.org/publications/abstract.cfm?articleID=3>
- Betz, J. W. (2001). Effect of partial-band interference on receiver estimation of C/N_0 : Theory. *Proc. of the 2001 National Technical Meeting of the Institute of Navigation*, Long Beach, CA, 817–828. <https://www.ion.org/publications/abstract.cfm?articleID=195>
- Borio, D. (2008). *A statistical theory for GNSS signal acquisition* [Doctoral dissertation, Politecnico di Torino]. https://www.ucalgary.ca/engo_webdocs/other/DBorio_Torino_Apr08.pdf
- Borio, D. (2017a). Myriad non-linearity for GNSS robust signal processing. *IET Radar, Sonar & Navigation*, 11(10), 1467–1476. <https://doi.org/10.1049/iet-rsn.2016.0610>
- Borio, D. (2017b). Robust signal processing for GNSS. *Proc. of the 2017 European Navigation Conference (ENC)*, Lausanne, Switzerland, 150–158. <https://doi.org/10.1109/EURONAV.2017.7954204>
- Borio, D., Camoriano, L., & Lo Presti, L. (2008). Two-pole and multi-pole notch filters: A computationally effective solution for GNSS interference detection and mitigation. *IEEE Systems Journal*, 2(1), 38–47. <https://doi.org/10.1109/JSYST.2007.914780>
- Borio, D., & Closas, P. (2018). Complex signum non-linearity for robust GNSS interference mitigation. *IET Radar, Sonar & Navigation*, 12(8), 900–909. <https://doi.org/10.1049/iet-rsn.2017.0552>
- Borio, D., & Closas, P. (2019). Robust transform domain signal processing for GNSS. *NAVIGATION*, 66(2), 305–323. <https://doi.org/10.1002/navi.300>
- Borio, D., Dovis, F., Kuusniemi, H., & Presti, L. L. (2016). Impact and detection of GNSS jammers on consumer grade satellite navigation receivers. *Proceedings of the IEEE*, 104(6), 1233–1245. <https://doi.org/10.1109/JPROC.2016.2543266>
- Borio, D., Li, H., & Closas, P. (2018). Huber's non-linearity for GNSS interference mitigation. *Sensors*, 18(7), 2217. <https://doi.org/10.3390/s18072217>
- Borio, D., O'Driscoll, C., & Fortuny, J. (2012). GNSS jammers: Effects and countermeasures. *Proc. of the 6th ESA Workshop on Satellite Navigation Technologies (NAVITEC 2012) & European Workshop on GNSS Signals and Signal Processing*, Noordwijk, Netherlands, 1–7. <https://doi.org/10.1109/NAVITEC.2012.6423048>
- Borio, D., O'Driscoll, C., & Fortuny, J. (2014). Tracking and mitigating a jamming signal with an adaptive notch filter. *Inside GNSS*. <https://insidengss.com/auto/marapr14-WP.pdf>

- Calatrava, H., Gusi-Amigó, A., Melman, F., & Closas, P. (2023). Effect of signal quantization on robust anti-jamming in snapshot receivers. *Proc. of the 36th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2023)*, Denver, CO, 2935–2948. <https://doi.org/10.33012/2023.19351>
- Díez-García, R., & Camps, A. (2019). Impact of signal quantization on the performance of RFI mitigation algorithms. *Remote Sensing*, *11*(17), 2023. <https://doi.org/10.3390/rs11172023>
- Ding, M., Chen, W., & Ding, W. (2023). Performance analysis of a normal GNSS receiver model under different types of jamming signals. *Measurement*, *214*, 112786. <https://doi.org/10.1016/j.measurement.2023.112786>
- Fernández-Hernández, I., & Borre, K. (2016). Snapshot positioning without initial information. *GPS Solutions*, *20*, 605–616. <https://doi.org/10.1007/s10291-016-0530-4>
- Fernández-Prades, C., Arribas, J., & Closas, P. (2016). Robust GNSS receivers by array signal processing: Theory and implementation. *Proceedings of the IEEE*, *104*(6), 1207–1220. <https://doi.org/10.1109/JPROC.2016.2532963>
- Fox, J. (2002). *An R and S-plus companion to applied regression*. Sage Publications. <https://collegepublishing.sagepub.com/products/an-r-and-s-plus-companion-to-applied-regression-1-17568>
- Gao, G. X. (2007). DME/TACAN interference and its mitigation in L5/E5 bands. *Proc. of the 20th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2007)*, Fort Worth, TX, 1191–1200. <https://www.ion.org/publications/abstract.cfm?articleID=7623>
- Gao, G. X., Sgammini, M., Lu, M., & Kubo, N. (2016). Protecting GNSS receivers from jamming and interference. *Proceedings of the IEEE*, *104*(6), 1327–1338. <https://doi.org/10.1109/JPROC.2016.2525938>
- García-Molina, J. A., & Crisci, M. (2016). Snapshot localisation of multiple jammers based on receivers of opportunity. *Proc. of the 8th ESA Workshop on Satellite Navigation Technologies and European Workshop on GNSS Signals and Signal Processing (NAVITEC)*, Noordwijk, Netherlands, 1–6. <https://doi.org/10.1109/NAVITEC.2016.7849324>
- Gioia, C., & Borio, D. (2021). Multi-layered multi-constellation GNSS interference mitigation. *NAVIGATION*, *70*(4). <https://doi.org/10.33012/navi.596>
- Grenier, A., Lohan, E. S., Ometov, A., & Nurmi, J. (2023). A survey on low-power GNSS. *IEEE Communications Surveys & Tutorials*, *25*(3), 1482–1509. <https://doi.org/10.1109/COMST.2023.3265841>
- Hegarty, C. J., Bobyn, D., Grabowski, J., & Van Dierendonck, A. (2020). An overview of the effects of out-of-band interference on GNSS receivers. *NAVIGATION*, *67*(1), 143–161. <https://doi.org/10.1002/navi.345>
- Janssen, T., Koppert, A., Berkvens, R., & Weyn, M. (2023). A survey on IoT positioning leveraging LPWAN, GNSS, and LEO-PNT. *IEEE Internet of Things Journal*, *10*(13), 11135–11159. <https://doi.org/10.1109/JIOT.2023.3243207>
- Kaplan, E. D., & Hegarty, C. (2005). *Understanding GPS/GNSS: Principles and applications* (3rd ed.). Artech House Publishers. <https://us.artechhouse.com/Understanding-GPSGNSS-Principles-and-Applications-Third-Edition-P1871.aspx>
- Li, H. (2023). *Robust signal processing with applications to positioning and imaging* [Doctoral dissertation, Northeastern University]. <https://doi.org/10.17760/D20487154>
- Li, H., Borio, D., & Closas, P. (2019). Dual-domain robust GNSS interference mitigation. *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2019)*, Miami, FL, 991–1002. <https://doi.org/10.33012/2019.16991>
- Li, H., Tang, S., Wu, P., & Closas, P. (2023). Robust interference mitigation techniques for direct position estimation. *IEEE Transactions on Aerospace and Electronic Systems*, *59*(6), 8969–8980. <https://doi.org/10.1109/TAES.2023.3312350>
- Li, S., Wang, F., Tang, X., Ni, S., & Lin, H. (2023). Anti-jamming GNSS antenna array receiver with reduced phase distortions using a robust phase compensation technique. *Remote Sensing*, *15*(17), 4344. <https://doi.org/10.3390/rs15174344>
- Linty, N. (2015). *Snapshot estimation algorithms for GNSS mass-market receivers* [Doctoral dissertation, Politecnico di Torino]. <https://doi.org/10.6092/polito/porto/2602171>
- Liu, X., Ribot, M., Gusi-Amigó, A., Closas, P., Rovira-García, A., & Sanz, J. (2020). RTK feasibility analysis for GNSS snapshot positioning. *Proc. of the 33rd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2020)*, 2911–2921. <https://doi.org/10.33012/2020.17768>
- Liu, X., Ribot, M. Á., Gusi-Amigó, A., Rovira-García, A., Sanz, J., & Closas, P. (2021). Cloud-based single-frequency snapshot RTK positioning. *Sensors*, *21*(11), 3688. <https://doi.org/10.3390/s21113688>
- Medina, D., Lass, C., Marcos, E. P., Ziebold, R., Closas, P., & García, J. (2019). On GNSS jamming threat from the maritime navigation perspective. *Proc. of the 2019 22nd International Conference on Information Fusion (FUSION)*, Ottawa, Ontario, 1–7. <https://doi.org/10.23919/FUSION43075.2019.9011348>

- Medina, D., Li, H., Vilà-Valls, J., & Closas, P. (2019). Robust statistics for GNSS positioning under harsh conditions: A useful tool? *Sensors*, 19(24), 5402. <https://doi.org/10.3390/s19245402>
- Medina, D., Ortega, L., Vilà-Valls, J., Closas, P., Vincent, F., & Chaumette, E. (2020). Compact CRB for delay, Doppler, and phase estimation—application to GNSS SPP and RTK performance characterisation. *IET Radar, Sonar & Navigation*, 14(10), 1537–1549. <https://doi.org/10.1049/iet-rsn.2020.0168>
- Meng, L., Yang, L., Yang, W., & Zhang, L. (2022). A survey of GNSS spoofing and anti-spoofing technology. *Remote Sensing*, 14(19), 4826. <https://doi.org/10.3390/rs14194826>
- Mitch, R. H., Dougherty, R. C., Psiaki, M. L., Powell, S. P., O'Hanlon, B. W., Bhatti, J. A., & Humphreys, T. E. (2011). Signal characteristics of civil GPS jammers. *Proc. of the 24th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2011)*, Portland, OR, 1907–1919. <https://www.ion.org/publications/abstract.cfm?articleID=9740>
- Mitch, R. H., Psiaki, M. L., Powell, S. P., & O'Hanlon, B. W. (2013). Signal acquisition and tracking of chirp-style GPS jammers. *Proc. of the 26th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2013)*, Nashville, TN, 2893–2909. <https://www.ion.org/publications/abstract.cfm?articleID=11419>
- Morales Ferre, R., De La Fuente, A., & Simona Lohan, E. (2019). Jammer classification in GNSS bands via machine learning algorithms. *Sensors*, 19(22), 4841. <https://doi.org/10.3390/s19224841>
- Morales Ferre, R., Richter, P., Falletti, E., De La Fuente, A., & Simona Lohan, E. (2020). A survey on coping with intentional interference in satellite navigation for manned and unmanned aircraft. *IEEE Communications Surveys & Tutorials*, 22(1), 249–291. <https://doi.org/10.1109/COMST.2019.2949178>
- Morton, Y. J., Van Diggelen, F., Spilker Jr, J. J., Parkinson, B. W., Lo, S., & Gao, G. (2021). *Position, navigation, and timing technologies in the 21st century: Integrated satellite navigation, sensor systems, and civil applications*. John Wiley & Sons. <https://doi.org/10.1002/9781119458449>
- Muthuraman, K., Brown, J., & Chansarkar, M. (2012). Coarse time navigation: Equivalence of algorithms and reliability of time estimates. *Proc. of the 2012 International Technical Meeting of the Institute of Navigation*, Newport Beach, CA, 1115–1138. <https://www.ion.org/publications/abstract.cfm?articleID=10010>
- Nardin, A., Minetto, A., Vouch, O., Mariani, M., & Dovis, F. (2023). Snapshot tracking of GNSS signals in space: A case study at lunar distances. *Proc. of the 36th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2023)*, Denver, CO, 3267–3281. <https://doi.org/10.33012/2023.19174>
- Parkinson, W. B., & Bartone, C. (2023). Enhancements enabled by multi-element antennas for GPS anti-jamming capabilities in civil applications. *Proc. of the 36th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2023)*, Denver, CO, 3731–3743. <https://doi.org/10.33012/2023.19368>
- Pullen, S., & Gao, G. (2012). GNSS jamming in the name of privacy: Potential threat to GPS aviation. *Inside GNSS*, 7(2), 33–43. <https://api.semanticscholar.org/CorpusID:14146225>
- Rousseeuw, P. J., & Croux, C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424), 1273–1283. <https://doi.org/10.1080/01621459.1993.10476408>
- Sun, Y., Chen, F., Lu, Z., & Wang, F. (2022). Anti-jamming method and implementation for GNSS receiver based on array antenna rotation. *Remote Sensing*, 14(19), 4774. <https://doi.org/10.3390/rs14194774>
- van der Merwe, J. R., Rügamer, A., Goicoechea, A. F.-D., & Felber, W. (2019). Blind spoofing detection using a multi-antenna snapshot receiver. *Proc. of the 2019 International Conference on Localization and GNSS (ICL-GNSS)*, Nuremberg, Germany, 1–7. <https://doi.org/10.1109/ICL-GNSS.2019.8752840>
- Van Diggelen, F. (2009). *A-GPS: Assisted GPS, GNSS, and SBAS*. Artech House. <https://us.artechhouse.com/A-GPS-Assisted-GPS-GNSS-and-SBAS-P1729.aspx>
- Van Diggelen, F., & Abraham, C. (2007). Coarse-time AGPS: Computing TOW from pseudorange measurements, and the effect on HDOP. *Proc. of the 20th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS 2007)*, Fort Worth, TX, 357–367. <https://www.ion.org/publications/abstract.cfm?articleID=7450>
- Vroom, A., Van den Oever, T., Gañez Fernandez, J., Van der Hijden, N., Zevenbergen, A., & Van der Hoeven, B. (2023). S-tracks: A secure snapshot-based solution for positioning and timing. *Engineering Proceedings*, 54(1), 31. <https://doi.org/10.3390/ENC2023-15457>
- Wang, X., & Poor, H. V. (1999). Robust multiuser detection in non-Gaussian channels. *IEEE Transactions on Signal Processing*, 47(2), 289–305. <https://doi.org/10.1109/78.740103>
- Wu, J., Tang, X., Huang, L., Yu, M., & Lin, H. (2023). Multipath suppression performance analysis of GNSS anti-jamming receiver. *Proc. of the 2023 6th International Conference on Information Communication and Signal Processing (ICICSP)*, Xi'an, China, 827–831. <https://doi.org/10.1109/ICICSP59554.2023.10390740>

- Wu, P., Calatrava, H., Imbiriba, T., & Closas, P. (2025). Federated learning of jamming classifiers: From global to personalized models. *NAVIGATION*, 72(1). <https://navi.ion.org/content/72/1/navi.688>
- Zhu, Y., Huang, S., Guo, K., Dan, Z., Fang, K., Wang, Z., & Wang, Y. (2023). An adaptive hybrid blanking algorithm to mitigate the DME pulse interference on BDS B2a receivers. *GPS Solutions*, 27(4), 164. <https://doi.org/10.1007/s10291-023-01495-5>
- Zidan, J., Adegoke, E. I., Kampert, E., Birrell, S. A., Ford, C. R., & Higgins, M. D. (2020). GNSS vulnerabilities and existing solutions: A review of the literature. *IEEE Access*, 9, 153960–153976. <https://doi.org/10.1109/ACCESS.2020.2973759>

How to cite this article: Calatrava, H., Gusi-Amigó, A., Melman, F., & Closas, P. (2025). Robust interference Mitigation in GNSS Snapshot Receivers. *NAVIGATION*, 72(2). <https://doi.org/10.33012/navi.699>