

Chasing Lightning: Detecting, Characterizing, and Identifying a Powerful Space-Based GNSS Interference Source

Zachary L. Clements*, Argyris Kriezis[†], and Todd E. Humphreys*

* *Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin*

[†] *Department of Aeronautics and Astronautics, Stanford University*

ABSTRACT

This paper analyzes and identifies a space-based Global Navigation Satellite System (GNSS) interference source that has caused scores of powerful transient wide-area interference events over continental Europe, Greenland, and Canada since 2019. While terrestrial or near-terrestrial sources are primarily responsible for the recent uptick in GNSS interference worldwide, space-based interferers are of special concern given their potential for vast geographic reach and their portent of a qualitative escalation in GNSS interference. Based on data collected between 2019 and 2026 from a network of terrestrial GNSS reference stations, this paper (1) develops a received-power-based detection framework; (2) details the spatial, temporal, and spectral patterns of wide-area interference events caused by the source; (3) presents and analyzes identification techniques that blend received-power and time-difference-of-arrival measurements; and (4) applies these techniques to confidently identify the GNSS interference source as a constellation of Russian early warning satellites in Molniya (“lightning”) orbits.

1 INTRODUCTION

Global Navigation Satellite Systems (GNSS) such as GPS provide meter-accurate positioning while offering global accessibility and all-weather, radio-silent operation. However, GNSS is fragile: its service is easily degraded by both deliberate (e.g., jamming and spoofing) and naturally occurring (e.g., multipath and atmospheric) interference (Scott, 2003; Humphreys et al., 2008; Humphreys, 2012, 2017). The past five years have seen a crescendo of GNSS disruptions in the aviation and maritime sectors (Osechas et al., 2022; GPS Spoofing Workgroup, 2024; Marcos et al., 2018; RIN Maritime GNSS Interference Working Group, 2026). Interference is especially problematic in these fields because GNSS is deeply interwoven into critical systems. Beyond simple navigational or timing displacement, cascading failures can occur when GNSS loss or corruption triggers follow-on problems in downstream systems (GPS Spoofing Workgroup, 2024; RIN Maritime GNSS Interference Working Group, 2026).

Fortunately, significant progress has been made over the past decades in GNSS interference mitigation and countermeasures (Psiaki and Humphreys, 2016, 2020; Borio et al., 2012; Clements et al., 2022). It has been shown that dedicated GNSS interference monitoring systems provide valuable spectrum situational awareness to enhance navigation security (Gunawardena et al., 2009; Gunawardena and Van Graas, 2011; Stader and Gunawardena, 2021; Kriezis et al., 2024; Sokolova et al., 2022; Morrison et al., 2023), with some demonstrating the ability to geolocate the interference source (Bhatti et al., 2012; Mitch et al., 2016). Furthermore, GNSS receivers situated in LEO enable terrestrial GNSS interference detection, characterization, and geolocation with worldwide coverage (Murrian et al., 2021; Clements et al., 2023, 2026).

Networks of terrestrial GNSS receivers can also be enlisted for GNSS interference monitoring (Abraha et al., 2024), including of space-based interference. York et al. (2014) revealed that GPS PRNs 24 and 27 (SVN65 and SVN66) were transmitting leakage tones at ± 10 MHz and ± 20 MHz from the GPS L5 frequency. The anomaly was subsequently fixed. More recently, a network of terrestrial reference stations was used to identify a BeiDou satellite (NORAD ID 40749) as a source of interference in the B3I band (1268.52 MHz) (Patil et al., 2023, 2024). The offending satellite continuously transmitted a tone centered at 1268.52 MHz and tones at ± 10 MHz and ± 20 MHz from this frequency. These tones are no longer present at the time of writing. Both of these interference cases were apparently due to faulty or improperly configured hardware.

This paper reports on the detection, characterization, and source identification of powerful transient GNSS interference originating from a non-GNSS satellite. The phenomenon was briefly mentioned in European Space Agency NAVISP-EL3-014 (2023), and more fully characterized in a preliminary conference publication by the current authors, Clements and Humphreys (2025). The current paper significantly extends both works. The interference reported here is different from that detailed in Patil et al. (2023, 2024); York et al. (2014) in several respects: (1) it is transient rather than continuous—the duration of interference events is less than 10 seconds; (2) it affects signals in the widely used GPS L1 frequency band rather than the less-common L5 and B3I bands; (3) it is much more powerful, causing drops in terrestrial receiver carrier-to-noise ratios (CNRs) of up to 10 dB; and

(4) it is not merely an anomaly of existing GNSS signals, nor does it originate from a GNSS satellite.

Obviously, high-powered interference with continental reach affecting the GPS L1 band—the primary band used for global aviation, shipping, and precise timing—is of serious concern. If deliberate, it portends a qualitative escalation in GNSS interference.

The effects of this interference are evident in public data from a network of terrestrial reference stations operated by the International GNSS Service (IGS) (Dow et al., 2009; Kouba, 2009; Montenbruck et al., 2017; Johnston et al., 2017). Most useful are data from a subset of IGS stations that produce high-rate (1-Hz) GNSS observables. These data may be retrieved from the Crustal Dynamics Data Information System archive, which is made available through NASA’s archive of space geodesy data (Noll, 2010). On scores of occasions since 2019, all tracked signals on the GPS L1 frequency at IGS reference stations across Europe, Greenland, and Canada simultaneously saw a sudden brief drop in CNR. The onset of the disruption was synchronous to within the 1-Hz sampling resolution, suggesting a single source per event. The affected terrestrial receivers span a geographic area so large that no single ground-based or aircraft-based source could reach them all; hence the space-based origin hypothesis.

While CNR variations across a receiver network are a valuable metric for space-based interference detection, it will be shown that they are not sufficiently informative for unique source identification. By contrast, a source’s position and velocity can be instantaneously estimated using time- and frequency-difference-of-arrival (T/FDOA) techniques if a spatially diverse network of four or more ground stations simultaneously captures raw broadband samples in the band affected by the interference. Moreover, by referencing a catalog of satellite ephemerides (assuming the source satellite is listed), even two stations are adequate to narrow the possibilities to a manageable few. An association framework that fuses both CNR and TDOA measurements can then allow a unique identification, as will be shown.

This paper makes six main contributions. First, it presents measurement models and a detection framework for transient wide-area interference. Second, it details the spatial, temporal, and spectral properties of multiple wide-area GNSS outage events from the space-based interference source and distinguishes these from a naturally occurring solar radio burst. Third, it presents a basic satellite identification strategy to narrow the candidate satellites and estimate the minimum satellite altitude at apogee. Fourth, it presents an advanced satellite associate framework using the generalized likelihood ratio test (GRLT) and applies the framework to a test scenario. Fifth, it presents a framework for instantaneously identifying an interference satellite based on a brief time history of TDOA measurements, and gives an error sensitivity analysis. Sixth, it combines IGS CNR data and raw wideband samples from two additional receivers in Europe to confidently identify the source, which is revealed to be a small constellation of Russian satellites in Molniya (“lightning”) orbits.

Compared to the preliminary conference version of this paper in Clements and Humphreys (2025), contributions two and three are extended and contributions four, five, and six are novel.

2 MEASUREMENT MODELS AND DETECTION

The IGS reference station network collects and provides observables in the Receiver Independent Exchange Format (RINEX) (Johnston et al., 2017). These GPS-time-tagged observables include carrier phase, pseudorange, Doppler, and CNR measurements for each tracked GNSS satellite. CNR observables from GPS L1 C/A signals at stations providing high-rate (1-Hz) GNSS observables are the focus of the following analysis.

2.1 Measurement Models

Let \mathcal{I} be the set of all relevant terrestrial reference stations (primarily IGS stations), and let \mathcal{J} be the set of all unique GNSS signals that can be tracked by such stations at a given center frequency (e.g., the GPS L1 frequency). By way of notation, let $|\mathcal{I}|$ denote the cardinality of \mathcal{I} , and similarly for all other sets. For station $i \in \mathcal{I}$ and signal $j \in \mathcal{J}$, let CNR_{ij} denote the true CNR, expressed as

$$\text{CNR}_{ij} = P_{\text{R}}^{ij} - N_0^i \quad (\text{dB-Hz}) \quad (1)$$

$$P_{\text{R}}^{ij} = P_{\text{T}}^j + G_{\text{T}}^j \left(\theta_{\text{T}}^{ij}, \phi_{\text{T}}^{ij} \right) + G_{\text{R}}^i \left(\theta_{\text{R}}^{ij}, \phi_{\text{R}}^{ij} \right) + L_{ij} \quad (\text{dBW}) \quad (2)$$

where P_{R}^{ij} is the received power (dBW), N_0^i is the thermal noise density (dBW/Hz), P_{T}^j is the transmit power (dBW), G_{T}^j is the transmitter antenna’s gain function (dB), G_{R}^i is the receiver antenna’s gain function (dB), θ_{R}^{ij} and θ_{T}^{ij} are the off-boresight angles at the receiver and transmitter, ϕ_{R}^{ij} and ϕ_{T}^{ij} are the azimuth angles at the receiver and transmitter, and L_{ij} is the path loss (dB),

defined as

$$L_{ij} = 20 \log_{10} \left(\frac{\lambda}{4\pi\rho_{ij}} \right) \quad (\text{dB}) \quad (3)$$

where λ is the carrier wavelength and ρ_{ij} is the range to the satellite transmitting signal j . The noise power density N_0^i also includes the effects of multi-access interference from GNSS signals other than the j th one on the same frequency.

CNR becomes the carrier-to-interference-and-noise ratio (CINR) when there is at least one interference signal present. This paper assumes a single interference signal and source at each epoch. Let CINR_{ij} denote the true CINR for $i \in \mathcal{I}$ and $j \in \mathcal{J}$, expressed as

$$\text{CINR}_{ij} = P_{\text{R}}^{ij} - 10 \log_{10} \left(\tilde{N}_0^i + \tilde{I}_0^i \right) \quad (\text{dB-Hz}) \quad (4)$$

where \tilde{N}_0^i is the linear unit equivalent of the noise power density N_0^i , satisfying $N_0^i = 10 \log_{10}(\tilde{N}_0^i)$, and \tilde{I}_0^i is the linear unit equivalent of the interference power density I_0^i (dBW/Hz), satisfying $I_0^i = 10 \log_{10}(I_0^i)$. Following Humphreys (2017), let $I^i(t)$ be the interference component of the product between the received signal (including the interference and thermal noise) and the local replica of the desired signal, and let $S_1^i(f)$ be its power spectral density (PSD, dBW/Hz). Then I_0^i may be defined as

$$I_0^i = S_1^i(0) \quad (\text{dBW/Hz}) \quad (5)$$

If $\tilde{S}_{\text{C}}^i(f)$ is the PSD of the local replica signal's spreading code, and $\tilde{S}_{\text{r}_1}^i(f)$ is the PSD of the received interference signal—both in linear units—and \hat{f}_{D}^i is the receiver's estimate of the desired signal's apparent Doppler frequency (Hz), then $\tilde{S}_1^i(f)$, the linear unit equivalent of $S_1^i(f)$, is formed by the convolution

$$\tilde{S}_1^i(f) = \tilde{S}_{\text{C}}^i(f) \star \tilde{S}_{\text{r}_1}^i(f + \hat{f}_{\text{D}}^i) \quad (6)$$

Similar to (2), let P_1^i denote the received interference signal's power at the i th station, given by

$$P_1^i = P_1 + G_1(\theta_1^i, \phi_1^i) + G_{\text{R}}^i(\theta_{\text{R}}^{i1}, \phi_{\text{R}}^{i1}) + L_{i1} \quad (\text{dBW}) \quad (7)$$

where P_1 is the transmitted interference power and G_1 is the interference source antenna's gain function. The arguments of G_1 and G_{R}^i , and the quantity L_{i1} are as before except that they apply to the interference source rather than to the j th GNSS signal.

At the k th time epoch, the receiver at the i th station reports the measurement $z_{ij}[k]$, which is either the measured CINR_{ij} under nominal operating conditions (the null hypothesis H_0), or is the measured CINR_{ij} under interference (the alternate hypothesis H_1), for all $j \in \mathcal{J}$. The reported measurement $z_{ij}[k]$ is modeled under H_0 and H_1 as

$$H_0 : z_{ij}[k] = \text{CINR}_{ij}[k] + w_{ij}[k] \quad (8)$$

$$H_1 : z_{ij}[k] = \text{CINR}_{ij}[k] + w_{ij}[k] \quad (9)$$

where $w_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2)$ is zero-mean additive white Gaussian noise (AWGN) that models measurement error due to thermal noise, atmospheric effects, multipath, and other minor effects.

2.2 Detection

Determining whether transient interference is present or not at each station, and across a network of stations, is an exercise in detection theory. The problem is to identify an interference event that typically persists for 3 to 5 seconds. Detection techniques from the nascent transient change detection literature are well suited to this problem (Guépié et al., 2012; Moustakides, 2014; Egea-Roca et al., 2018). During a transient interference event, the measurement $z_{ij}[k]$ transitions from H_0 to H_1 —due to a sudden injection of I_0^i —and then back to H_0 as the event concludes. An optimal transient change detector, such as proposed in Guépié et al. (2012), is matched to the interference-induced pattern in $z_{ij}[k]$. But in cases like the present one in which the interference interval is unknown *a priori*, no optimal solution is available. Egea-Roca et al. (2018) recommend a windowed solution, and derive certain performance bounds. The current paper approximates the windowed solution with a simple

differencing technique that is more amenable to analysis and becomes optimal for single-epoch change durations (Moustakides, 2014).

Let T_{\max} be the maximum duration of the transient event, and Δt be the measurement epoch ($\Delta t = 1$ second in the current context). Set $l = \lceil T_{\max}/\Delta t \rceil + 1 \in \mathbb{N}$ as the detector's index stride. Then for the k th time-epoch, i th station, and j th tracked GNSS signal, let the signal-specific detection statistic be given by

$$\xi_{ij}[k] = \frac{1}{2} (z_{ij}[k+l] - 2z_{ij}[k] + z_{ij}[k-l]) \quad (10)$$

Under H_0 and H_1 , $\xi_{ij}[k]$ takes on distributions

$$H_0 : \xi_{ij}[k] \sim \mathcal{N}(0, \frac{3}{2}\sigma_{ij}^2) \quad (11)$$

$$H_1 : \xi_{ij}[k] \sim \mathcal{N}(\mu_i, \frac{3}{2}\sigma_{ij}^2) \quad (12)$$

where the mean μ_i is

$$\mu_i = 10 \log_{10} \left(\frac{\tilde{N}_0^i + \tilde{I}_0^i}{\tilde{N}_0^i} \right) \quad (\text{dB}) \quad (13)$$

and can be calculated if all relevant parameters are known.

Let $\mathcal{J}_i[k] \subset \mathcal{J}$ be the set of GNSS signals tracked by station i at epoch k . Then the station-specific interference detection statistic $\Lambda_i[k]$ is the across-signal average of $\xi_{ij}[k]$:

$$\Lambda_i[k] = \frac{1}{|\mathcal{J}_i[k]|} \sum_{j \in \mathcal{J}_i[k]} \xi_{ij}[k] \quad (14)$$

The distribution of $\Lambda_i[k]$ under H_0 and H_1 is

$$H_0 : \Lambda_i[k] \sim \mathcal{N}(0, \sigma_i^2) \quad (15)$$

$$H_1 : \Lambda_i[k] \sim \mathcal{N}(\mu_i, \sigma_i^2) \quad (16)$$

The quantity μ_i amounts to the average drop in observed CNR at the i th station due to the interference event, in dB. In practice, the parameters that contribute to \tilde{I}_0^i are unknown *a priori*, leaving μ_i unknown. But because $\tilde{I}_0^i > 0$, it follows that $\mu_i > 0$, which allows it to be cancelled from each site's detection statistic, making a test based on (14) uniformly most powerful (Van Trees, 2001). The value σ_i^2 can be estimated at the i th station using historical data, and from this estimate a detection threshold ν_i can be calculated for a constant false alarm rate (CFAR). The hypothesis test for the i th station then becomes

$$\Lambda_i[k] \underset{H_0}{\overset{H_1}{\geq}} \nu_i \quad (17)$$

Null-hypothesis (H_0) detection statistic distributions from a day's worth of data for two stations, METG and MATE, are shown in Fig. 1. Both appear to be zero mean and approximately Gaussian distributed. The modest difference in σ_i^2 can be attributed to local noise, different receivers and antennas, and quantization of the reported CNR.




Figure 1: The distributions of Λ_i under H_0 (no interference present) for stations METG ($i = 1$) and MATE ($i = 2$).

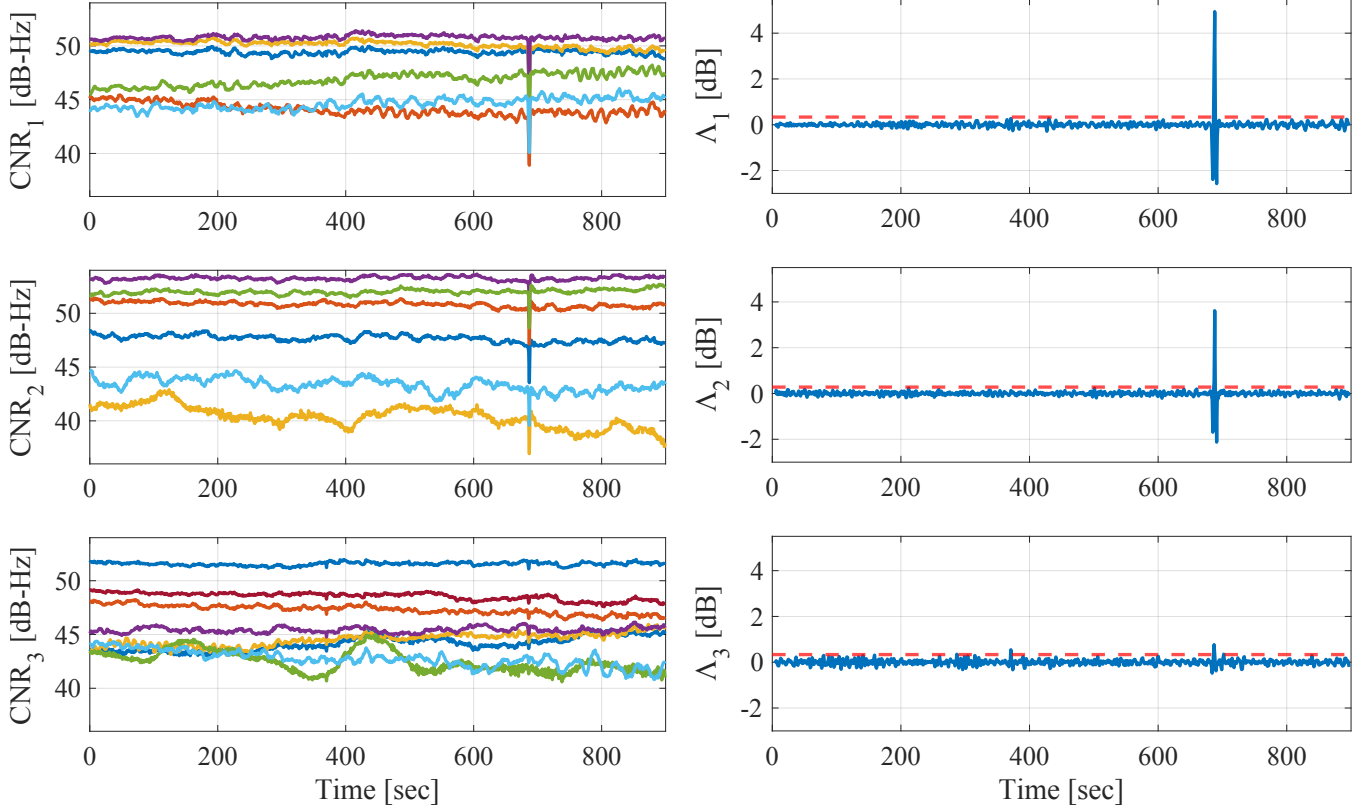


Figure 2: Reported CNR (left) and the detection statistic Λ_i (right) for IGS stations METG ($i = 1$, Finland), MATE ($i = 2$, Italy), and THU2 ($i = 3$, Greenland) over a 15-minute interval on day 160 of year 2021. The dashed red line is the detection threshold with a 10^{-4} probability of false alarm.

The remainder of this section focuses on day 160 of year 2021. Fig. 2 shows the time history of $\{z_{ij}\}$ and $\{\Lambda_i\}$ over a 15-minute interval for three stations with $l = 3$, amounting to a 3-second stride at the 1-Hz measurement rate. The dashed red lines in the right plots indicate the detection thresholds corresponding to a 10^{-4} probability of false alarm. Around the 700-second mark, interference is detected simultaneously at all three stations. While individual interference detections are not uncommon because reference stations occasionally experience local interference, time alignment of an interference event across geographically distant sites in Finland, Italy, and Greenland is noteworthy. Moreover, the interference exhibits the same characteristic signature at each location: a CNR drop lasting approximately 3 seconds. This suggests that all three stations were affected by a common interference source.

The detection test was conducted continuously across stations spanning Europe, Greenland, and Canada on day 160 of 2021. Fig. 3 shows the number of stations registering detections during this period with a 10^{-4} probability of false alarm. A lower-power interference event was detected by 21 stations, followed by a more powerful event detected by 58 stations. The stronger event corresponds to the 700-second mark shown in Fig. 2. The recorded CNR observables at each station align with the patterns depicted in Fig. 2, providing further evidence that these stations were affected by the same interference source. A heat map showing the spatial distribution of the CNR reduction is presented in Fig. 4. During the more powerful event, the tracked CNR of GPS L1 C/A signals exhibited drops as large as 6 dB, while the weaker event produced drops up to 1.5 dB.




Figure 3: The number of stations that detected interference during day 160 of year 2021. The expanded view on the right shows a lower-power event detected by 21 stations, followed by a higher-power event detected by 58 stations.

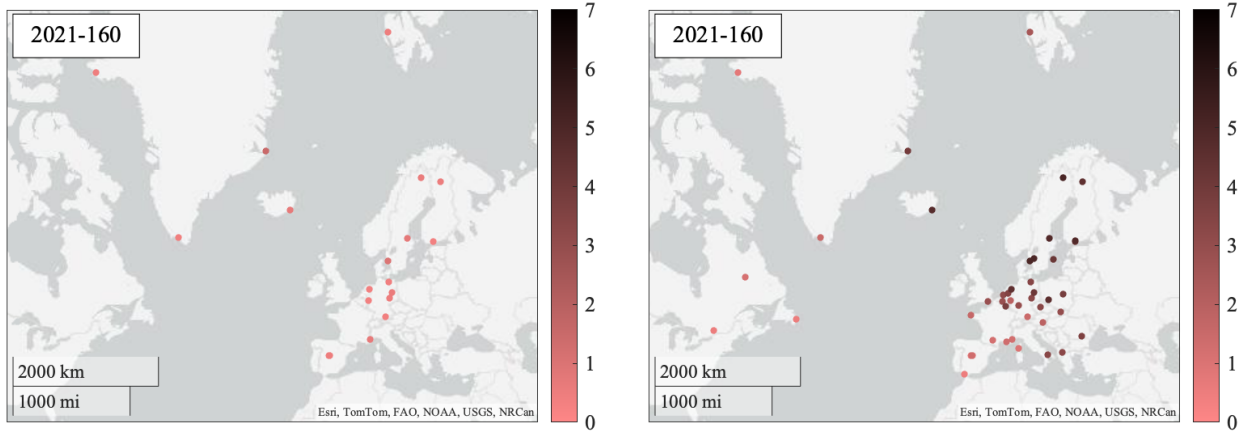


Figure 4: Heat map of the test statistic (in dB) at triggering stations for the lower-power event (left) and the higher-power event (right). The drop in GPS L1 C/A CNR during the more powerful interference event was as large as 6 dB, centered near the Baltic region.

3 INTERFERENCE PROPERTIES

This section details the transient space-based interference’s temporal, spatial, and spectral properties over the seven-year period from January 2019 to April 2026. For comparison, it also presents data from a naturally occurring solar radio burst.

3.1 Temporal Patterns

High-rate (1-Hz) data from 165 reference stations collected over the seven-year period were retrieved and analyzed. The detection hypothesis test described earlier was applied to every station on an epoch-by-epoch basis from January 1, 2019 to May 4, 2026. If interference was detected simultaneously at several stations, then an event detection was declared. A total of 75 days were identified with at least one wide-area transient GNSS interference event on the GPS L1 frequency during which at least one station experienced a CNR drop of 5 dB or greater. The annual number of such occurrences is summarized in Table 1. The earliest detection of a significant transient wide-area interference event within this period occurred in October 2019. Fig. 5 presents a histogram showing the distribution of high-power interference events by day of the week and hour of the week. Notably, these predominantly occurred during business days and business hours (UTC time), which suggests human involvement: a purely random phenomenon would tend to exhibit a temporally uniform distribution.

There were an 47 additional days with weaker widespread transient interference events. On these days, the largest CNR drop did not exceed 5 dB. If these days are included in Fig. 5, the distributions shown do not change appreciably. The weaker interference events are left out of this paper’s analysis to avoid false alarms at the periphery of the affected area.

Table 1: The number of days per year that saw at least one wide-area transient GNSS interference event on the GPS L1 frequency with at least one station suffering a drop of 5 dB or greater.

	2019	2020	2021	2022	2023	2024	2025	2026	Total
Count	8	17	5	7	15	9	12	2	75




Figure 5: Distribution of the day of the week and hour of day (with respect to UTC) during which interference events with at least one station suffering a drop of 5 dB or greater occurred. Clearly, the high-power interference events typically occur during business days and business hours.

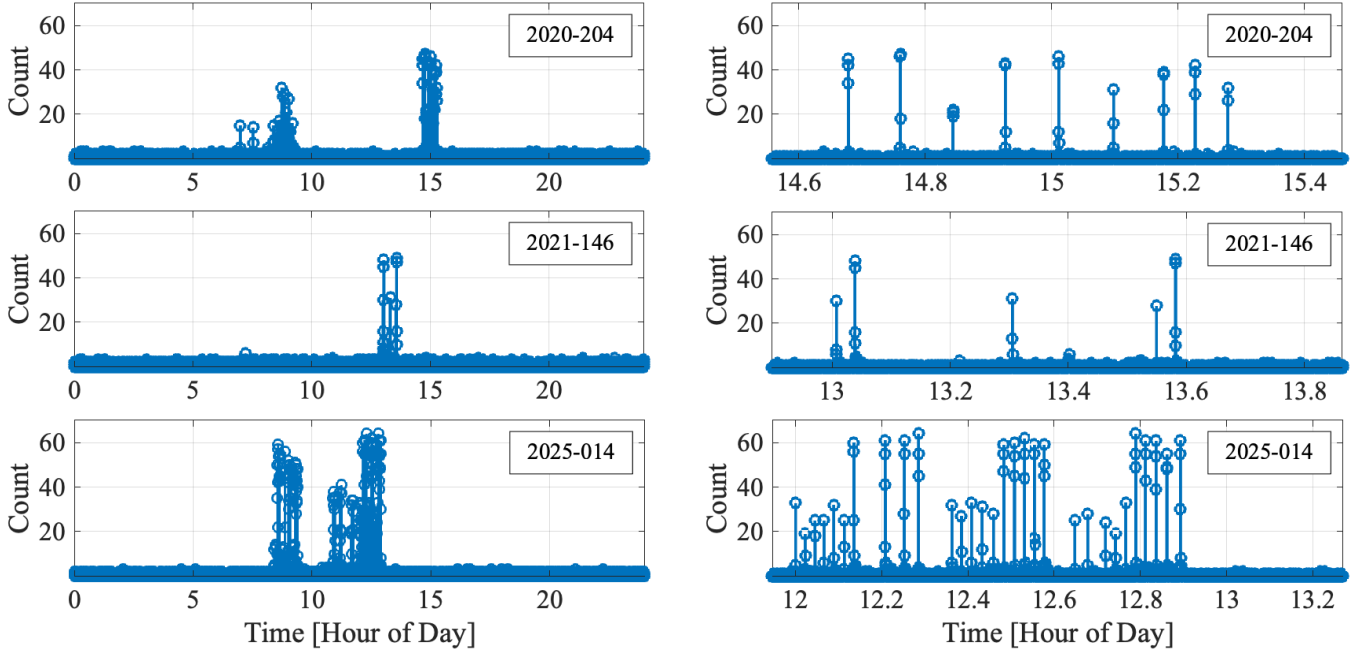


Figure 6: The number of stations that detected interference during three different days. Interference events can occur several times a day. The left column shows the entire day, while the right column is an expanded sub-interval.

Fig. 6 presents the number of daily detections across three additional example days. Unlike the case in Fig. 3, which shows a single high-power interference event, some days experienced multiple high-power events. Notably, day 146 of 2021 exhibited an interference signature similar to that of day 160 of the same year, with a low-power burst followed by a high-power burst. This happened twice on day 146, with the strong interference events being separated by approximately 32.6 minutes. The time delay between the low-power and high-power burst was 317 seconds on day 160, and 115 seconds for both events on day 146. This signature was observed on several other days. Day 204 of 2020 and day 014 of year 2025 recorded numerous high-power interference bursts. The timing of the bursts within a single day is typically periodic, with large bursts often spaced by an integer multiple of 150 seconds. The daily temporal patterns shown here of the high-power interference bursts are broadly representative of the phenomenon across all 75 days.

3.2 Spatial Patterns

Across all 75 days on which high-power events occurred, GNSS receivers in Europe were the most affected, with the Baltic region consistently experiencing the largest CNR drops. Fig. 7 shows example heat maps from day 146 of 2021 and day 014 of 2025. On day 146, two high-power interference events occurred, both producing nearly identical spatial patterns. Day 014 of 2025 also saw multiple high-power bursts, all with consistent patterns matching the example shown. The largest CNR drop across all events was 10 dB recorded at the LAMA station in Poland in 2025. Notably, during wide-area interference events in Europe, no similar disruptions were detected elsewhere in the world.

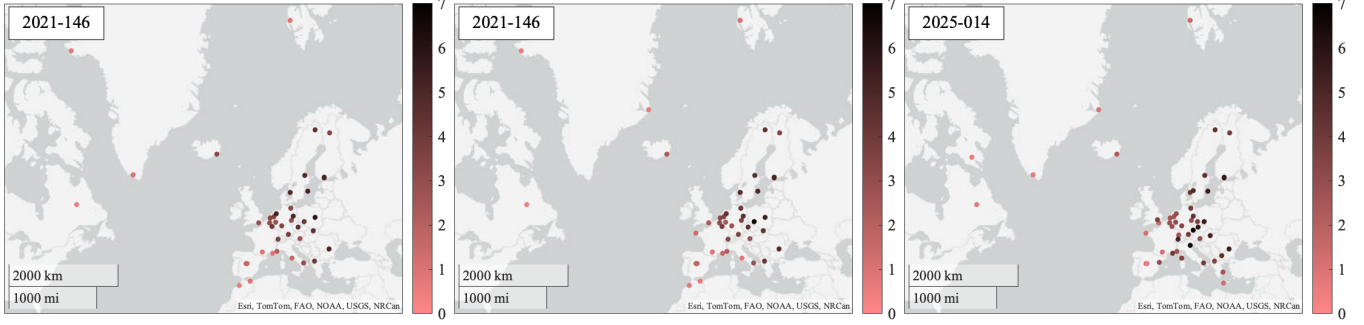


Figure 7: Detection statistic heat maps, in dB, for day 146 of year 2021 and day 014 of year 2025. GNSS receivers in the Baltic region are impacted most, which is generally representative of nearly all of the interference events.

Although the overwhelming majority of interference events saw receivers in the Baltic region impacted most, day 204 of 2020 exhibited a distinct interference pattern compared to other days. On this day, there was progressive movement in the geographic center of interference over multiple events, starting in the Baltic Sea and then moving into Germany and on to the Norwegian Sea, all over a 20-minute interval (Clements and Humphreys, 2025). This deviation may be attributed to satellite motion, to a change in the interference source’s beam pointing vector, or to multiple active satellite sources.

3.3 Spectral Properties

Spectral data were collected using a u-blox F9P GNSS receiver connected to a Trimble GNSS antenna installed in Gdynia, Poland, operating as an experimental radio frequency interference (RFI) monitoring station. Direct estimation of the absolute received interference power at the antenna reference point is not possible because of the unknown loss between the antenna and the receiver front end. Accordingly, the analysis here focuses on relative comparisons of the PSD across interference events and against interference-free baseline conditions. The u-blox receiver provides uncalibrated 1-Hz spectral observations through the “SPAN” message. Although SPAN output measurements are dimensionless, they are nonetheless suitable for comparative analysis. In previous work, a methodology was developed to transform these uncalibrated SPAN observations into a spectrally adjusted power density metric (Kriezis et al., 2025).

Fig. 8 presents the raw SPAN spectral output under nominal (RFI-free) conditions (black line) and during 48 strong interference events from 2024 to 2025. The interference spectrum exhibits a consistent spectral shape across all recorded events on days for which data are available. The interference peak is centered at 1577.5 MHz, about 2 MHz above the GPS L1 center frequency of 1575.42 MHz, and has an approximate bandwidth of 5 MHz.

In addition to tracked GPS L1 C/A signals, tracked Galileo E1 and BeiDou B1C/B1A signals also exhibited a concurrent drop in CNR during interference events. The magnitude of the decrease at each station closely matched the drop of the GPS L1 C/A signals, which might be expected because they share the same center frequency. The drop in CNR is not identical because the GPS, Galileo, and BeiDou spreading codes each have a different PSD $\tilde{S}_C^i(f)$ which, as shown in (6), leads to a different interference component PSD $\tilde{S}_I^i(f)$ and ultimately to a different interference power density I_0^i .

During strong 1577.5-MHz-centered interference events, a small but noticeable drop in CNR is also observed on tracked BeiDou B1I signals, which are centered at 1561.098 MHz. This is a puzzling because the interference spectrum shown in Fig. 8 has no spectral overlap with the BeiDou B1I signal, from which it follows that $I_0^i = S_i^i(0) = 0$. Apparently, the method by which receivers in the IGS network estimate the noise floor N_0^i that they use to report CNR values is somewhat sensitive to interference near L1.

Occasionally, an initial CNR drop and recovery in tracked GPS L1 C/A signals was followed by an equal-magnitude CNR drop and recovery in tracked BeiDou B1I signals. This pattern was observed on 15 of the 75 days, the first in June 2020. This indicates that the interference source can generate signals at 1577.5 MHz and near 1561.098 MHz. Fig. 8 only shows the interference centered at 1577.5 MHz, but Fig. 9, which shows PSDs derived from raw wideband samples captured in February 2026 by a receiver in the Netherlands clearly shows both interference bands, with the lower band centered at 1558.5 MHz. Interestingly, these two interference bands never appear to be active at the same time.

No wide-area transient interference has yet been observed near the GPS L2 or L5 bands.




Figure 8: Uncalibrated average PSD near the GPS L1 band during nominal operation (thick black line) and 1-Hz PSD estimates during 48 transient interference events, as recorded in Gdynia, Poland.




Figure 9: Power spectrum derived from raw wideband samples captured in Amsterdam, Netherlands, during an interference event on February 11, 2026. An initial burst of interference at 1577.5 MHz was followed by a burst at 1558.5 MHz.

3.4 Comparison with Solar Radio Burst

It is worth noting that naturally occurring phenomena can also cause significant CNR reduction over a large geographic area. Cerruti et al. (2006, 2008) examined the impact of intense solar radio bursts on tracked GNSS signals using observational data from the IGS network during solar events in the mid 2000s. It was demonstrated that solar radio bursts can cause CNR degradation by as much as 25 dB across the sunlit side of Earth. But solar radio bursts are qualitatively different from the transient interference studied in this paper: They are typically broadband, evolve more slowly than the rapid onset observed in Fig. 3, and the associated CNR reduction typically persists for a longer duration.

On November 11, 2025, a strong solar flare occurred with a magnitude of X5.1, causing a geomagnetic storm on Earth reaching category G4 (“severe”) (European Space Agency, 2025). IGS stations positioned on the sunlit side of Earth experienced drops in CNR on all GNSS frequencies. Fig. 10 shows the CNR time history of tracked GPS L1, L2, and L5 signals from the SUTM reference station in South Africa. Signals on L2 and L5 were degraded as much as 17 dB for hundreds of seconds. The CNR of tracked signals at other IGS stations were degraded in a similar manner. Fig. 10 shows the magnitude of the CNR drop on the tracked L2 signals across the globe. Clearly, the effects of solar radio bursts manifest differently in the IGS data compared to the transient phenomenon studied here.

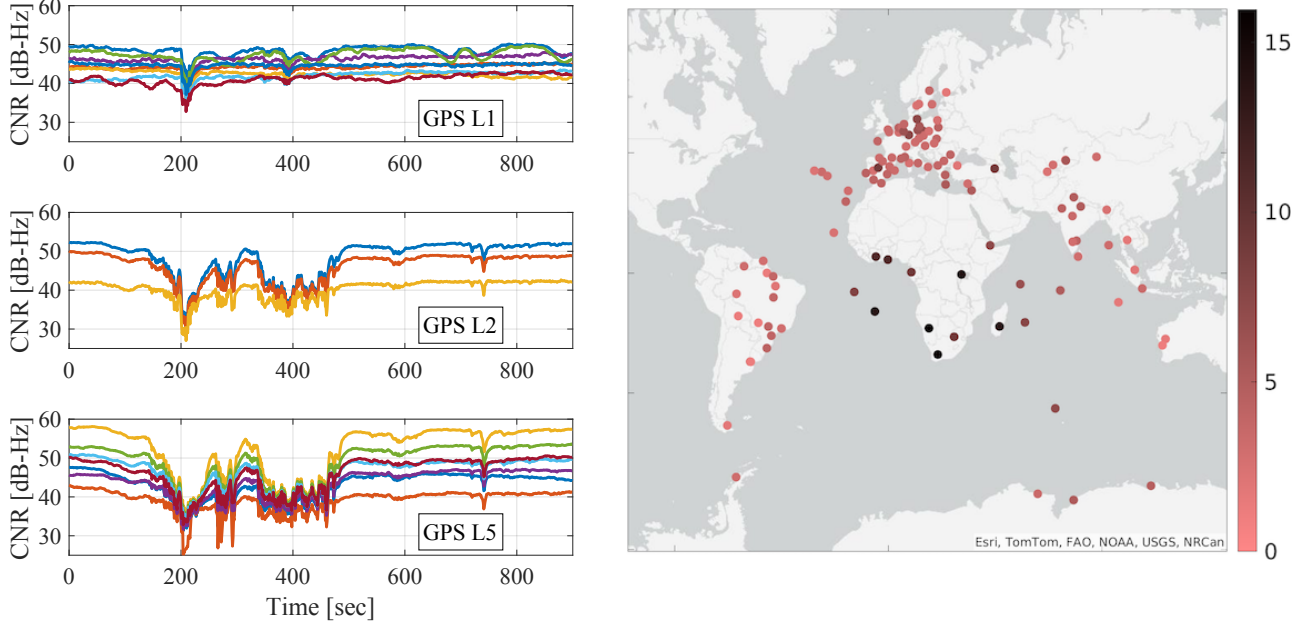


Figure 10: Left: CNR time history of tracked GPS L1, L2, and L5 signals produced by IGS station SUTM in South Africa during the solar radio burst on November 11, 2025. Right: Heat map of maximum CNR degradation at GPS L2 during the solar radio burst.

4 ELEVATION-MASK-BASED INTERFERENCE SOURCE IDENTIFICATION

A simple technique for winnowing the list of possible candidate satellites causing interference is to determine which satellites were overhead the affected region during a given interference event. Two-Line Elements (TLEs) for publicly tracked orbiting objects can be obtained from space-track.org, a database maintained by the United States Space Force to support spaceflight safety. TLEs can be used to compute the approximate position of an object at any time reasonably close to the TLE epoch. While the accuracy of these estimates is roughly 1 km at epoch and degrades with the age of the TLE (Komodromos et al., 2025; Morgan et al., 2025), their precision is more than sufficient for a preliminary candidacy test.

Let \mathcal{S} be the set of all space objects with public TLEs, and $\mathcal{I}_d \subseteq \mathcal{I}$ be the set of all reference stations that detected a given interference event. The elevation angle of satellite $s \in \mathcal{S}$ as seen from the GNSS antenna at station $i \in \mathcal{I}_d$ is denoted α_{is} . This can be calculated assuming an ellipsoidal model for Earth such as WGS84. Assuming an elevation mask α_0 , satellite $s \in \mathcal{S}$ is considered a valid candidate if $\alpha_{is} \geq \alpha_0$ for all $i \in \mathcal{I}_d$.

Under this framework, one can readily determine the minimum altitude at apogee for candidate satellites. Let $\mathbf{r} \in \mathbb{R}^3$ be an arbitrary position in Earth-centered, Earth-fixed (ECEF) coordinates, and $\alpha_i(\mathbf{r})$ be its elevation angle as seen from station $i \in \mathcal{I}_d$. The minimum-radius position \mathbf{r}^* satisfying all elevation mask constraints may be found by solving

$$\mathbf{r}^* = \underset{\mathbf{r} \in \mathbb{R}^3}{\operatorname{argmin}} \|\mathbf{r}\| \quad \text{such that} \quad \alpha_i(\mathbf{r}) \geq \alpha_0 \quad \text{for all} \quad i \in \mathcal{I}_d \quad (18)$$

A solution, which is unique for $\alpha_0 \geq 0$, can be obtained straightforwardly using a numerical optimizer. The minimum satellite altitude at apogee is then given by $\zeta^* = \|\mathbf{r}^*\| - r_E$, where r_E is Earth's equatorial radius.

One can also construct the feasible region of satellite positions, which corresponds to the intersection of all interior ($\alpha_0 > 0$) or exterior ($\alpha_0 < 0$) points of each reference station's feasibility cone. In the special case of $\alpha_0 = 0$, feasible region construction becomes a nonlinear optimization problem with linear constraints formulated by enforcing $\mathbf{r} \in \mathbb{R}^3$ to be above the tangent plane at each reference station. The general case is similar but with nonlinear constraints—those in (18).

One might assume that the reference stations not detecting interference could be used to further constrain the feasible region. However, this assumption does not hold when the antenna gain pattern of the interference source is unknown. For example, if the interference source is equipped with a narrow-beam antenna, many reference stations may not observe the interference even if the satellite satisfies the elevation mask constraint.

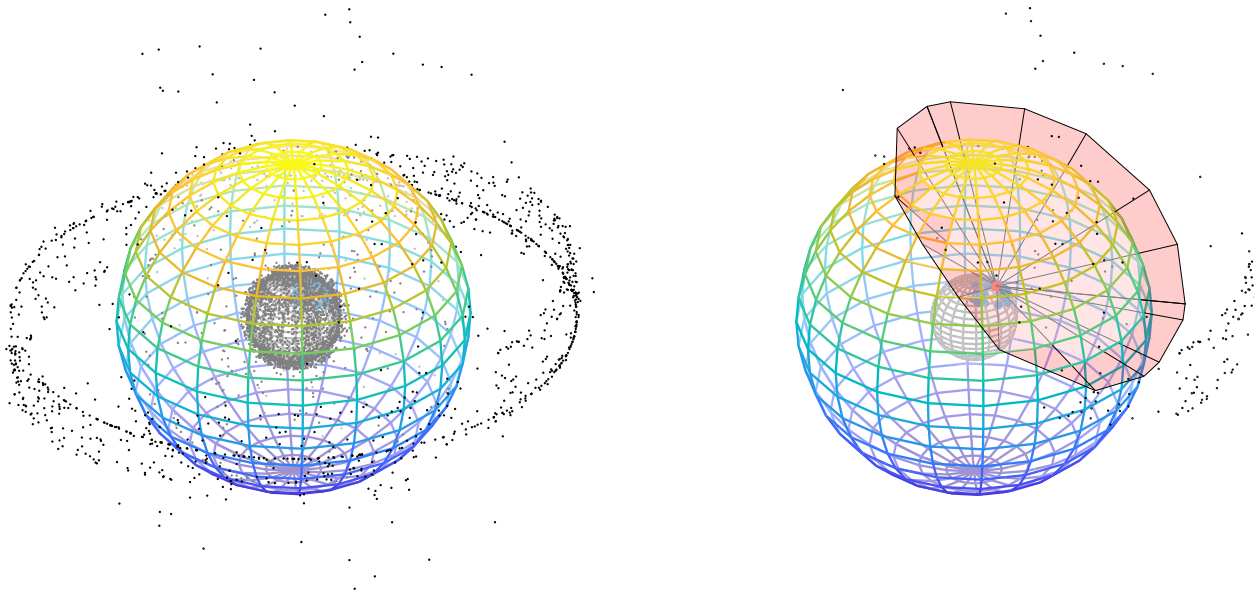


Figure 11: Left: The position of all tracked objects during the high-power interference burst on day 160 year 2021. Right: The position of all satellites that satisfy a 0° elevation mask, excluding debris and rocket bodies. The feasible region in which the interference source could have been positioned is interior to the red surface. For reference, the colored spherical shell corresponds to medium Earth orbit (20,000 km altitude).

Table 2: The minimum altitude and number of satellites that satisfy various elevation masks for the event on day 160 of 2021

α_0	-5°	-3°	-1°	0°	1°	3°	5°	10°
ζ^* [km]	800	957	1,120	1,212	1,290	1,468	1,655	2,173
Number of valid satellites	381	303	229	201	176	145	113	87

Fig. 11 illustrates the positions of all publicly tracked objects during the high-power interference burst on day 160 of 2021. Its right figure shows the subset of satellites satisfying $\alpha_0 = 0$, excluding known debris and rocket bodies. The interior of the quasi-conical red surface represents the feasible region in which the interference satellite could have been located. Table 2 shows ζ^* for various elevation mask angles for this event, as well as the number of satellites that satisfy the elevation mask, e.g., there are 201 satellites that satisfy the 0° elevation mask.

If one assumes that a single satellite is responsible for the observed interference events, the list of candidate satellites can be further narrowed by eliminating those that do not satisfy the elevation mask across multiple events. As shown later, however, the single-satellite assumption does not hold for this paper’s transient interference.

5 MULTI-HYPOTHESIS CNR-BASED INTERFERENCE SOURCE IDENTIFICATION

This section develops a more rigorous CNR-based technique to associate measurements to a satellite than the simple elevation-mask-based method of the previous section.

5.1 Generalized Likelihood Ratio Test

A binary Neyman-Pearson hypothesis testing framework provides the optimal test for simple cases in which the relevant probability density functions under both the null (H_0) and alternative (H_1) hypotheses are fully specified. The test is optimal in the sense of maximizing the probability of detection for a fixed probability of false alarm. In real-world scenarios, however, some parameters are often unknown, resulting in composite tests. If the unknown parameters’ *a priori* probability distributions are also unknown, then no known optimal test exists (Van Trees, 2001). Such conditions apply to this paper’s attempt to identify the source of space-based interference using received-power measurements because parameters such as the source’s transmit power are unknown.

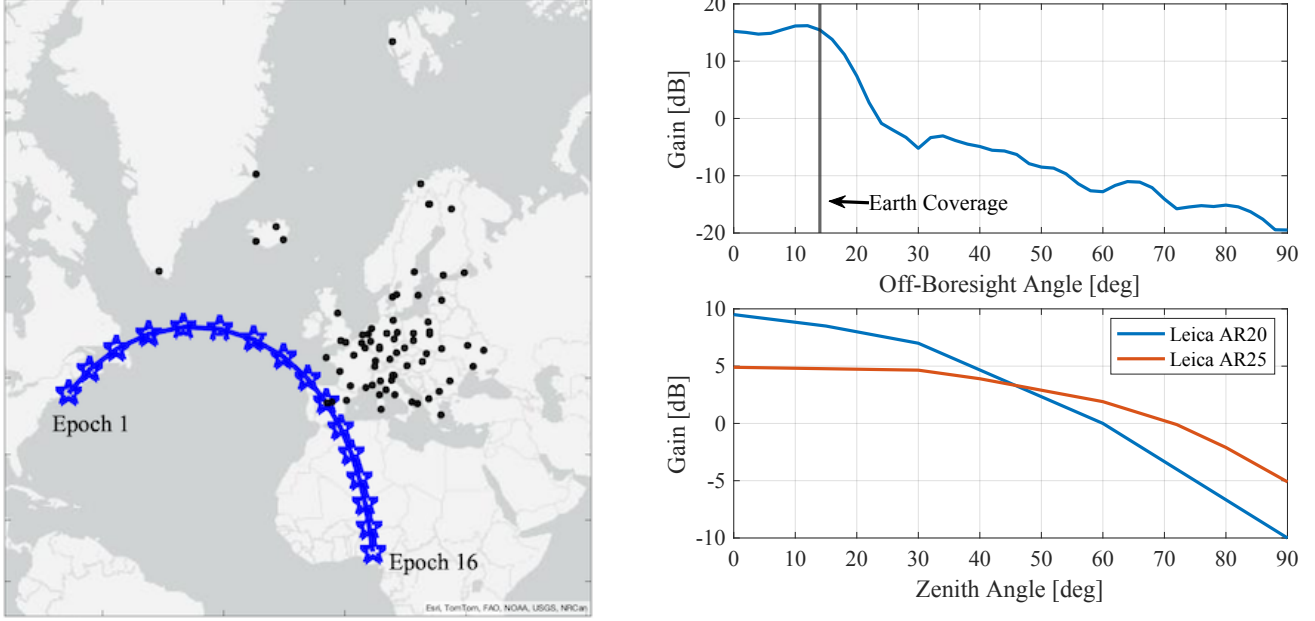


Figure 12: Simulation study setup. Left: The geometry of the true source GNSS satellite over 16 epochs at 15-minute intervals, with terrestrial reference stations shown as black dots. Right: Transmitter (top) and receiver (bottom) gain patterns. The receiver gain patterns are from a model provided by the manufacturer and verified empirically.

The Generalized Likelihood Ratio Test (GLRT) is often employed for composite detection problems with unknown prior distributions. It amounts to a two-step procedure. First, the unknown parameter vector \mathbf{x} is estimated from the measurements \mathbf{z} via maximum likelihood estimation under each hypothesis. Second, these estimates are used to evaluate likelihoods as if they were the correct values of \mathbf{x} under the respective hypotheses (Van Trees, 2001). In the current context, the hypothesis test is M-ary rather than binary, with H_s being the hypothesis that satellite $s \in \mathcal{S}$ is the interference source. The vector \mathbf{x}_s contains parameters assumed unknown *a priori* that are used to model the candidate satellite’s transmit power, antenna gain pattern, and pointing direction (beam vector), or some subset of these.

5.2 Simulation Study

The performance of a multi-hypothesis satellite association technique may be illustrated by a simulation study. This section’s study applies an M-ary version of the GLRT not to the interference source identification problem itself but to a closely related problem: identifying which satellite is the source of a GNSS signal whose CNR measurements are taken from a network of reference stations. This related problem is chosen because (1) it yields insights pertinent to the interference identification problem, and (2) the technique developed and demonstrated in simulation can be tested on real-world CNR measurements from the IGS network to demonstrate identification of a known source GNSS satellite.

For the i th station and j th GNSS signal at epoch k , let $z_{ij}[k]$ be an interference-free CNR measurement, modeled as in (8). The measurement noise variance $\sigma_{ij}^2[k] = f(\theta_R^{ij}[k])$ is modeled as a nonlinear function of the received zenith angle, with $\sigma_{ij} = 0.25$ dB at zenith ($\theta_R^{ij}[k] = 0$) and $\sigma_{ij} = 1.25$ dB at the horizon ($\theta_R^{ij}[k] = 90^\circ$). Both receiver and transmitter antenna gain patterns are modeled as azimuthally symmetric— $G_R^i(\theta, \phi) = G_R^i(\theta)$, $G_T^j(\theta, \phi) = G_T^j(\theta)$ —which is approximately true for both GNSS receiver and transmitter antennas. Fig. 12 shows the assumed reference station locations (black dots) and the position of an exemplar GNSS satellite for 16 time epochs spaced by 15-minute intervals. It also shows the corresponding transmitter and receiver gain patterns. In this study, all transmitter antennas are nadir-pointing, and all reference stations are assumed to utilize either Leica AR20 or AR25 antennas.

At epoch k , let $\mathcal{I}_j[k] \subseteq \mathcal{I}$ be the set of participating stations tracking the j th GNSS signal, and $\mathcal{S}_e[k]$ be the set of candidate source satellites that have all passed elevation mask vetting. Let the measurement vector $\mathbf{z}_j[k]$ contain ordered CNR measurements of signal $j \in \mathcal{J}$ from those receivers tracking the signal at epoch k , i.e., measurements from the set $\{z_{ij}[k] \mid i \in \mathcal{I}_j[k]\}$.

A GLRT-based association test is used to assess whether a candidate satellite $s \in \mathcal{S}_e[k]$ is consistent with $\mathbf{z}_j[k]$ given varying

degrees of prior knowledge about the satellite's transmitter. The unknown parameter vector $\mathbf{x}_s[k]$ may include any combination of transmit power $P_s[k]$, beamwidth $\beta_s[k]$, and beam vector $\mathbf{v}_s[k] = [\theta_s[k], \phi_s[k]]^\top$, where $\theta_s[k]$ and $\phi_s[k]$ are the off-boresight and azimuth angle relative to the nadir vector and a reference azimuth vector in the satellite frame. In the fully unknown case, \mathbf{x}_s includes all three parameters:

$$\mathbf{x}_s[k] = [P_s[k], \beta_s[k], \mathbf{v}_s^\top[k]]^\top \quad (19)$$

Under hypothesis H_s , satellite $s \in \mathcal{S}_e[k]$ is the source of signal $j \in \mathcal{J}$, in which case the measurement vector $\mathbf{z}_j[k]$ is distributed as

$$H_s : \quad \mathbf{z}_j[k] \sim \mathcal{N}(\bar{\mathbf{z}}_j[k, \mathbf{x}_s], R_j[k, \mathbf{x}_s]), \quad s \in \mathcal{S}_e[k] \quad (20)$$

Here, $\bar{\mathbf{z}}_j[k, \mathbf{x}_s]$ is the vector of expected CNR measurements for signal $j \in \mathcal{J}$ emanating from satellite $s \in \mathcal{S}_e[k]$, and $R_j[k, \mathbf{x}_s]$ is a square diagonal matrix with diagonal elements $\sigma_{ij}^2[k]$ for $i \in \mathcal{I}_j[k]$, where the angle $\theta_R^{ij}[k]$, on which $\sigma_{ij(s)}^2[k]$ depends, is calculated assuming signal j comes from satellite s .

This association problem amounts to a multi-hypothesis composite general Gaussian problem (Van Trees, 2001). Hypotheses may be compared based on their negative log likelihood function evaluated at the maximum-likelihood value for the parameter vector \mathbf{x}_s , which, for the model in (20), leads to a cost function of the form

$$J_j[k, \mathbf{x}_s^*] = (\mathbf{z}_j[k] - \bar{\mathbf{z}}_j[k, \mathbf{x}_s^*])^\top (R_j[k, \mathbf{x}_s^*])^{-1} (\mathbf{z}_j[k] - \bar{\mathbf{z}}_j[k, \mathbf{x}_s^*]), \quad s \in \mathcal{S}_e[k] \quad (21)$$

where \mathbf{x}_s^* is the value of \mathbf{x}_s that minimizes $J_j[k, \mathbf{x}_s]$. Thus, for $a, s \in \mathcal{S}_e[k]$, if $J_j[k, \mathbf{x}_a^*] < J_j[k, \mathbf{x}_s^*]$, then H_a is more likely than H_s . Since $\mathcal{S}_e[k]$ is assumed to include the true source for signal j , then for the case of hypotheses with equally likely priors, the probability that H_s is true is given by (Bar-Shalom et al., 2001)

$$P(s|\mathbf{z}_j[k]) = \frac{|R_j[k, \mathbf{x}_s^*]|^{-1/2} \exp(-\frac{1}{2}J_j[k, \mathbf{x}_s^*])}{\sum_{a \in \mathcal{S}_e[k]} |R_j[k, \mathbf{x}_a^*]|^{-1/2} \exp(-\frac{1}{2}J_j[k, \mathbf{x}_a^*])}, \quad s \in \mathcal{S}_e[k] \quad (22)$$

Based on this expression, one could choose the highest-probability satellite, or the top N most probable, or all those for which $P(s|\mathbf{z}_j[k])$ exceeds a threshold. Nearly equivalent to the latter approach, one may retain all $s \in \mathcal{S}_e[k]$ for which $J_j[k, \mathbf{x}_s^*] \leq \nu_j[k]$ for a given threshold $\nu_j[k]$. A principled choice for $\nu_j[k]$ can be obtained by limiting the probability $P(J_j[k, \mathbf{x}_u^*] > \nu_j[k])$ to a specified size, where $u \in \mathcal{S}_e$ is the true source satellite. It is straightforward to show that if \mathbf{x}_u^* is an accurate estimate then $J_j[k, \mathbf{x}_u^*]$ is distributed as a chi-square random variable with $|\mathcal{I}_j[k]|$ degrees of freedom:

$$J_j[k, \mathbf{x}_u^*] \sim \chi_{|\mathcal{I}_j[k]|}^2 \quad (23)$$

From this one can readily calculate $\nu_j[k]$ to ensure $P(J_j[k, \mathbf{x}_u^*] > \nu_j[k])$ is below a chosen value.

5.2.1 Simulation

1000 Monte Carlo trials were conducted at each time epoch depicted in Fig. 12. A trajectory for GPS PRN 15 from February 2020 was chosen for the true source satellite s . At each epoch, (1) a subset $\mathcal{S}_e[k]$ of candidate satellites was generated by applying an elevation mask $\alpha_0 = 0$; (2) the subset $\mathcal{I}_j[k]$ was determined for signal j ; (3) the noise-free CNR $_{ij}$ at each reference station $i \in \mathcal{I}_j[k]$ was computed based on the true parameter vector \mathbf{x}_s under the model (1); (4) a realization of the measurement $\mathbf{z}_{ij}[k]$ was generated according to (8) assuming $\sigma_{ij}^2[k] = f(\theta_R^{ij}[k])$ for each $i \in \mathcal{I}_j[k]$.

Costs $J_j[k, \mathbf{x}_s^*]$ were calculated for all $s \in \mathcal{S}_e[k]$, with \mathbf{x}_s estimated under five scenarios reflecting varying levels of *a priori* knowledge of the transmitter: (S1) all parameters known; (S2) unknown transmit power P_s ; (S3) unknown P_s and beam width β_s ; (S4) unknown P_s and beam vector \mathbf{v}_s ; and (S5) unknown P_s, β_s , and \mathbf{v}_s . The maximum likelihood (minimum cost) estimates β_s^* , and \mathbf{v}_s^* were found via grid search, while P_s^* was taken as the mean of the elements of the difference $\mathbf{z}_j[k] - \bar{\mathbf{z}}_j[k, \tilde{\mathbf{x}}_s^*]$, with $\tilde{\mathbf{x}}_s^* = [P_s = 0, \beta_s^*, \mathbf{v}_s^{*\top}]^\top$.

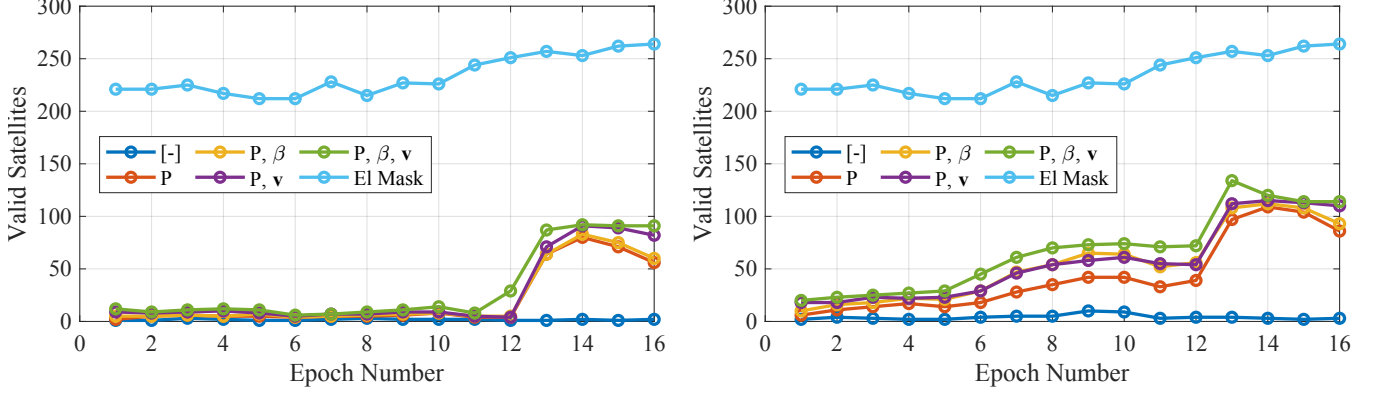


Figure 13: The number of satellites that, for at least 5% of the Monte Carlo trials, remained viable candidates after association testing with $P(J_j[k, \mathbf{x}_s^*] > \nu_j[k]) < 10^{-3}$. The legends indicate which parameters are assumed unknown. The number of satellites satisfying all elevation masks at each epoch, $|\mathcal{S}_e[k]|$, is shown in light blue. Left: All reference stations modeled with Leica AR20 antennas. Right: All reference stations modeled with Leica AR25 antennas.

5.2.2 Results and Discussion

Fig. 13 shows the number of satellites that, for at least 5% of the Monte Carlo trials, were considered viable, i.e., had costs $J_j[k, \mathbf{x}_s^*]$ below $\nu_j[k]$ when $P(J_j[k, \mathbf{x}_u^*] > \nu_j[k])$ was limited to under 10^{-3} . As shown, the analysis was performed across various epochs, transmitter scenarios, and receiver antenna types. The results reveal that the GLRT-based multi-hypothesis technique is far superior to multi-station elevation masking in narrowing the number of candidate satellites. When transmitter power, beamwidth, and beam vector are known *a priori*, only the true satellite and a small number of others are consistent with the data at any given epoch, as shown by the dark blue trace along the bottom of both panels in Fig. 13. Even when none of these parameters is known *a priori* and so they must be estimated from the network data (green traces in Fig. 13), the method offers useful discrimination. But it is clear that, at the geometries arising during many of the epochs, the CNR-based association method alone would not be sufficient for confident single-epoch satellite identification when there are unknown parameters.

The two panels in Fig. 13 also show that the receiver antenna gain pattern substantially influences the test’s discrimination power. Specifically, the Leica AR20 (left panel) exhibits a more pronounced variation in gain over the zenith angle compared to the Leica AR25 (right panel), thereby enhancing the sensitivity of measurements to changes in zenith angle and facilitating the exclusion of false associations.

The impact of transmitter satellite geometry on correct association is also evident. Note how the number of viable satellites increases when the true satellite approaches the equator during the last few epochs. Fig. 14 offers an explanation: Compared to epoch $k = 1$, which is far from the equator and saw relatively few false associations, at epoch $k = 16$ many more satellites are falsely associated because a whole band of geostationary satellites are also consistent with the data when any transmitter parameters are unknown *a priori*.

5.3 GNSS Satellite Identification with Real CNR Data

The GNSS satellite identification technique developed here and demonstrated by simulation was also tested using real CNR measurements taken from reference stations in \mathcal{I} . Let $\mathcal{J}_g \subset \mathcal{J}$ be the subset of GPS L1 C/A signals. Based on historical data from each site $i \in \mathcal{I}$ and on the known azimuthally symmetric antenna gain pattern $G_T^j(\theta)$ and transmit power P_T^j for $j \in \mathcal{J}_g$ from a select number of GPS satellites $\mathcal{S}_g \subset \mathcal{S}_e$, empirical estimates of the combined receiver antenna gain pattern and receiver noise floor

$$G_{\text{R0}}^i(\theta) = G_{\text{R}}^i(\theta) - N_0^i \quad (24)$$

were obtained for all $i \in \mathcal{I}$ across zenith angles $\theta \in [0, 90^\circ]$, assuming azimuthal symmetry. Accurate estimates of $G_T^j(\theta)$ for $j \in \mathcal{J}_g$ and $s \in \mathcal{S}_g$ are available in public documents (Marquis and Reigh, 2015), and corresponding estimates for P_T^j were obtained by observation with a calibrated high-gain antenna. Normalized estimates of $G_{\text{R0}}^i(\theta)$ for stations MEDI (Leica AR20) and NICO (Leica AR25) are shown in Fig. 15.

The GLRT-based multi-hypothesis association test was then applied based on measured CNR_{ij} values for $i \in \mathcal{I}$ and $j \in \mathcal{J}_g$ for

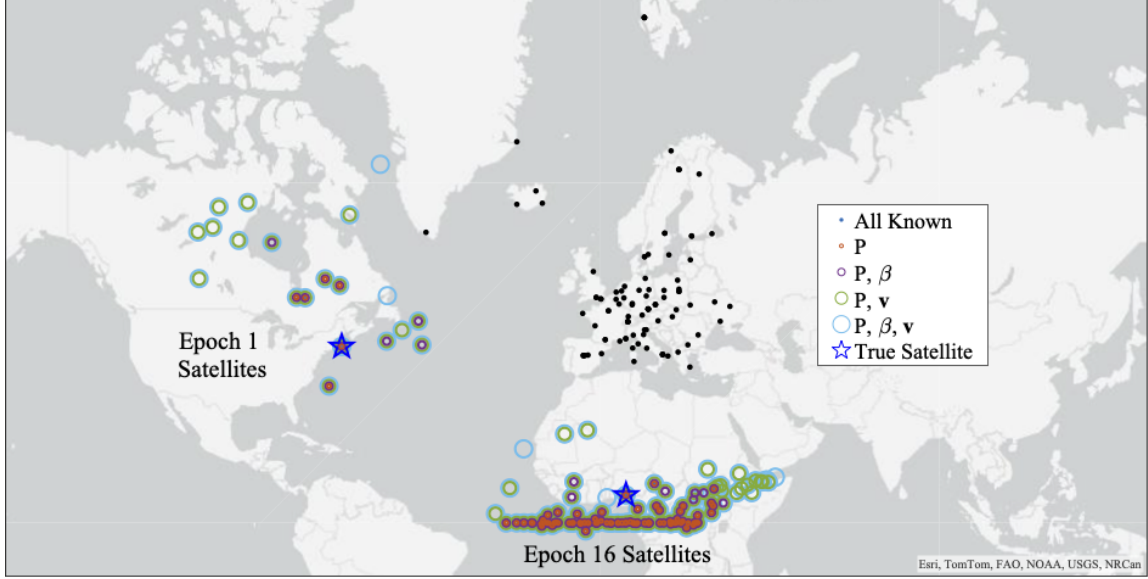


Figure 14: Satellites that remained viable candidates on more than 5% of Monte Carlo trials for epoch $k = 1$ and epoch $k = 16$ when the reference stations are modeled as using the Leica AR25 antenna. The legend gives symbols for the true satellite and for valid candidates under scenarios with an increasing number of unknown transmitter parameters.




Figure 15: Normalized estimates of $G_{\text{R}0}^i(\theta)$ for stations MEDI (Leica AR20, $i = 4$) and NICO (Leica AR25, $i = 5$). The blue dots represent collected values of $\text{CNR}_{ij} - G_{\text{T}}^j - L_{ij}$, while the red line indicates a weighted third-order polynomial fit, which becomes $G_{\text{R}0}^i(\theta)$, all normalized to $G_{\text{R}0}^i(0) = 0$. As expected, measurement noise increases significantly as the elevation angle decreases.

$s \in \mathcal{S}_{\text{g}}$, according to (8) with $\sigma_{ij}^2 = f_i(\theta_{\text{R}}^{ij})$. The site-specific function f_i was determined by the zenith-angle-dependent spread in the empirical data underlying each station's antenna model (e.g., the blue dots in Fig. 15). The method was tested under full *a priori* knowledge of the transmit power, beamwidth, and beam vector of $s \in \mathcal{S}_{\text{g}}$, and under a case where the transmit power P_s was unknown. The results of this empirical study were qualitatively consistent with those of the simulation study: the true $s \in \mathcal{S}_{\text{g}}$ was often uniquely identifiable in a single epoch when its transmitter parameters were all known *a priori*, but false associations increased in geometry-dependent ways when P_s had to be estimated from the data.

5.4 Implications for CINR-Based Interference Source Identification

Based on CNR and CINR measurements alone, interference source identification is more challenging than GNSS satellite identification in two key respects. First, the transmitter gain pattern may be wholly unknown, not just unknown in beamwidth β_s as in the simulation study. Second, and more importantly, the receiver antenna gain pattern $G_{\text{R}}^i(\theta)$ and noise floor N_0^i must be known independently, rather than combined as in (24), because the model for CINR_{ij} given in (4) involves an isolated value of \tilde{N}_0^i . Yet it is not possible to independently estimate $G_{\text{R}}^i(\theta)$ and N_0^i for $i \in \mathcal{I}$ strictly from historical CNR measurements because CNR is a linear combination of the two quantities, as evident in (1) and (2). Separate estimates of N_0^i can be obtained for $\mathcal{I}_{\text{a}} \in \mathcal{I}$, the subset of reference stations for which an accurate antenna gain pattern $G_{\text{R}}^i(\theta)$ is available, e.g., one provided by the manufacturer. But the authors of the current paper found \mathcal{I}_{a} was limited to stations with Leica AR20 or AR25 antennas. For all other stations (all $i \in \mathcal{I} \setminus \mathcal{I}_{\text{a}}$), either $G_{\text{R}}^i(\theta)$ was not available or the available model did not match the shape of the empirically derived combined model $G_{\text{R}0}^i(\theta)$. (Note that, whereas the IGS provides extremely accurate models of antenna

phase center variation as a function of θ and ϕ (Krzan et al., 2020), it provides no such models for antenna gain.) Importantly, the geographic distribution of \mathcal{I}_a is not as extensive as that shown in Fig. 12, leading to weaker observability of the unknown transmitter parameters, and thus weaker discrimination of candidate satellites.

In view of these challenges, one can expect the current performance of CNR-based multi-hypothesis association for interference source identification to fall somewhere between the light blue and green curves shown in Fig. 13, tending towards the light blue curve when the unknown parameters multiply. It becomes clear that, to confidently identify a GNSS interference source, more powerful techniques are required. Nonetheless, CNR-based multi-hypothesis association could offer useful prior probabilities for other interference source identification techniques if accurate gain patterns $G_R^i(\theta)$ were available for more stations in \mathcal{I} . Moreover, many stations provide CNR measurements only as integer-quantized values. Reporting these measurements at a higher resolution would significantly enhance measurement accuracy.

Assuming these desiderata for IGS stations will eventually be fulfilled, the mapping from the GNSS satellite identification problem to the interference source identification problem is briefly outlined. First, supposing that $G_R^i(\theta)$ is known, and leveraging the known transmit gain pattern $G_T^j(\theta)$ and transmit power P_T^j for $j \in \mathcal{J}_g$, N_0^i is estimated from historical CNR data for all $i \in \mathcal{I}$ so that the models (4) and (9) can be applied. Then, assuming that at epoch k there is an interference event, let $\mathcal{I}_d[k] \subset \mathcal{I}$ be the subset of stations that detect the event, and let the vector $\mathbf{z}[k]$ contain ordered station-specific detection statistics at epoch k , i.e., values from the set $\{\Lambda_i[k] \mid i \in \mathcal{I}_d[k]\}$. For every satellite $s \in \mathcal{S}[k]$ and for an assumed parameter vector \mathbf{x}_s , one can develop a hypothesis model for $\mathbf{z}[k]$ like the one in (20) and find the subset of $\mathcal{S}[k]$ consistent with $\mathbf{z}[k]$, just as with the GNSS satellite identification problem. Let the resulting probabilities, analogous to $P(s|\mathbf{z}_j[k])$ in (22), be denoted $P(s|\mathbf{z}[k])$. The following section shows how these probabilities can act as prior probabilities for a more powerful identification technique.

6 TDOA-BASED INTERFERENCE SOURCE IDENTIFICATION

This section presents a satellite association framework based on TDOA measurements, together with a TDOA analysis based on recently captured raw in-phase and quadrature (IQ) samples from two spatially separated receivers in Europe during a continent-scale interference event in February 2026.

6.1 TDOA Measurement Models

Consider two receivers at fixed locations \mathbf{g}_1 and \mathbf{g}_2 in the ECEF reference frame. Let $\mathcal{K} = \{1, \dots, K\}$ be the set of TDOA measurement epochs over an observation interval, \mathcal{S}_e be the set of candidate satellites that satisfy elevation mask vetting for all $k \in \mathcal{K}$, and $\mathbf{r}_s[k]$ be the ECEF position of satellite $s \in \mathcal{S}_e$ at epoch k . For model simplicity, $\mathbf{r}_s[k]$ is assumed to remain constant over the signal's time-of-flight for each measurement. Denoting by $\rho_{is}[k] = \|\mathbf{g}_i - \mathbf{r}_s[k]\|$ the range between receiver $i \in \{1, 2\}$ and satellite $s \in \mathcal{S}_e$, the modeled TDOA for s is the range difference between the two receivers: $\bar{y}_s[k] = \rho_{2s}[k] - \rho_{1s}[k]$. Assume that $u \in \mathcal{S}_e$ is the true interference source. Then the actual TDOA measurement is modeled as

$$y[k] = \bar{y}_u[k] + w[k], \quad k \in \mathcal{K} \quad (25)$$

where the AWGN noise term $w[k] \sim \mathcal{N}(0, \sigma^2)$ accounts for the aggregate effects of thermal noise and quantization error from nearest-neighbor sampling. Epoch-wise measurements are stacked as $\mathbf{y} = [y[1], \dots, y[K]]^\top$, and modeled values for $s \in \mathcal{S}_e$ as $\bar{\mathbf{y}}_s = [\bar{y}_s[1], \dots, \bar{y}_s[K]]^\top$.

The location $\mathbf{r}_s[k]$ of each $s \in \mathcal{S}_e$ is not known exactly but can be approximated with the public TLE data available for s . Let $\tilde{\mathbf{r}}_s[k]$ be the TLE-based orbit for s propagated to epoch k , and \mathbf{e}_s be the error in $\tilde{\mathbf{r}}_s[k]$, with $\mathbf{e}_s \sim \mathcal{N}(\mathbf{0}, Q_s)$. Modeling \mathbf{e}_s as constant is acceptable for observation intervals limited to a few seconds as in the current case. It is assumed that \mathbf{e}_s and $w[k]$ are mutually independent for all $s \in \mathcal{S}_e$ and $k \in \mathcal{K}$. The true satellite location is related to $\tilde{\mathbf{r}}_s[k]$ and \mathbf{e}_s by

$$\mathbf{r}_s[k] = \tilde{\mathbf{r}}_s[k] + \mathbf{e}_s, \quad s \in \mathcal{S}_e, \quad k \in \mathcal{K} \quad (26)$$

The framework established in Komodromos et al. (2025); Morgan et al. (2025) can be leveraged to investigate the effects of \mathbf{e}_s on the TDOA measurement. Let the range Jacobian of the i th receiver with respect to $\mathbf{r}_s[k]$ be given by

$$\Gamma_i[k] = \frac{\partial \rho_{is}[k]}{\partial \mathbf{r}_s[k]} = -\frac{(\mathbf{g}_i - \mathbf{r}_s[k])^\top}{\|\mathbf{g}_i - \mathbf{r}_s[k]\|}, \quad i \in \{1, 2\}; \quad k \in \mathcal{K} \quad (27)$$

This is approximated as constant over the short observation interval: $\Gamma_i[k] = \Gamma_i[1]$, $i \in \{1, 2\}$. The modeled TDOA for $s \in \mathcal{S}_e$ can be approximated by retaining up to the linear terms in a Taylor series expansion of $\bar{y}_s[k]$ about $\tilde{\mathbf{r}}_s[k]$:

$$\begin{aligned}\bar{y}_s[k] &= \rho_{2s}[k] - \rho_{1s}[k] \\ &= \|(\tilde{\mathbf{r}}_s[k] + \mathbf{e}_s) - \mathbf{g}_2\| - \|(\tilde{\mathbf{r}}_s[k] + \mathbf{e}_s) - \mathbf{g}_1\| \\ &= \|\tilde{\mathbf{r}}_s[k] - \mathbf{g}_2\| - \|\tilde{\mathbf{r}}_s[k] - \mathbf{g}_1\| + (\Gamma_2 - \Gamma_1) \mathbf{e}_s \\ &= \tilde{y}_s[k] + \eta_s\end{aligned}\tag{28}$$

One thus observes that the random ephemeris error \mathbf{e}_s manifests as a random TDOA measurement bias $\eta_s = (\Gamma_2 - \Gamma_1)\mathbf{e}_s$ over all $k \in \mathcal{K}$, with $\eta_s \sim \mathcal{N}(0, \sigma_s^2)$ and $\sigma_s^2 = (\Gamma_2 - \Gamma_1)Q_s(\Gamma_2 - \Gamma_1)^\top$. One may also lump into η_s any residual time synchronization error between the reference stations, which, like \mathbf{e}_s can be considered constant over \mathcal{K} . Finally, let $\tilde{\mathbf{y}}_s = [\tilde{y}_s[1], \dots, \tilde{y}_s[K]]^\top$ be the TLE-based modeled TDOA vector for $s \in \mathcal{S}_e$.

6.2 TDOA-Based Association Framework

Based on the foregoing measurement models, an association framework can be developed with TLE-derived satellite positions. The objective is to assess whether a candidate satellite $s \in \mathcal{S}_e$ is consistent with the measurement \mathbf{y} . For each s , the elements of $\tilde{\mathbf{y}}_s$ can be readily calculated as in (28). Let $\mathbf{1}$ be the all-ones vector of appropriate size, and let $\mathbf{w} = [w[1], \dots, w[K]]^\top \sim \mathcal{N}(\mathbf{0}, R)$, with $R = \sigma^2 I$. Then the TDOA measurement residual with respect to the TLE-modeled satellite s is

$$\gamma_s = \mathbf{y} - \tilde{\mathbf{y}}_s = \bar{\mathbf{y}}_u - \tilde{\mathbf{y}}_s + \eta_s \mathbf{1} + \mathbf{w}\tag{29}$$

Note that if $s = u$, i.e., if s is the true satellite, then γ_s is zero mean. A cost can be assigned to each satellite $s \in \mathcal{S}_e$ under a Bayesian framework that treats η_s as an unknown random parameter to be estimated. This is similar to how unknown transmitter parameters are treated in Section 5 except that η_s is modeled as a random variable with a known prior distribution: $\eta_s \sim \mathcal{N}(0, \sigma_s^2)$. The cost function follows from the negative log of the *a posteriori* probability $p(\eta_s | \mathbf{y})$:

$$J_s(\eta_s) = (\gamma_s - \eta_s \mathbf{1})^\top R^{-1} (\gamma_s - \eta_s \mathbf{1}) + \left(\frac{\eta_s}{\sigma_s}\right)^2\tag{30}$$

For satellite s , the maximum *a posteriori* estimate of η_s is the value that minimizes the cost: $\eta_s^* = \operatorname{argmin}_{\eta_s} J_s(\eta_s)$. Under the true hypothesis H_u , the minimum cost is distributed as

$$H_u : J_u(\eta_u^*) \sim \chi_K^2\tag{31}$$

This distribution arises because the total number of independent measurements is $K + 1$ (K from \mathbf{y} and one from the prior constraint on η_u), but one degree of freedom is lost to estimate η_u (Bar-Shalom et al., 2001).

Thus, as in Section 5.5.4, the satellite identification problem reduces to a multi-hypothesis association problem where comparisons between hypotheses are made on the basis of a cost function. In fact, the two techniques can be combined by taking the probability $P(s|\mathbf{z})$ from Section 5.5.4 as the prior probability for the TDOA-based identification problem. Then, since \mathcal{S}_e is assumed to include the true interference satellite, and approximating σ_s as equal for all $s \in \mathcal{S}_e$, the probability that H_s is true is

$$P(s|\mathbf{y}, \mathbf{z}) = \frac{\exp\left(-\frac{1}{2}J_s(\eta_s^*)\right) P(s|\mathbf{z})}{\sum_{a \in \mathcal{S}_e} \exp\left(-\frac{1}{2}J_a(\eta_a^*)\right) P(a|\mathbf{z})}, \quad s \in \mathcal{S}_e\tag{32}$$

Here, the vector \mathbf{z} contains the ordered station-specific detection statistics for those stations \mathcal{I}_d that detect the interference event associated with the TDOA measurement, i.e., values from the set $\{\Lambda_i \mid i \in \mathcal{I}_d\}$. As before, one can choose the highest-probability satellite, or the top N most probable, or all those for which $P(s|\mathbf{y}, \mathbf{z})$ exceeds a threshold.

6.3 A Wideband Raw-Signal Capture

On February 11, 2026, several wide-area transient GNSS interference events occurred across Europe. Fig. 16 shows the CNR of the tracked GPS L1 C/A and BeiDou B1I signals at the METG station in Finland for an interval on this day during which two




Figure 16: CNR time history of tracked GPS L1 C/A (left) and BeiDou B1I (right) signals at the IGS station METG in Finland. The GPS L1 C/A signals are first affected by each interference event, followed by the BeiDou B1I signals. The black line indicates the 2.3-second time-overlapped raw IQ capture interval.

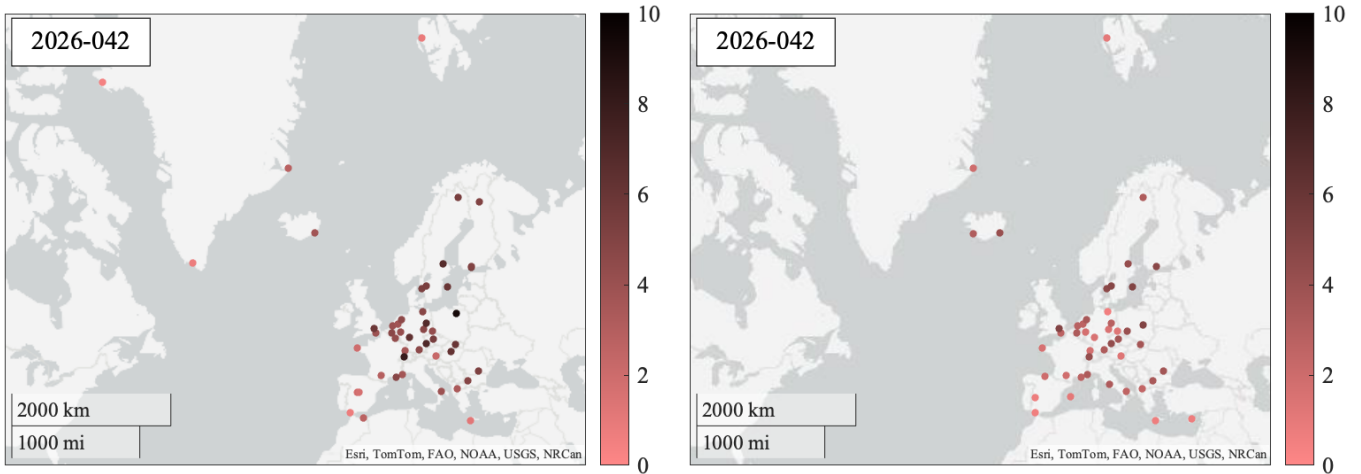


Figure 17: Interference heat maps for day 42 of year 2026 from tracked GPS L1 C/A (left) and BeiDou B1I (right) signals.

events occurred. For each event, the GPS L1 C/A signals first dropped by 5 dB, then recovered, immediately following which the BeiDou B1I signals experienced two cycles of 5 dB degradation and recovery. The CNR drops of the BeiDou B1I signals persisted for ten seconds—more than twice as long as the drops observed on the GPS L1 C/A signals. The interference heat maps shown in Fig. 17 reveal a spatial pattern consistent with previously observed interference events.

The remainder of this section analyzes contemporaneous raw IQ samples captured by a receiver located in Amsterdam, Netherlands (R1), and one located in Trondheim, Norway (R2). R1 employed a complex sampling rate of 60 MHz, while R2 sampled at 75 MHz. Both captures were centered at 1585 MHz, utilized three-bit quantization, and were driven by high-stability oven-controlled crystal oscillators (OCXOs). These receivers are part of the Advanced RFI Detection, Analysis, and Alerting System, which focuses on the capture and collection of GNSS interference signals (Sokolova et al., 2022; Morrison et al., 2023). The system was developed under the European Space Agency Navigation, Innovation, and Support Program (NAVISP).

R1 captured the interference signals responsible for the spatial degradation pattern shown in Fig. 17 and the power spectrum shown in Fig. 9. The interference waveform that caused CNR drops on the GPS L1 C/A signals was centered at 1577.5 MHz, which is consistent with the interference observed in Fig. 8. This signal exhibited cyclostationary properties, with repeating structure every 12 μ s. The interference waveform that caused the CNR drops on the BeiDou B1I signals was centered at 1558.5 MHz. It too exhibited cyclostationary properties, at first with repeating structure every 255.8 μ s, then later at 292.9 μ s.

On February 11, 2026 at 05:36:30 UTC, R1 and R2 recorded 2.3 seconds of time-overlapped raw IQ samples during the 1558.5-MHz-centered interference event. Several GPS and Galileo signals were contained within the captured bandwidth, which permitted processing and GPS time registration to better than 30 ns using the GRID software-defined GNSS receiver (Nichols et al., 2022; Pany et al., 2024).




Figure 18: Left: Example normalized cross correlation at the optimal FDOA. Right: TDOA measurement time history.

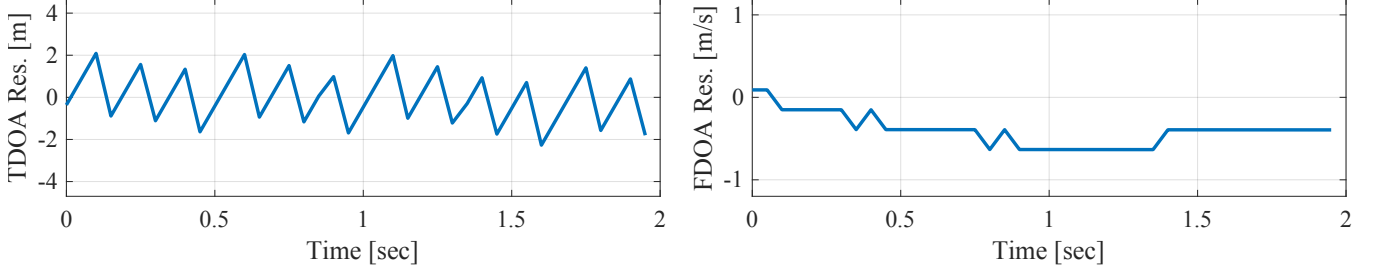


Figure 19: TDOA (left) and FDOA (right) residuals for Cosmos 2546 (NORAD ID 45608). The effects of nearest-sample quantization and Doppler search quantization are evident in the respective plots.

6.4 TDOA Measurement and Method Validation

The IQ samples from the two receivers were synchronized in both time and frequency and the samples from R1 were upsampled to 75-MHz to match those of R2. To obtain the maximum likelihood estimates of TDOA and frequency-difference of arrival (FDOA), the complex ambiguity function (CAF) was computed (Stein, 1981). The TDOA and FDOA values that maximize the CAF magnitude correspond to the maximum likelihood estimates. The left panel of Fig. 18 shows the magnitude of the normalized cross correlation C at the optimal FDOA estimate (155 Hz) for a 50 ms integration interval. Similar functions were repeatedly calculated every 50 ms over two seconds to obtain the 40 TDOA measurements shown as a time history in the right panel of Fig. 18. The TDOA measurements are magnitude-bounded by 4.25 ms, which corresponds to the delay of a hypothetical light pulse traveling between the two receivers. The first 0.3 seconds of the time-overlapped IQ samples were excluded because near-equal-magnitude cyclostationary peaks every 255.8 μ s caused ambiguity in the CAF. Apparently, the interference signal was purely periodic during this interval. After 0.3 seconds, significant repetition remained present, but the CAF peak became unambiguous, as shown in the left panel of Fig. 18.

To validate the TDOA measurement technique, TDOAs were measured for each available GPS L1 C/A signal captured in the wideband data and compared against the expected TDOA for the corresponding overhead GPS satellite. Expected TDOA values were modeled using the true satellite positions as provided by IGS final ephemerides. Besides differential time-of-flight, the model also accounted for tropospheric delay via the Saastamoinen model, ionospheric delay via the Klobuchar model, and ECEF frame rotation during signal time of flight. The resulting measured and modeled TDOAs exhibited excellent agreement, remaining within 15 ns of each other during the 2.3-second overlapping interval of the wideband captures.

6.5 Satellite Association

The measurement vector \mathbf{y} was formed from the time history shown in Fig. 18. Values for the TDOA measurement noise variance σ^2 and for the TLE-based position error variance Q_s were chosen as follows for all $s \in \mathcal{S}_c$. Because the interference signal has a high interference-to-noise ratio and a bandwidth spanning several MHz, the TDOA estimation errors are dominated by nearest-sample quantization error, the variance of which, denoted σ_n^2 , can be determined by calculating the variance of a uniform distribution with a width of one sample, or $\sigma_n^2 = 0.76^2 \text{ m}^2$ for 75-MHz sampling. Alternatively, a conservative choice of σ^2 can be determined by selecting the width of half a sample, $\sigma^2 = 2^2 \text{ m}^2$ for 75-MHz sampling, which is an over-bound when compared to σ_n^2 and more amenable to the developed Gaussian framework. Thus, $\sigma^2 = 2^2 \text{ m}^2$.

Position estimates derived from propagated TLEs typically exhibit kilometer-level inaccuracies, which can escalate to 10 km

under volatile orbital conditions (Vetter, 2007). Error is typically concentrated in the along-track direction and increases with the age of ephemeris. A conservative approach was taken to over-bound the ephemeris position error. For all $s \in \mathcal{S}_e$, the position covariance in ECEF was set to $Q_s = 10000^2 I$. The TLEs of satellites identified as either “payload” or “unknown” within seven days of the interference event were retrieved from space-track.org. For each satellite, the TLE with the epoch nearest to the measurement time was used to calculate position. The set \mathcal{S}_e was composed of all satellites satisfying elevation angle masking with $\alpha_0 = 0$ for all stations that detected the event.

For each satellite $s \in \mathcal{S}_e$, the cost $J_s(\eta_s^*)$ was computed. For a threshold ν chosen such that $P(J_u(\eta_u^*) > \nu) \leq 10^{-3}$, only one satellite satisfied $J_s(\eta_s^*) \leq \nu$: Cosmos 2546 (NORAD ID 45608). This is one of six Molniya-orbit satellites that compose the Russian Edinaya Kosmicheskaya Sistema (EKS), an early warning constellation. For Cosmos 2546, even assuming equally probable priors $P(s|z) = 1/|\mathcal{S}_e|$ for $s \in \mathcal{S}_e$, the probability $P(s|\mathbf{y}, \mathbf{z})$ calculated according to (32) for Cosmos 2546 was numerically indistinguishable from unity. Cosmos 2546 also satisfied an $\alpha_0 = 35^\circ$ elevation mask with respect to every $i \in \mathcal{I}_d$. This is consistent with its highly elliptical orbit, which positioned it far above the northern Atlantic Ocean during the event captured by the wideband data.

T/FDOA residuals for Cosmos 2546 are shown in Fig. 19. The estimated value of the TLE error was $\eta_s^* = 188$ m and the standard deviation of the TDOA residuals was 1.2 m, indicating a highly plausible match. The projected TDOA standard deviation due to TLE error was $\sigma_s = 370$ m for this satellite given the assumed Q_s . The FDOA residuals shown in the right panel of Fig. 19 also indicate tight consistency. FDOA measurements were not used for association but could easily be incorporated into an augmented association framework.

6.6 Further Analysis and Discussion

When all space objects are considered, not only those in \mathcal{S}_e , one additional satellite satisfies $J_s(\eta_s^*) < \nu$ —one from SpaceX’s Starlink constellation. But this satellite is not a viable candidate: during the interval of the TDOA measurement it was located over the Pacific Ocean west of South America, yielding elevation angles far below zero for all stations detecting the event. One may conclude from these false associations that fusing TDOA-based association with elevation angle considerations or CNR-based priors is necessary for a confident unique association when data are limited to a single short TDOA time history.

The TLE for Cosmos 2546 closest to the TDOA measurement time is 40 hours offset, which motivates further examination of its expected accuracy. An analysis was performed based on the full historical archive of TLEs for Cosmos 2546 since its launch in 2020. Let $\mathcal{L} = \{1, \dots, L\}$ be the index set for this TLE archive, and let t_l represent the epoch for the l th TLE. Let $\mathbf{r}_k(t_l)$ denote the position of Cosmos 2546 based on the k th TLE but propagated to the epoch of the l th TLE. For all $l \in \mathcal{L}$ and all $k \in \mathcal{L} \setminus l$, one may assume that $\mathbf{r}_l(t_l)$ is more accurate than $\mathbf{r}_k(t_l)$. Let $\mathcal{L}_l \subset \mathcal{L}$ be the set of TLEs within a 60-hour window of t_l , and let $\mathcal{R}_l = \{\|\mathbf{r}_k(t_l) - \mathbf{r}_l(t_l)\| \mid k \in \mathcal{L}_l\}$ be the corresponding residual set. Statistical analysis of $\mathcal{R} = \bigcup_{l \in \mathcal{L}} \mathcal{R}_l$ (after outlier removal) revealed that the maximum per-dimension residual standard deviation was below 4 km and occurred in the along-track direction. This value represents an appropriate upper bound for Cosmos 2546’s *a priori* position error and falls below 10 km, the determinant of Q_s assumed in the association framework. This methodology for gauging orbital accuracy mirrors established techniques for verifying the accuracy of the IGS’s real-time service products (Hadas and Bosy, 2015; Griffiths and Ray, 2009).

Strictly speaking, identification of Cosmos 2546 as a source for the transient wide-area interference holds only for the 1558.5 MHz interference during the time of wideband IQ capture. But due to the time proximity with the 1577.5 MHz interference, and to their similar time-domain and spatial patterns, interference in the two bands around the time of the IQ capture almost surely originated from the same satellite.

Two other assumptions should be acknowledged as underlying this paper’s results: (1) the TLEs for all interference sources are available on space-track.org, and (2) neither Cosmos 2546 nor any satellite near the quasi-hyperboloid surface defined by the TDOA measurement time history in Fig. 18 maneuvered so significantly between its nearest TLE epoch and the TDOA measurement as to alter the association findings. Both assumptions seem reasonable. In fact, it was found that the TLE error standard deviation in Q_s could be expanded from 10 km to 500 km before any additional satellites became consistent with the TDOA data.

Note that Cosmos 2546 was launched in May 2020 and so cannot be responsible for the interference events that occurred in 2019. Moreover, Cosmos 2546 was not over Europe during some interference events after May 2020. But during all events on the 75 days shown in Table 1 there was at least one EKS satellite above a 35° elevation angle with respect to every reference station that observed the interference. Thus, it is highly probable that the EKS constellation is collectively responsible for the wide-area transient GNSS interference events noted since 2019.

7 CONCLUSION

This paper presented a comprehensive analysis of a novel GNSS interference phenomenon: wide-area transient interference from a space-based source causing up to 10-dB GNSS degradation across Europe since 2019 in the important L1 band. The interference's spatial, temporal, and spectral properties were detailed. A framework was developed to detect events using 1-Hz carrier-to-noise ratio observables from a network of 165 reference stations. A total of 75 days were identified since 2019 on which at least one wide-area interference event occurred. Three techniques were developed to identify candidate source satellites: (1) a simple technique that finds all satellites satisfying a given elevation mask (e.g., 0°) for each station detecting the interference; (2) a more advanced technique based on each station's detection statistic rather than on a binary decision, and on estimation of unknown parameters such as the interference transmit power; and (3) a technique based on time-difference-of-arrival measurements made possible by capture of wideband (e.g., 60-MHz) raw samples during an interference event. By a combination of these techniques the satellite Cosmos 2546 (NORAD ID 45608) was identified with high confidence as one source of the interference. Further analysis pointed to the Russian Edinaya Kosmicheskaya Sistema, an early warning constellation to which Cosmos 2546 belongs, as collectively responsible for the wide-area transient interference causing GNSS degradation across Europe since 2019.

ACKNOWLEDGMENTS

Research was supported by the U.S. Department of Transportation under Grant 69A3552348327 for the CARMEN+ University Transportation Center, and by affiliates of the 6G@UT center within the Wireless Networking and Communications Group at The University of Texas at Austin. Special thanks to Aiden Morrison from SINTEF Digital for providing dual-station wideband-sampled data, and to Michael Meurer and Steffen Thölert from the German Aerospace Center for calibrated estimates of GPS transmitter power.

REFERENCES

- Abraha, K. E., Frisk, A., and Wiklund, P. (2024). GNSS interference monitoring and detection based on the Swedish CORS network SWEPOS. *Journal of Geodetic Science*, 14(1):20220157.
- Bar-Shalom, Y., Li, X. R., and Kirubarajan, T. (2001). *Estimation with Applications to Tracking and Navigation*. John Wiley and Sons, New York.
- Bhatti, J. A., Humphreys, T. E., and Ledvina, B. M. (2012). Development and demonstration of a TDOA-based GNSS interference signal localization system. In *Proceedings of the IEEE/ION PLANS Meeting*, pages 1209–1220.
- Borio, D., O'Driscoll, C., and Fortuny, J. (2012). GNSS jammers: Effects and countermeasures. In *2012 6th ESA Workshop on Satellite Navigation Technologies (Navitec 2012) & European Workshop on GNSS Signals and Signal Processing*, pages 1–7. IEEE.
- Cerruti, A. P., Kintner, P. M., Gary, D. E., Lanzerotti, L. J., de Paula, E. R., and Vo, H. B. (2006). Observed solar radio burst effects on GPS/Wide Area Augmentation System carrier-to-noise ratio. *Space Weather*, 4(S10006).
- Cerruti, A. P., Kintner, P. M., Gary, D. E., Mannucci, A. J., Meyer, R. F., Doherty, P., and Coster, A. J. (2008). Effect of intense December 2006 solar radio bursts on GPS receivers. *Space Weather*, 6(10).
- Clements, Z., Ellis, P., and Humphreys, T. E. (2023). Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit. In *Proceedings of the IEEE/ION PLANS Meeting*, pages 458–469, Monterey, CA.
- Clements, Z., Yoder, J. E., and Humphreys, T. E. (2022). Carrier-phase and IMU based GNSS spoofing detection for ground vehicles. In *Proceedings of the ION International Technical Meeting*, pages 83–95, Long Beach, CA.
- Clements, Z. L., Ellis, P. B., Murrian, M. J., Psiaki, M. L., and Humphreys, T. E. (2026). Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit. *NAVIGATION*, 73(1).
- Clements, Z. L. and Humphreys, T. E. (2025). Transient space-based GNSS interference: Observations and analysis. In *Proceedings of the ION GNSS+ Meeting*, pages 71–81, Baltimore, MD.
- Dow, J. M., Neilan, R. E., and Rizos, C. (2009). The international GNSS service in a changing landscape of global navigation satellite systems. *Journal of geodesy*, 83:191–198.

- Egea-Roca, D., López-Salcedo, J. A., Seco-Granados, G., and Poor, H. V. (2018). Performance bounds for finite moving average tests in transient change detection. *IEEE Transactions on Signal Processing*, 66(6):1594–1606.
- European Space Agency (2025). Lessons from the November 2025 solar storm. https://www.esa.int/Space_Safety/Space_weather/Lessons_from_the_November_2025_solar_storm.
- European Space Agency NAVISP-EL3-014 (2023). CGI GNSS event notification service final presentation. <https://navisp.esa.int/project/details/116/show>.
- GPS Spoofing Workgroup (2024). GPS spoofing: Final report of the GPS spoofing workgroup. Technical report, OPSGROUP.
- Griffiths, J. and Ray, J. R. (2009). On the precision and accuracy of IGS orbits. *Journal of Geodesy*, 83(3):277–287.
- Guépié, B. K., Fillatre, L., and Nikiforov, I. (2012). Sequential detection of transient changes. *Sequential Analysis*, 31(4):528–547.
- Gunawardena, S. and Van Graas, F. (2011). Multi-channel wideband GPS anomalous event monitor. In *Proceedings of the ION GNSS+ Meeting*, pages 1957–1968.
- Gunawardena, S., Zhu, Z., de Haag, M. U., and van Graas, F. (2009). Remote-controlled, continuously operating GPS anomalous event monitor. *NAVIGATION*, 56(2):97–113.
- Hadas, T. and Bosy, J. (2015). IGS RTS precise orbits and clocks verification and quality degradation over time. *GPS solutions*, 19(1):93–105.
- Humphreys, T. E. (2012). Statement on the vulnerability of civil unmanned aerial vehicles and other systems to civil GPS spoofing. *United States House of Representatives Committee on Homeland Security: Subcommittee on Oversight, Investigations, and Management*.
- Humphreys, T. E. (2017). Interference. In *Springer Handbook of Global Navigation Satellite Systems*, pages 469–503. Springer International Publishing.
- Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O’Hanlon, B. W., and Kintner, Jr., P. M. (2008). Assessing the spoofing threat: Development of a portable GPS civilian spoofer. In *Proceedings of the ION GNSS Meeting*, Savannah, GA. Institute of Navigation.
- Johnston, G., Riddell, A., and Hausler, G. (2017). The International GNSS Service. In *Springer Handbook of Global Navigation Satellite Systems*, pages 967–982. Springer International Publishing.
- Komodromos, Z. M., Morgan, S. C., Clements, Z. L., Qin, W., Morrison, W. J., and Humphreys, T. E. (2025). Network-aided pseudorange-based LEO PNT from OneWeb. In *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT.
- Kouba, J. (2009). A guide to using international GNSS service (IGS) products.
- Kriezis, A., Chen, Y.-H., Akos, D., Lo, S., and Walter, T. (2024). GNSS RFI detection and impact characterization in various interference environments using low-cost receivers. In *Proceedings of the ION GNSS+ Meeting*, pages 3348–3360.
- Kriezis, A., Chen, Y.-H., Akos, D., Lo, S., and Walter, T. (2025). GNSS jamming and spoofing monitoring using low-cost COTS receivers. *NAVIGATION*. Submitted for review.
- Krzan, G., Dawidowicz, K., and Wielgosz, P. (2020). Antenna phase center correction differences from robot and chamber calibrations: the case study LEIAR25. *GPS Solutions*, 24(2):44.
- Marcos, E. P., Caizzone, S., Konovaltsev, A., Cuntz, M., Elmarissi, W., Yinusa, K., and Meurer, M. (2018). Interference awareness and characterization for GNSS maritime applications. In *Proceedings of the IEEE/ION PLANS Meeting*, page 908–919. IEEE.
- Marquis, W. A. and Reigh, D. L. (2015). The GPS block IIR and IIR-M broadcast L-band antenna panel: its pattern and performance. *NAVIGATION*, 62(4):329–347.
- Mitch, R., Psiaki, M., and Ertan, T. (2016). Chirp-style GNSS jamming signal tracking and geolocation. *NAVIGATION*, 63(1):15–37.

- Montenbruck, O., Steigenberger, P., Prange, L., Deng, Z., Zhao, Q., Perosanz, F., Romero, I., Noll, C., Stürze, A., Weber, G., et al. (2017). The multi-GNSS experiment (MGEX) of the international GNSS service (IGS)—achievements, prospects and challenges. *Advances in space research*, 59(7):1671–1697.
- Morgan, S. C., Komodromos, Z. M., Qin, W., Clements, Z. L., Graff, A. M., Morrison, W. J., and Humphreys, T. E. (2025). A mock implementation of fused LEO GNSS. In *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT.
- Morrison, A., Sokolova, N., Gerrard, N., Rødningsby, A., Rost, C., and Ruotsalainen, L. (2023). Radio-frequency interference considerations for utility of the Galileo E6 signal based on long-term monitoring by ARFIDAAS. *NAVIGATION*, 70(1).
- Moustakides, G. V. (2014). Multiple optimality properties of the shewhart test. *Sequential Analysis*, 33(3):318–344.
- Murrian, M. J., Narula, L., Iannucci, P. A., Budzien, S., O’Hanlon, B. W., Psiaki, M. L., and Humphreys, T. E. (2021). First results from three years of GNSS interference monitoring from low Earth orbit. *NAVIGATION*, 68(4):673–685.
- Nichols, H. A., Murrian, M. J., and Humphreys, T. E. (2022). Software-defined GNSS is ready for launch. In *Proceedings of the ION GNSS+ Meeting*, Denver, CO.
- Noll, C. E. (2010). The crustal dynamics data information system: A resource to support scientific analysis using space geodesy. *Advances in Space Research*, 45(12):1421–1440.
- Osechas, O., Fohlmeister, F., Dautermann, T., and Felux, M. (2022). Impact of GNSS-band radio interference on operational avionics. *NAVIGATION*, 69(2).
- Pany, T., Akos, D., Arribas, J., Bhuiyan, M. Z. H., Closas, P., Dovis, F., Fernandez-Hernandez, I., Fernández-Prades, C., Gunawardena, S., Humphreys, T., Kassas, Z. M., Salcedo, J. A. L., Nicola, M., Psiaki, M. L., Rügamer, A., Song, Y.-J., and Won, J.-H. (2024). GNSS software defined radio: History, current developments, and standardization efforts. *NAVIGATION*, 71(1).
- Patil, A., Phelts, R. E., Chen, Y.-H., Lo, S., and Walter, T. (2023). Detecting space based interference on GNSS signals. In *Proceedings of the ION GNSS+ Meeting*, pages 1232–1244.
- Patil, A., Phelts, R. E., Walter, T., and Thaelert, S. (2024). Detecting and localizing space based interference on GNSS signals using machine learning. In *Proceedings of the ION International Technical Meeting*, pages 532–545.
- Psiaki, M. L. and Humphreys, T. E. (2016). GNSS spoofing and detection. *Proceedings of the IEEE*, 104(6):1258–1270.
- Psiaki, M. L. and Humphreys, T. E. (2020). *Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications*, volume 1, chapter Civilian GNSS Spoofing, Detection, and Recovery, pages 655–680. Wiley-IEEE.
- RIN Maritime GNSS Interference Working Group (2026). The impacts of GNSS interference on maritime safety report. Technical report, Royal Institute of Navigation.
- Scott, L. (2003). Anti-spoofing and authenticated signal architectures for civil navigation systems. In *Proceedings of the ION GNSS Meeting*, pages 1542–1552.
- Sokolova, N., Morrison, A., and Diez, A. (2022). Characterization of the GNSS RFI threat to DFMC GBAS signal bands. *Sensors*, 22(22).
- Stader, J. and Gunawardena, S. (2021). Leveraging worldwide, publicly-available data to create an automated satnav interference detection system. In *Proceedings of the ION International Technical Meeting*, pages 69–83.
- Stein, S. (1981). Algorithms for ambiguity function processing. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 29(3):588–599.
- Van Trees, H. L. (2001). *Detection, Estimation, and Modulation Theory*. Wiley.
- Vetter, J. R. (2007). Fifty years of orbit determination. *Johns Hopkins APL technical digest*, 27(3):239.
- York, J., Joplin, A., Bratton, M., and Munton, D. (2014). A detailed analysis of GPS live-sky signals without a dish. *Navigation: Journal of The Institute of Navigation*, 61(4):311–322.