Technical Report Documentation Page

1. Report No. USDOT assigns	2. Government Accession No. N/A	3. Recipient's Catalog No. N/A
4. Title and Subtitle Ensuring PNT Resilience		5. Report Date September 17, 2026 6. Performing Organization Code N/A
7. Author(s) Todd E. Humphreys https://orcid.org/0000-0003-0749-6988 Zacharias M. Komodromos Samuel C. Morgan Zachary L. Clements Wenkai Qin W. Jeremy Morrison		8. Performing Organization Report No. N/A
9. Performing Organization Name and Address The University of Texas at Austin Austin, Texas, 78712		10. Work Unit No. (TRAIS) N/A 11. Contract or Grant No. 69A3552348327
12. Sponsoring Agency Name and Address U.S. Department of Transportation (SPR)		13. Type of Report and Period Covered Final (January 6, 2026 to August 30, 2026) 14. Sponsoring Agency Code N/A

15. Supplementary Notes

Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.

16. Abstract

The work reported here involves two primary thrusts. In the first thrust, we identify and characterize dedicated pilot symbols and other predictable elements embedded within the Starlink Ku-band downlink waveform. Exploitation of these predictable elements enables precise opportunistic positioning, navigation, and timing using compact, low-gain receivers by maximizing the signal processing gain available for signal acquisition and time-of-arrival (TOA) estimation. In the second thrust, we analyze and identify 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.

17. Key Words	18. Distribution Statement
Resilient positioning, navigation, and timing. Protection of radio bands designated RNSS. Augmentations to GNSS for PNT.	No restrictions. This document is available to the public through NTIS: National Technical Information Service Springfield, Virginia 22161

19. Security Classif.(of this report)	20. Security Classif.(of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	48	N/A

Form DOT F 1700.7 (8-72) **Reproduction of completed page authorized**

CARMEN+

Center for Automated Vehicles Research with Multimodal Assured Navigation

USDOT University Transportation Centers Program

Final Report: Ensuring PNT Resilience

P.I.	Project Info:
Todd E. Humphreys The University of Texas at Austin Aerospace Engineering	Grant No. 69A3552348327
	DUNS: 832127323
	EIN #: 31-6025986
	Project Effective: January 6, 2026 Project End: August 30, 2026 Submission: September 17, 2026

Consortium Members:

The Ohio State University College of Engineering · The University of Texas at Austin · UCI University of California, Irvine · North Carolina Agricultural and Technical State University

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under grant number 69A3552348327 from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Table of Contents

After its preamble sections 1-10, this report includes two separate papers published during the course of the project.

Paper 1: Qin, Wenkai, Mark L. Psiaki, John R. Bowman, and Todd E. Humphreys. "Pilots and other predictable elements of the Starlink Ku-band downlink." npj Wireless Technology 2, no. 1 (2026): 69.

Paper 2: Clements, Zachary L., Argyris Kriezis, and Todd E. Humphreys. "Chasing Lightning: Detecting, Characterizing, and Identifying a Powerful Space-Based GNSS Interference Source." arXiv preprint arXiv:2606.03673 (2026).

Abstract

The work reported here involves two primary thrusts. In the first thrust, we identify and characterize dedicated pilot symbols and other predictable elements embedded within the Starlink Ku-band downlink waveform. Exploitation of these predictable elements enables precise opportunistic positioning, navigation, and timing using compact, low-gain receivers by maximizing the signal processing gain available for signal acquisition and time-of-arrival (TOA) estimation. In the second thrust, we analyze and identify 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.

Acknowledgements

Research was supported by the U.S. Department of Transportation under Grant 69A3552348327 for the CARMEN+ University Transportation Center. Cost sharing was supported by affiliates of the 6G@UT center within the Wireless Networking and Communications Group at The University of Texas at Austin.

Executive Summary

The two thrusts in this project are summarized separately below.

Thrust 1: Characterizing low-entropy elements of the Starlink Ku-band signal

We identify and characterize dedicated pilot symbols and other predictable elements embedded within the Starlink Ku-band downlink waveform. Exploitation of these predictable elements enables precise opportunistic positioning, navigation, and timing using compact, low-gain receivers by maximizing the signal processing gain available for signal acquisition and time-of-arrival (TOA) estimation. We develop a post-processing acquisition and demodulation framework to decode Starlink frames and disclose the explicit sequences of the edge pilots—bands of 4QAM symbols located at both edges of each Starlink channel that apparently repeat identically across all frames, beams, channels, and satellites. We further reveal that the great majority of QPSK-modulated symbols do not carry high-entropy user data but instead follow a regular tessellated structure super-imposed on a constant reference template. We demonstrate that exploiting frame-level predictable elements yields a processing gain of approximately 48 dB, thereby enabling low-cost, compact receivers to extract precise TOA measurements even from low-SNR Starlink side beams.

Thrust 2: Determining the source of an important continental-wide interference phenomenon

This work 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 work (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 and Problem Statement

Position, navigation, and timing (PNT) is a crucial utility whose applications span vast sectors of modern life. From agriculture to aviation to financial markets, Global Navigation Satellite Systems (GNSS)---the prevailing technology for PNT---provide the nanosecond-level timing and meter-level positioning necessary for efficient logistics, safe transportation, synchronized telecommunications, and precise financial transactions. Yet legacy GNSS face both a set of critical vulnerabilities and an expanding number of threats. Because legacy civil GNSS transmit weak L-band signals from medium Earth orbit and operate as open-access systems, users risk denial of service and deception by jamming and spoofing attacks. These risks are intensifying as the space domain and electromagnetic spectrum emerge as decisive arenas of 21st-century warfare, with predictable and escalating dangers for civilian users of GNSS-based PNT services.

2. Project Objectives

(1) Offer the fullest account yet of predictable structure in the Starlink waveform.

(a) Reveal the exact information symbols of the edge pilots.

(b) Report a remarkable discovery: after the first few OFDM symbols in each frame, which appear to contain header information, almost all symbols modulated as QPSK are nearly perfectly predictable and thus low-entropy. While this fact is not obvious from a cursory examination of the data, application of certain transformations described in detail herein reveals a highly consistent structure with extended segments of predictable symbols. Fully exploiting these low-entropy signal elements allows matched filtering with sufficient processing gain to support robust PNT with low-cost compact feedhorn receivers.

(2) Determine the source of continental-scale interference affecting Europe since 2019:

(a) Present measurement models and a detection framework for transient wide-area interference.

(b) Detail 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.

(c) Present a basic satellite identification strategy to narrow the candidate satellites and estimate the minimum satellite altitude at apogee.

(d) Present an advanced satellite associate framework using the generalized likelihood ratio test (GLRT) and apply the framework to a test scenario.

(e) Present a framework for instantaneously identifying an interference satellite based on a brief time history of TDOA measurements, and gives an error sensitivity analysis.

(f) Combine 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.

3. Technical Approach and Methodology

These are described in Papers 1 and 2 below.

4. Results and Findings

These are described in Papers 1 and 2 below.

5. Implications for Transportation

These are described in Papers 1 and 2 below.

6. Workforce Development

This work was led by two graduate research assistants. One graduate student (Zach Clements) graduated with his Ph.D. during the course of this study. He was awarded the Brad Parkinson Best Dissertation award from the Institute of Navigation. The other graduate student, Wenkai Qin, will be defending his dissertation in late September 2026.

Zach has taken employment with a U.S. government contractor supporting signals intelligence. Wenkai has taken employment with SpaceX.

7. Publications and Presentations

Qin, Wenkai, Mark L. Psiaki, John R. Bowman, and Todd E. Humphreys. "Pilots and other predictable elements of the Starlink Ku-band downlink." npj Wireless Technology 2, no. 1 (2026): 69. https://doi.org/10.1038/s44459-026-00075-6

Clements, Zachary L., Argyris Kriezis, and Todd E. Humphreys. "Chasing Lightning: Detecting, Characterizing, and Identifying a Powerful Space-Based GNSS Interference Source." arXiv preprint arXiv:2606.03673 (2026). https://arxiv.org/pdf/2606.03673

8. Data Management and Dataset Statement

The data and code supporting the findings of this study are available in the Supplementary Material, hosted at The University of Texas at Austin (https://rnl-data.ae.utexas.edu/datastore/supplementaryMaterial/qin-starlink-pilots/) and archived at Zenodo

(https://doi.org/10.5281/zenodo.21404931). The repository includes (i) decodedframe data and derived products used for the analyses reported in this paper, and (ii) MATLAB code for generating the reference template and edge pilot replica, correlating signals against the edge pilots, and reproducing the diagnostic plots from the included decoded exemplar frames.

The Supplementary Material includes MATLAB code for generating the reference template and edge pilot replica, correlating signals against the edge pilots, and reproducing the diagnostic plots from the included decoded exemplar frames.

9. Matching Funds Statement

Amazon Leo provided matching funds in the amount of $60,000 for this project. The funds supported data collection and equipment purchased to enable data collection (a 1-meter collection dish).

10. Conclusions and Recommendations

These are described in Papers 1 and 2 below.

Pilots and Other Predictable Elements of the Starlink Ku-Band Downlink

Wenkai Qin[1,*], Mark L. Psiaki[2], John R. Bowman[2], & Todd E. Humphreys[1]

[1]*Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin*, Austin, TX, USA

[2]*Department of Aerospace and Ocean Engineering, Virginia Tech*, Blacksburg, VA, USA

[*]Corresponding author: wqin@utexas.edu

***Abstract*—We identify and characterize dedicated pilot symbols and other predictable elements embedded within the Starlink Ku-band downlink waveform. Exploitation of these predictable elements enables precise opportunistic positioning, navigation, and timing using compact, low-gain receivers by maximizing the signal processing gain available for signal acquisition and time-of-arrival (TOA) estimation. We develop a post-processing acquisition and demodulation framework to decode Starlink frames and disclose the explicit sequences of the *edge pilots*—bands of 4QAM symbols located at both edges of each Starlink channel that apparently repeat identically across all frames, beams, channels, and satellites. We further reveal that the great majority of QPSK-modulated symbols do not carry high-entropy user data but instead follow a regular tessellated structure superimposed on a constant reference template. We demonstrate that exploiting frame-level predictable elements yields a processing gain of approximately 48 dB, thereby enabling low-cost, compact receivers to extract precise TOA measurements even from low-SNR Starlink side beams.**

Introduction

Position, navigation, and timing (PNT) is a crucial utility whose applications span vast sectors of modern life. From agriculture to aviation to financial markets, Global Navigation Satellite Systems (GNSS)—the prevailing technology for PNT—provide the nanosecond-level timing and meter-level positioning necessary for efficient logistics, safe transportation, synchronized telecommunications, and precise financial transactions. Yet legacy GNSS face both a set of critical vulnerabilities and an expanding number of threats. Because legacy civil GNSS transmit weak L-band signals from medium Earth orbit and operate as open-access systems, users risk denial of service and deception by jamming and spoofing attacks. These risks are intensifying as the space domain and electromagnetic spectrum emerge as decisive arenas of 21st-century warfare, with predictable and escalating dangers for civilian users of GNSS-based PNT services [1]–[4].

Meanwhile, advances in reusable rocket technology have led to a tenfold reduction in space launch costs, laying the groundwork for a rapidly expanding low Earth orbit (LEO) space economy [5]. Recently launched LEO communications systems have already attracted millions of subscribers and currently generate billions of dollars in annual revenue. These systems have drawn the interest of PNT researchers seeking a robust, precise, and broadly accessible alternative to legacy GNSS: the consistent availability, high power, and wide bandwidth of LEO communications signals confer inherent resilience to adversarial interference and great potential for PNT precision. But despite the significant potential and low marginal downlink capacity cost of fusing a PNT service with an existing broadband LEO communications service [6], no operator of the leading LEO communications systems—Eutelsat OneWeb, SpaceX Starlink, and Amazon Leo—has formally committed to provision of such a service. Meanwhile, researchers have developed techniques to opportunistically exploit LEO communications signals for PNT, with efforts from several groups already bearing fruit [7]–[15]. While many constellations' signals have been shown to be exploitable [7], [9], [15], Starlink remains the most attractive candidate for opportunistic navigation: with a constellation of over 10,000 satellites, Starlink's dense, geometry-rich coverage reaches nearly the entire globe and offers a promising foundation for multilaterated PNT [6].

Nonetheless, several challenges remain before Starlink-based PNT can be a practically viable alternative to legacy GNSS. One major hurdle concerns the size, weight, and power (SWaP) of effective Starlink PNT receiver antennas. Supposing current Starlink phased array terminal form-factors are indicative of a future PNT-capable terminal, even the smallest option currently offered by SpaceX (Starlink Mini) measures $30\,\mathrm{cm} \times 26\,\mathrm{cm}$, weighs $1.1\,\mathrm{kg}$, and requires 25–$40\,\mathrm{W}$ to operate [16], rendering it impractical for many SWaP-constrained applications such as small unmanned aerial vehicles, compact quadrupedal robots, consumer electronics, and asset trackers. While not addressing the most demanding of SWaP constraints, a compact, mass-produced low-noise block (LNB) with an integrated hemispherical feedhorn (such as the unit shown in Fig. 1) offers a more attractive profile, with maximal dimension $12\,\mathrm{cm}$, mass under $200\,\mathrm{g}$, and power consumption of only 1–$2\,\mathrm{W}$.

The use of such a compact feedhorn receiver for Starlink-based PNT would confer an additional benefit: its wide beamwidth enables simultaneous reception of multiple satellite transmissions, which allows multilateration, provided that each transmitting satellite can be uniquely associated with its corresponding component in the received signal. Indeed, the systems presented in [11], [17], [18] exclusively employ compact feedhorn receivers to track so-called pilot tones for

Copyright © 2026 by Wenkai Qin, Mark L. Psiaki, John R. Bowman, and Todd E. Humphreys

Starlink-based Doppler positioning. Unfortunately, these pilot tones' strength and availability have significantly decreased since 2023, rendering this approach ineffective [10].

Similar systems not reliant upon the pilot tones have been introduced, as well. The system presented in [14] geolocates a compact feedhorn receiver to within $6.47\,\mathrm{m}$, with a 95% confidence ellipse whose semi-major and semi-minor axes are $11.96\,\mathrm{m}$ and $6.31\,\mathrm{m}$. This approach, however, relies on a stationary reference receiver that relays high-bitrate data captured using a 35-dBi, 90-cm parabolic dish antenna to aid signal tracking and disambiguation. In [13], researchers describe a pseudorange-based Starlink PNT system with an approximately 20-cm feedhorn (yielding $\sim$15 dBi of gain) coupled to a wideband LNB. Yet the resulting position solutions remain relatively coarse, with miss distances of $56.4\,\mathrm{m}$ or more. The authors of [10] claim meter-level positioning with a single compact feedhorn receiver and pseudorange-rate observables. But pseudorange-rate methods are subject to the same $0.1\,\mathrm{ms}$ timing precision limitation [19]–[21] as Doppler-based methods, making them unsuitable for applications requiring precise timing, such as cellular networks and financial markets.

Thus, although the feasibility of Starlink-based PNT has been repeatedly demonstrated, achieving precise PNT with a compact, low-cost antenna and minimal side-channel aiding remains a challenge.

Ultimately, the practicality and performance of a Starlink-based PNT system are subject to constraints on the receiver's antenna gain, which must provide a sufficient signal-to-noise ratio (SNR) to ensure reliable signal acquisition and satisfactory PNT precision. For instance, consider a Ku-band horn antenna with a circular $54\,\mathrm{mm}$ diameter aperture. Supposing an operating frequency of $12\,\mathrm{GHz}$ and antenna efficiency of $\eta = 0.7$, this antenna's gain is theoretically limited to approximately $15.1\,\mathrm{dBi}$ with an 32° half-power beamwidth, as calculated using standard antenna theory [22, Section 2-11], [23, Eq. (12-29a)]. Assuming clear-sky conditions, an LNB noise figure of $1\,\mathrm{dB}$, and published Starlink Ku-band power flux density (PFD) of -145 to $-138\,\mathrm{dBW}$ per $4\,\mathrm{kHz}$, this antenna yields a maximum pre-correlation SNR in the range of $1\,\mathrm{dB}$ to $8\,\mathrm{dB}$ [23, Sec. 2.16–2.18], [24].

Our own measurements of Starlink pre-correlation SNR from an assigned beam signal (one directed to the receiver's service cell) captured through a compact feedhorn antenna with an integrated LNB lie several decibels below this range, as shown in Fig. 1. This is because, in practice, SNR is degraded by off-axis reception of the satellite beam, polarization mismatch, receiver implementation losses, atmospheric attenuation, and co-channel interference from other beams and satellites.

Nonetheless, the clearly defined correlation peaks in Fig. 1 show that matched filtering of an assigned beam's signal against a local replica composed of the coherent combination of the primary and secondary synchronization sequences (PSS and SSS—described in [25]) can provide ample processing gain for signal acquisition and tracking, assuming the signal is captured with adequate bandwidth and under low-interference conditions. But this result does not necessarily extend to side

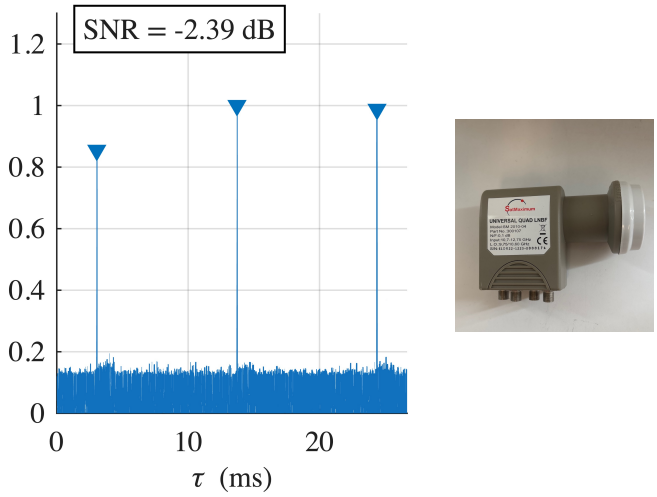


Fig. 1: Normalized matched-filter correlation against a coherent combination of the known primary and secondary synchronization sequences (PSS and SSS) with a signal captured from a compact feedhorn receiver of the type shown. This receiver, intended for consumer-grade satellite television reception in the 10.7–$12.75\,\mathrm{GHz}$ frequency band, provides a 50–$60\,\mathrm{dB}$ gain and a noise figure below $1\,\mathrm{dB}$. The SNR value shown indicates the pre-correlation SNR of the received signal associated with the strongest post-correlation peak. For this capture, which spans the full $240\,\mathrm{MHz}$ bandwidth of a single Starlink channel, the theoretical processing gain by which the post-correlation SNR exceeds the pre-correlation SNR is $33\,\mathrm{dB}$.

beams—those directed to a service cell other than the one in which the receiver resides. The authors of [26] observe that a composite signal's trackable side beams may have SNRs that are 5 to $18\,\mathrm{dB}$ below that of the assigned beam—too low to extract useful time-of-arrival (TOA) measurements solely with the PSS+SSS local replica, as will be shown in "SNR-Dependent TOA Precision". The ability to extract sub-meter-level-precise TOA measurements from several side beams would improve geometric dilution of precision (GDOP) and shorten time-to-first-fix.

The key to unlocking robust and practical Starlink-based PNT with a compact receiver is to boost the processing gain well beyond that provided by the PSS+SSS combination. Processing gain, a fundamental concept in direct-sequence spread-spectrum (DSSS) systems such as code-division multiple access (CDMA) and GNSS, can be extended to opportunistic exploitation of orthogonal frequency division multiplexing (OFDM) waveforms such as Starlink's, in which context the processing gain scales as the number of known information symbols. With adequate processing gain, the side beam signals required for a fast and precise Starlink PNT solution can achieve sufficient post-correlation SNR to permit precise TOA measurement.

We note that, in principle, an opportunistic PNT user could apply the maximum-likelihood (ML) or decision-directed

(DD) methods of [27] to recover TOA measurements despite lack of knowledge of the OFDM information symbols. But these methods break down catastrophically in the low-SNR regime even for low-order modulation (see Fig. 1 in [27]), and so are inapt for Starlink signals received through compact antennas.

While known information symbols are the most straightforward features to exploit *a priori*, knowledge of any invariant property of the Starlink frame likewise increases the processing gain—for instance, the range of modulation types, the temporal arrangement of modulation, or whether individual OFDM symbols are simple in their modulation (e.g., purely QPSK) or composite (e.g., a mixture of QPSK and 16QAM). Accordingly, we aim to identify key predictable features present in the Starlink Ku-band downlink waveform, thereby boosting the achievable processing gain.

Prior work already discovered several predictable features. Reference [25] revealed the Starlink frame structure, including the PSS and SSS and their explicit values. SpaceX's 2024 patent [28] confirmed this frame structure and additionally disclosed the existence of two bands of pilot symbols not detected by [25] near the edges of each Starlink channel's allocated bandwidth. Pilot symbols are complex-valued coefficients that modulate OFDM subcarriers just like other information symbols, but are entirely predictable to the receiver and so can be used for channel estimation. Reference [10] empirically confirmed these pilot bands, hereafter called the edge pilots, and reported additional inter-frame-correlated data not disclosed by [28], but implied that these additional correlated data appear in all frames. This, as we explain herein, does not appear to be true. Moreover, because a local replica constructed via the frame averaging methods presented in [10], [29] retains residual noise, it does not enable the most potent method for signal detection—matched filtering against a noise-free local signal replica.

The purpose of the current paper is to offer the fullest account yet of predictable structure in the Starlink waveform. Not only do we reveal the exact information symbols of the edge pilots, but we also report a remarkable discovery: after the first few OFDM symbols in each frame, which appear to contain header information, almost all symbols modulated as QPSK are nearly perfectly predictable and thus low-entropy. While this fact is not obvious from a cursory examination of the data, application of certain transformations described in detail herein reveals a highly consistent structure with extended segments of predictable symbols. Fully exploiting these low-entropy signal elements allows matched filtering with sufficient processing gain to support robust PNT with low-cost compact feedhorn receivers.

Methods

SNR-Dependent TOA Precision

To assess the conditions under which precise TOA estimation is feasible, we examine the Ziv-Zakai Bound (ZZB), which characterizes estimator performance across a given SNR range by explicitly modeling the transition from small-error behavior in the high-SNR regime, where the Cramér-Rao Bound (CRB) is asymptotically tight, to large-error behavior dominated by noise and ambiguity. In contrast to the CRB, which assumes unbiased estimation in the high-SNR limit, the ZZB captures threshold effects that lead to rapid degradation in TOA precision at low SNR. Reference [30] describes in detail the ZZB formulation and its application to OFDM-based TOA estimation.

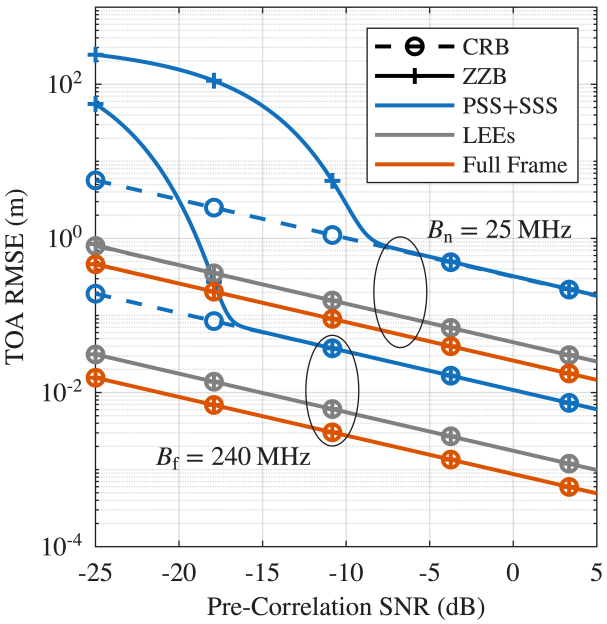


Fig. 2: Analytical lower bounds on single-frame TOA root mean squared error (RMSE) vs. pre-correlation SNR for processing with three representative local replicas: (1) the PSS+SSS combination (blue), (2) low-entropy elements (LEEs) of the frame (gray), and (3) the full Starlink frame (red). Results are shown for two representative capture bandwidths, 25 and 240 MHz. The Ziv-Zakai bounds (ZZBs; solid lines) are tighter than the Cramér-Rao bounds (CRBs; dashed lines) at low SNR because they capture the threshold effects that lead to rapid degradation in TOA precision in that regime. The SNR range shown covers the relevant values for trackable signals from assigned beams (high SNR) to side beams (low SNR). The bounds based on LEEs are discussed further in "Processing Gain and TOA Bounds".

Fig. 2 shows analytical TOA precision bounds as functions of pre-correlation SNR for representative Starlink processing strategies. When correlating against the relatively short PSS+SSS local replica using the full-channel bandwidth $B_f = 240\,\mathrm{MHz}$, the ZZB departs from the CRB for pre-correlation SNR below approximately $-16.5\,\mathrm{dB}$. Narrowing the capture bandwidth to $B_n = 25\,\mathrm{MHz}$ shifts this divergence point to a substantially higher SNR of approximately $-8.5\,\mathrm{dB}$. Below these thresholds, TOA precision degrades rapidly.

In contrast, when correlating against a full-frame local replica—one that spans all subcarriers and OFDM symbols of the Starlink frame—the substantially increased number of coherently accumulated samples yields a greater processing

gain and suppresses ambiguity sidelobes in the correlation function. As a result, the ZZB remains coincident with the CRB across the entire SNR range shown, even when the capture bandwidth is limited to B_{n}. These results indicate that full-frame processing maintains CRB-like TOA estimation over a wider SNR range than PSS+SSS-only processing.

One can think of the PSS+SSS-only and full-frame strategies as extreme cases that bookend reasonable coherent processing for TOA measurements from Starlink signals: There is little reason to process less than the coherent combination of the PSS and SSS, which are known and easily constructed [25], whereas coherent processing beyond one full frame is complicated by inter-frame carrier phase discontinuities, which, so far, have resisted modeling. Thus, Fig. 2 covers all reasonable cases between short-duration processing with narrow-bandwidth captures (the PSS+SSS case with B_{n}) and full-frame processing with full-channel-bandwidth captures (the full-frame case with B_{f}).

Fig. 2 makes it clear that a receiver seeking to suppress the ZZB-predicted breakdown in TOA precision at low SNR must adopt strategies that provide greater processing gain than that afforded by a PSS+SSS matched filter. Such strategies extend the SNR range over which TOA estimation remains asymptotic and enable the extraction of accurate TOA measurements from weaker side beams that would otherwise be unusable.

Processing Gain

Here we define processing gain, extend it to OFDM, and show how any invariant property of the Starlink frame may be exploited to increase it.

For any signal and processing strategy, define the processing gain as

$$L \triangleq \frac{\text{SNR}_{\text{post}}}{\text{SNR}_{\text{pre}}} \tag{1}$$

where SNR_{pre} and SNR_{post} denote pre- and post-correlation SNR, respectively. The former is the ratio of the total signal power to the total noise power within the signal bandwidth. The latter is the ratio of the signal-to-noise power within the de-spread bandwidth, or, equivalently, the SNR of the complex product that results from correlation against whatever local replica is used for detection. For DSSS despreading, this definition recovers the familiar relation $L = B_{\text{spread}}/B_{\text{despread}}$, where B_{spread} and B_{despread} are the spread and despread bandwidths, respectively.

Consider a sequence of complex information symbols subject to additive noise

$$r_k = x_k + n_k, \quad k \in \{1, \ldots, N\} \tag{2}$$

where the kth information signal $x_k \in \mathcal{C}$ is drawn uniformly from the modulation constellation $\mathcal{C} = \{c^{(1)}, c^{(2)}, \ldots, c^{(M)}\}$ and the symbols $\{x_k\}_{k=1}^N$ are statistically independent. The constellation is assumed to be normalized and zero-mean such that $\mathbb{E}\left[|x_k|^2\right] = 1$ and $\mathbb{E}\left[x_k\right] = 0$ for all $k \in \{1, \ldots, N\}$. The additive noise samples $\{n_k\}_{k=1}^N$ are complex Gaussian random variables with $n_k \sim \mathcal{CN}(0, \sigma_n^2)$, mutually independent for all $k \in \{1, \ldots, N\}$, and independent of $\{x_k\}_{k=1}^N$.

Define the signal detection statistic as the complex correlation product

$$S \triangleq \sum_{k=1}^{N} r_k^* l_k = \sum_{k=1}^{N} x_k^* l_k + \sum_{k=1}^{N} n_k^* l_k = S_x + S_n \tag{3}$$

where l_k is the kth element of a local replica sequence and $(\cdot)^*$ denotes the complex conjugate. Pre- and post-correlation SNR can then be defined as

$$\text{SNR}_{\text{pre}} \triangleq \frac{\mathbb{E}\left[x_k^* x_k\right]}{\mathbb{E}\left[n_k^* n_k\right]} = \frac{1}{\sigma_n^2}, \quad \text{SNR}_{\text{post}} \triangleq \frac{\mathbb{E}\left[S_x^* S_x\right]}{\mathbb{E}\left[S_n^* S_n\right]} \tag{4}$$

We additionally define the following correlation parameters for which expectations are taken over the joint distribution of x_k and l_k across all $k \in \{1, \ldots, N\}$:

$$\alpha \triangleq \mathbb{E}\left[|x_k|^2 |l_k|^2\right], \quad \rho \triangleq \mathbb{E}\left[|l_k|^2\right], \quad \mu \triangleq \mathbb{E}\left[l_k^* x_k\right] \tag{5}$$

Now consider L under each of three conditions for $\{x_k\}_{k=1}^N$: (i) completely unknown, (ii) completely known, or (iii) partially known. These conditions can be modeled through the *a priori* probability mass function (PMF) of x_k over $\mathcal{C}$. In the unknown case, the PMF is uniform over $\mathcal{C}$. In the known case, it collapses to a point mass at the transmitted symbol. In the partially known case, side information or invariant properties of the given waveform may be exploited to restrict the PMF to a subset of the available constellation points or otherwise bias it away from a uniform distribution. In Supplementary Note S1 we evaluate L under each condition and provide a general result.

For the special case of constant-modulus constellations (e.g., QPSK) and unit-modulus replicas, we have $|x_k| = |l_k| = 1$, from which it follows that $\alpha = \rho = 1$ and

$$L = 1 + (N-1)|\mu|^2 \tag{6}$$

This result interpolates smoothly between the incoherent limit $L = 1$ for completely unknown data ($|\mu| = 0$), and the fully coherent limit $L = N$ for completely known data ($l_k = x_k$, $|\mu| = 1$).

Note that it makes no difference whether the complex information symbols in (2) represent samples in the time domain, as in DSSS, or the frequency domain, as in OFDM. Since the discrete Fourier transform (DFT) is a linear, unitary transformation that preserves inner products and signal energy, the detection statistic (3) and its properties are preserved through the transformation. Consequently, correlating in the frequency domain across subcarriers yields the same processing gain as correlating in the time domain across samples.

In general, any invariant property of an OFDM waveform can be exploited during the construction of $\{l_k\}_{k=1}^N$ to increase $|\mu|$ and thus L. For example, knowledge that certain OFDM symbols are simple rather than composite in their modulation, or that their information symbols are merely a low-entropy transformation of a known sequence, or knowledge of the guard (cyclic prefix; CP) length N_{g}, can be embedded in the index-dependent PMF for x_k and thus exploited to raise L.

For a generic OFDM waveform, calculation of L is complicated by the CP, which introduces correlation between the information symbols $\{x_k\}_{k=1}^N$, whether viewed in the

frequency or time domain. For the particular case of the Starlink Ku-band downlink waveform, calculation is further complicated by time-domain correlation arising due to the repetitive structure of the PSS and to the mid-channel gutter in which four adjacent subcarriers have zero energy. Ignoring the minor effects of such correlation, one can approximate the effective accumulation length N in (6) for full-frame Starlink processing as

$$N \approx N_{\mathrm{sf}}(N_{\mathrm{s}} + N_{\mathrm{g}}) = 318912 \tag{7}$$

where $N_{\mathrm{sf}} = 302$ is the number of nonzero symbols in a frame, $N_{\mathrm{s}} = 1024$ is the number of subcarriers, and $N_{\mathrm{g}} = 32$ is the CP length. Thus, for the case of full-frame correlation with completely known symbols, $L = N \approx 55$ dB.

Data Corpus

The signal models and observations about data structures that follow are supported by a large corpus of Starlink Ku-band captures spanning many observation sessions and satellites, across which the reported properties were observed consistently. The corpus spans all generations of Starlink satellites to date (blocks 1.0, 1.5, and 2.0-mini) and covers captures from October 2024 through June 2025. Details about each capture are given in Table S1 of the Supplementary Information. To encourage follow-on work and ensure reproducibility, the Supplementary Material contains decoded frame data for a high-quality subset of 1009 decoded frames that we call the *exemplar frames*, which are from a single capture of high-SNR data from one Starlink satellite. The reported low-entropy properties were verified across every decoded frame in the corpus and are likely to generalize across the full constellation, given their consistency across satellite hardware generations, frequency bands, and observation periods.

Signal Model

To facilitate understanding of the OFDM signal processing techniques presented in the sequel, this section presents a joint time- and frequency-domain signal model for the Starlink Ku-Band downlink. The model assumes the basic frame structure detailed in [25]. In brief, Starlink Ku-band downlink data are packaged into frames transmitted at a maximum rate of 750 Hz. Each frame consists of $N_{\mathrm{sf}} = 302$ OFDM symbol slots of nominal length $T_{\mathrm{sym}} = 4.4\,\mu\mathrm{s}$ followed by an unoccupied frame guard interval $T_{\mathrm{fg}} = T_{\mathrm{sym}} + T_{\mathrm{g}}$, where $T_{\mathrm{g}} = 2/15\ \mu\mathrm{s}$ is the CP interval. A sequence of frames from a given Starlink satellite vehicle (SV) to a given service cell is conveyed by a downlink beam whose SV-to-cell assignment remains fixed over a 15-second fixed assignment interval (FAI) [26].

Clock Models: Unlike GNSS and other DSSS-based waveforms, Starlink's symbols, frames, and underlying carrier are not driven directly by the same clock—at least not for block 1.0, 1.5, and 2.0 SVs [26]. Accordingly, we adopt the cascading clock model shown schematically in Fig. 3. The GNSS-disciplined base time t_{d} drives the downstream frame, carrier, and sample clocks. These in turn drive signal generation onboard Starlink SVs. The frame clock sets the start time of frames within an FAI. The carrier clock drives

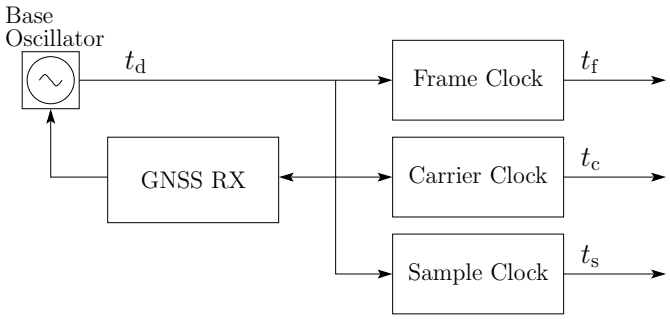


Fig. 3: Cascading clock model for Starlink transmitter showing the relationship between base time t_{d}, frame clock time t_{f}, carrier clock time t_{c}, and sample clock time t_{s}.

generation of the carrier signal used to mix the baseband OFDM waveform to radio frequency (RF). The sample clock drives the digital-to-analog converter used to reconstruct the digital waveform as an analog baseband waveform.

Time according to these three clocks is related to true time t, or time according to an ideal clock, by

$$t = t_{\mathrm{f}}(t) - \delta t_{\mathrm{f}}(t) \tag{8}$$
$$t = t_{\mathrm{c}}(t) - \delta t_{\mathrm{c}}(t) \tag{9}$$
$$t = t_{\mathrm{s}}(t) - \delta t_{\mathrm{s}}(t) \tag{10}$$

where $\delta t_{\mathrm{f}}(t)$, $\delta t_{\mathrm{c}}(t)$, and $\delta t_{\mathrm{s}}(t)$ are the frame, carrier, and sample clock offsets, respectively. Given that this paper addresses the signal processing of individual frames rather than the timing behavior of frame sequences as in [26], the frame clock is not discussed in further detail.

The time derivatives of $\delta t_{\mathrm{c}}(t)$ and $\delta t_{\mathrm{s}}(t)$ are referred to as the carrier and sample clock drifts. For the temperature-controlled crystal oscillators (TCXOs) apparently used in Starlink SVs [26], these may be modeled as constant over a frame interval $T_{\mathrm{f}} = 1/750$ s. This gives rise to a linear clock offset model that is valid for $t_0 \le t \le t_0 + T_{\mathrm{f}}$:

$$\delta t_{\mathrm{s}}(t) = \delta t_{\mathrm{s0}} + \dot{\delta t}_{\mathrm{s}}(t - t_0) \tag{11}$$
$$\delta t_{\mathrm{c}}(t) = \delta t_{\mathrm{c0}} + \dot{\delta t}_{\mathrm{c}}(t - t_0) \tag{12}$$

Here, δt_{c0} and δt_{s0} are the carrier and sample clock offsets at time t_0, while $\dot{\delta t}_{\mathrm{c}}$ and $\dot{\delta t}_{\mathrm{s}}$ are the respective clock drifts, modeled as piecewise constant from frame to frame. Substituting (12) into (9) and solving for $t_{\mathrm{c}}(t)$ yields

$$t_{\mathrm{c}}(t) = t + \delta t_{\mathrm{c0}} + \dot{\delta t}_{\mathrm{c}}(t - t_0) \tag{13}$$
$$= (1 + \dot{\delta t}_{\mathrm{c}})t + (\delta t_{\mathrm{c0}} - \dot{\delta t}_{\mathrm{c}} t_0) \tag{14}$$

Similarly, $t_{\mathrm{s}}(t)$ becomes

$$t_{\mathrm{s}}(t) = (1 + \dot{\delta t}_{\mathrm{s}})t + (\delta t_{\mathrm{s0}} - \dot{\delta t}_{\mathrm{s}} t_0) \tag{15}$$

Transmitted Signal Model: As mentioned, each Starlink frame contains $N_{\mathrm{sf}} = 302$ occupied OFDM symbol intervals of length T_{sym}. We refer to the time-domain signal that applies over the ith symbol interval of the mth frame as $x_{mi}(t)$, where $i \in \mathcal{I} \triangleq \{0, 1, \ldots, N_{\mathrm{sf}} - 1\}$. The first symbol interval is occupied by $x_{m0}(t)$, a known time-domain sequence referred

to as the PSS; the second is occupied by $x_{m1}(t)$, a known time-domain OFDM symbol referred to as the SSS. The remaining $\{x_{mi}(t) \mid i \in [2, N_{\text{sf}} - 1]\}$ are time-domain OFDM symbols containing a mix of header information, pilots, user data, and—as will be shown—low-entropy codes.

Each $x_{mi}(t)$ for $i \in [1, N_{\text{sf}} - 1]$ consists of $N_{\text{s}} = 1024$ mutually orthogonal data subcarriers spread across the channel bandwidth $F_{\text{s}} = 240\,\text{MHz}$, resulting in a carrier spacing of $F = F_{\text{s}}/N_{\text{s}}$. Let $k \in \mathcal{K} \triangleq \{0, 1, \ldots, N_{\text{s}} - 1\}$ represent a single subcarrier's assigned index. We adopt the typical OFDM subcarrier index-to-frequency offset mapping $d[k]$, defined as

$$d[k] = \begin{cases} k, & k \in [0, \frac{N_{\text{s}}}{2} - 1] \\ k - N_{\text{s}}, & k \in [\frac{N_{\text{s}}}{2}, N_{\text{s}} - 1] \end{cases} \tag{16}$$

Each subcarrier is modulated by a complex-valued information symbol $X_{mik} \in \mathcal{C}_{mi}$, which encodes one or more bits of information depending on the modulation scheme corresponding to the constellation $\mathcal{C}_{mi}$ (e.g., BPSK, 4QAM, etc., or a mixture of these). Four mid-channel subcarriers, those with indices in $\mathcal{K}_{\text{g}} = \{0, 1, 1022, 1023\}$, form a gutter in which $X_{mik} = 0$ for $k \in \mathcal{K}_{\text{g}}$. Let $\mathcal{I}_1 \triangleq \mathcal{I} \setminus 0$ be the symbol index set excluding the PSS, which is not a standard OFDM symbol. Then for $i \in \mathcal{I}_1$ the baseband time-domain signal over $0 \le t < T_{\text{sym}}$ can be written as

$$x_{mi}(t) = \frac{1}{\sqrt{N_{\text{s}}}} \sum_{k \in \mathcal{K}} X_{mik} \exp\left(j2\pi d[k] F (t - T_{\text{g}})\right) \tag{17}$$

where $T_{\text{g}} = N_{\text{g}}/F_{\text{s}}$ is the CP interval, $T_{\text{sym}} = T + T_{\text{g}}$ is the full symbol interval, and $T = N_{\text{s}}/F_{\text{s}}$ is the useful (non-cyclic) symbol interval. Note that despite the introduction of $d[k]$ in (17), our definition of X_{mik} is identical to that of [25] for sampling at $t = 0, 1/F_{\text{s}}, \ldots, T_{\text{sym}} F_{\text{s}}$. Let the OFDM symbol support function be defined as

$$g_{\text{s}}(t) = \begin{cases} 2t/T_{\text{g}}, & 0 \le t < T_{\text{g}}/2 \\ 1, & T_{\text{g}}/2 \le t < T + T_{\text{g}}/2 \\ 1 - 2(t - T - T_{\text{g}}/2)/T_{\text{g}}, & T + T_{\text{g}}/2 \le t < T_{\text{sym}} \\ 0, & \text{otherwise} \end{cases} \tag{18}$$

Then, assuming a start time $t_0 = 0$, the baseband time-domain signal over the mth frame may be written as

$$x_m(t) = \sum_{i \in \mathcal{I}} x_{mi}(t - iT_{\text{sym}}) g_{\text{s}}(t - iT_{\text{sym}}), \quad 0 \le t < T_{\text{f}} \tag{19}$$

Starlink OFDM signals are transmitted over one of eight channels spanning the 10.7–12.7 GHz band [25]. In what follows, F_{c} represents an arbitrary channel center frequency selected from among the eight. Assume $x_m(t)$ undergoes digital-to-analog conversion driven by $t_{\text{s}}(t)$, and mixing from baseband to center frequency F_{c} driven by $t_{\text{c}}(t)$. Then, assuming $t_0 = 0$, we arrive at the transmitted mth-frame RF signal model

$$\begin{aligned} \tilde{x}_m(t) &= x_m[t_{\text{s}}(t)] \exp\left[j2\pi F_{\text{c}} t_{\text{c}}(t)\right] && (20) \\ &= x_m\Big[(1 + \delta\dot{t}_{\text{s}})t + \delta t_{\text{s0}}\Big] && (21) \\ &\quad \times \exp\Big\{j2\pi F_{\text{c}}\Big[(1 + \delta\dot{t}_{\text{c}})t + \delta t_{\text{c0}}\Big]\Big\} \end{aligned}$$

Received Signal Model: This section develops a discrete-time received signal model for the LEO-to-Earth channel that includes sample frequency offset (SFO) and carrier frequency offset (CFO) parameters.

Let v_{los} be the scalar velocity between the SV and receiver along the line-of-sight, defined such that $v_{\text{los}} < 0$ when the SV is moving towards the receiver. Define the Doppler parameter $\beta \triangleq v_{\text{los}}/c$, where c is the free-space speed of light. For signals from a LEO SV traveling at an orbital altitude of 340 km with a minimum transmission elevation angle of 10°—corresponding to the lowest possible values currently authorized by the FCC—to a stationary ground receiver, $|\beta|$ is bounded to below 25.6 ppm [31], [32]. As the frame interval $T_{\text{f}} \approx 1.33\,\text{ms}$ is short compared to the timescale over which the SV-receiver geometry evolves, β is modeled as constant over each frame.

Let τ_{los} be the line-of-sight propagation delay for the leading edge of the frame and $w(t)$ be receiver noise, modeled as complex-valued zero-mean additive white Gaussian noise (AWGN) with two-sided power spectral density $N_0/2$ per dimension and uncorrelated with the transmitted signal. Assume that $\beta, \delta\dot{t}_{\text{s}}$, and $\delta\dot{t}_{\text{c}}$ are small enough that their second-order products are negligible. Then the analog RF received signal may be modeled as

$$\begin{aligned} \tilde{y}_m(t) &= \tilde{x}_m\left[(1 - \beta)(t - \tau_{\text{los}})\right] + w(t) && (22) \\ &= x_m\Big[(1 + \delta\dot{t}_{\text{s}} - \beta)(t - \tau_{\text{los}}) + \delta t_{\text{s0}}\Big] && (23) \\ &\quad \times \exp\Big\{j2\pi(1 + \delta\dot{t}_{\text{c}} - \beta)F_{\text{c}}(t - \tau_{\text{los}}) + j2\pi F_{\text{c}}\delta t_{\text{c0}}\Big\} \\ &\quad + w(t) \end{aligned}$$

Let the effects of τ_{los}, $\delta\dot{t}_{\text{s}}$, and δt_{s0} be combined into an equivalent timing offset τ_m, and those of τ_{los}, $\delta\dot{t}_{\text{c}}$, δt_{c0} be combined into an equivalent phase offset ϕ_m:

$$\tau_m = \tau_{\text{los}} - \frac{\delta t_{\text{s0}}}{1 + \delta\dot{t}_{\text{s}} - \beta} \tag{24}$$

$$\phi_m = 2\pi F_{\text{c}}\delta t_{\text{c0}} - 2\pi(1 + \delta\dot{t}_{\text{c}} - \beta)F_{\text{c}}\,\tau_{\text{los}} \tag{25}$$

Then the received signal model simplifies to

$$\begin{aligned} \tilde{y}_m(t) &= x_m\Big[(1 + \delta\dot{t}_{\text{s}} - \beta)(t - \tau_m)\Big] && (26) \\ &\quad \times \exp\Big\{j\Big[2\pi(1 + \delta\dot{t}_{\text{c}} - \beta)F_{\text{c}}t + \phi_m\Big]\Big\} + w(t) \end{aligned}$$

For the data processing described in the sequel, signals were captured using a highly directional antenna and sampled at a complex sampling rate of 250 Msps after mixing to baseband, then resampled at F_{s} such that all subsequent operations may proceed at the native Starlink information symbol rate. The receiver clock driving our capture equipment was a GNSS-disciplined OCXO whose time can be taken as synonymous with true time t. To reduce complexity, we aggregate the clock drift and Doppler parameters into an effective SFO factor $\beta_{\text{s}} \triangleq -\delta\dot{t}_{\text{s}} + \beta$ and an effective CFO factor $\beta_{\text{c}} \triangleq -\delta\dot{t}_{\text{c}} + \beta$. Let t_m be the true start time of the mth frame as received, which is governed by Doppler and by the satellite frame clock. Then, for a sampling interval $T_{\text{s}} = 1/F_{\text{s}}$ and receiver sample index

$n \in \mathbb{Z}$ on $t_m \le nT_\mathrm{s} < t_m + T_\mathrm{f}$, we arrive at the baseband discrete-time signal model

$$\begin{aligned}\tilde{y}_m[n] = x_m[(1-\beta_\mathrm{s})\,(n-n_m)T_\mathrm{s}] \\ \times \exp(-j2\pi\,\beta_\mathrm{c}F_\mathrm{c}\,nT_\mathrm{s} + j\phi_m) + w(nT_\mathrm{s})\end{aligned} \tag{27}$$

where $n_m \in \mathbb{R}$ satisfies $n_m T_\mathrm{s} = \tau_m$. A continuous sequence of received samples across both occupied and unoccupied frame slots is given by

$$\tilde{y}[n] = \sum_{m\in\mathbb{Z}} \tilde{y}_m[n - t_m/T_\mathrm{s}]g_\mathrm{f}(nT_\mathrm{s} - t_m), \quad n \in \mathbb{Z} \tag{28}$$

where

$$g_\mathrm{f}(t) = \begin{cases}1, & 0 \le t < T_\mathrm{f} \\ 0, & \text{otherwise}\end{cases} \tag{29}$$

is the frame support function. Note that, as regards the sample index and frame support function limits, the wide frame guard interval T_fg—during which no signal is present—allows us to ignore small deviations from T_f in the received signal's frame period.

Demodulation

This section describes the post-processing demodulation procedure used to correct residual synchronization errors and recover the transmitted information symbols from a single received frame. Here, we follow the indexing convention established in "Signal Model": the received, equalized, phase-aligned, and hard-decoded counterparts to X_{mik} are denoted $\bar{Y}_{mik}$, $\tilde{Y}_{mik}$, Y_{mik}, and $\hat{X}_{mik}$, respectively, with corresponding OFDM-symbol-level vectors $\bar{\boldsymbol{Y}}_{mi}$, $\tilde{\boldsymbol{Y}}_{mi}$, $\boldsymbol{Y}_{mi}$, $\hat{\boldsymbol{X}}_{mi} \in \mathbb{C}^{|\mathcal{K}|}$.

In brief overview, the demodulation process proceeds as follows. Frame m is detected and coarsely acquired by matched filtering against a Doppler-adjusted PSS+SSS local replica, whereupon standard OFDM processing yields, for each symbol index $i \in \mathcal{I}$, a noisy received vector $\bar{\boldsymbol{Y}}_{mi}$. Residual SFO and CFO errors are then estimated and compensated. Finally, hard-decoding decisions are made to produce the recovered information symbol vectors $\hat{\boldsymbol{X}}_{mi}$.

Coarse Frame Acquisition: Individual Starlink frame detection and acquisition is accomplished through matched filtering with a Doppler-adjusted local baseband replica. Define the coherent concatenation of the time-domain PSS and SSS functions from [25] as

$$x_{m01}(t) \triangleq \begin{cases}x_{m0}(t), & 0 \le t < T_\mathrm{sym} \\ x_{m1}(t - T_\mathrm{sym}), & T_\mathrm{sym} \le t < 2T_\mathrm{sym} \\ 0, & \text{otherwise}\end{cases} \tag{30}$$

Because the PSS and SSS are identical for all $m \in \mathbb{Z}$ and all SVs, $x_{m01}(t)$ can be used for signal detection on all captured data. The full discrete-time local replica used for correlation is the sampled product of the Doppler-adjusted $x_{m01}(t)$ and a complex exponential:

$$\tilde{y}_{01}[n,\beta] \triangleq x_{m01}\,[nT_\mathrm{s}(1-\beta)]\exp\left(-j2\pi\beta F_\mathrm{c} nT_\mathrm{s}\right) \tag{31}$$

Sampling of the local replica is limited to $n \in [0, M(\beta)-1]$ with $M(\beta) = \lceil 2T_\mathrm{sym}/[(1-\beta)T_\mathrm{s}]\rceil$. For $\tilde{y}[n]$ given as in (28), the complex ambiguity function is formed as

$$R[k,\beta] = \sum_{n=0}^{M(\beta)-1} \tilde{y}[n+k]\,y_{01}^*[n,\beta] \tag{32}$$

The nearest-sample timing offset estimate $\hat{n}_m$ and initial Doppler estimate $\hat{\beta}_m$ are obtained via a two-dimensional search over lag and Doppler hypotheses:

$$(\hat{n}_m, \hat{\beta}_m) = \underset{k,\beta}{\operatorname{argmax}}\,|R[k,\beta]| \tag{33}$$

These estimates are accepted if $|R[\hat{n}_m, \hat{\beta}_m]|$ exceeds a detection threshold, whereupon the mth Starlink frame is extracted from the continuous stream as

$$\tilde{y}_m[n] = \tilde{y}[n + \hat{n}_m], \quad n = 0, 1, \ldots, N_\mathrm{sf}(N_\mathrm{s}+N_\mathrm{g}) - 1 \tag{34}$$

Using the initial Doppler estimate $\hat{\beta}_m$, the extracted frame is coarsely Doppler compensated via time resampling and carrier frequency compensation. The resulting signal can be modeled in terms of residual synchronization errors as

$$\begin{aligned}\bar{y}_m[n] = x_m\left[(1-\delta\beta_\mathrm{s})\,(n-\delta n)T_\mathrm{s}\right] \\ \times \exp\left[-j2\pi\,\delta\beta_\mathrm{c}F_\mathrm{c}\,nT_\mathrm{s} + j\phi_m\right] + w(nT_\mathrm{s})\end{aligned} \tag{35}$$

where $\delta n = n_m - \hat{n}_m$ denotes the fractional residual timing offset, $\delta\beta_\mathrm{s} = \beta_\mathrm{s} - \hat{\beta}_m$ the residual SFO factor, and $\delta\beta_\mathrm{c} = \beta_\mathrm{c} - \hat{\beta}_m$ the residual CFO factor. We assume $|\delta n|\,T_\mathrm{s} < T_\mathrm{g}$ and $|\delta\beta_\mathrm{c}|\,F_\mathrm{c} \ll F$ so that CP removal and frequency transformation yield negligible inter-symbol interference (ISI) and inter-carrier interference (ICI). Next, $\tilde{y}_m[n]$ is segmented into N_sf blocks of $N_\mathrm{s}+N_\mathrm{g}$ samples each. Each block undergoes CP removal and conversion to the frequency domain via the FFT. In the frequency domain, the ith OFDM symbol's kth received data symbol $\bar{Y}_{mik}$ can be modeled as the following function of the transmitted data symbol X_{mik}:

$$\bar{Y}_{mik} = \alpha_{mik}X_{mik} + V_{mik} \tag{36}$$

$$\alpha_{mik} = Z_m\,H_k\,\exp(-j2\pi d[k]F\tau_{mi} + j\phi_{mi}) \tag{37}$$

Here, V_{mik} denotes frequency-domain AWGN, H_k is the frequency-selective channel response at subcarrier k normalized such that $H_0 = 1$, and $Z_m = \sqrt{g}\exp(j\theta_m)$ is a frame-specific complex constant, independent of k, with channel gain g and initial phase offset θ_m. The phase offset θ_m does not appear to be arbitrary from frame to frame, but our efforts at modeling it have so far not successfully generalized to all frames and all Starlink SVs. The per-symbol time delay τ_{mi} and phase offset ϕ_{mi} relate to $\delta\beta_\mathrm{s}$ and $\delta\beta_\mathrm{c}$ as follows:

$$\tau_{mi} = \delta nT_\mathrm{s} + iT_\mathrm{sym}\left(\frac{\delta\beta_\mathrm{s}}{1-\delta\beta_\mathrm{s}}\right) \approx \delta nT_\mathrm{s} + iT_\mathrm{sym}\delta\beta_\mathrm{s} \tag{38}$$

$$\phi_{mi} = \phi_m - 2\pi iT_\mathrm{sym}\delta\beta_\mathrm{c}F_\mathrm{c} \tag{39}$$

Channel Equalization: The channel transfer function H_k may not have a flat amplitude response or a linear phase response over the entire 240-MHz band of a Starlink channel, in which case channel equalization is helpful before proceeding with demodulation. As a native time-domain signal, the PSS

is excluded from this process; thus, our concern is the symbol index set $\mathcal{I}_1$.

Let $\mathcal{K}_\text{l} \triangleq \mathcal{K} \setminus \mathcal{K}_\text{g}$ denote the set of loaded (non-gutter) subcarriers. The channel transfer function can be identified by using the known transmitted SSS symbols X_{m1k} for all $k \in \mathcal{K}_\text{l}$. Suppose that the corresponding received demodulated values $\bar{Y}_{m1k}$ have been determined using the techniques from the preceding subsection, and suppose we define $\phi_{m1} \triangleq 0$. Then the following relationship holds true for all $k \in \mathcal{K}_\text{l}$:

$$\bar{Y}_{m1k} = Z_m H_k \exp(-j2\pi d[k] F \tau_{m1}) X_{m1k} + V_{m1k} \quad (40)$$

Under the constraint $H_0 = 1$, a pre-equalization maximum-likelihood estimate of τ_{m1} can be obtained from $\bar{Y}_{m1k} X^*_{m1k}$, $k \in \mathcal{K}_\text{l}$ in the same way that frequency is estimated in [33]. This allows the exponential factor in (40) to be eliminated, yielding

$$\bar{Y}_{m1k} = Z_m H_k X_{m1k} + V_{m1k} \quad (41)$$

One can then estimate H_k (to within the constant Z_m) by computing $\breve{H}_k = \bar{Y}_{m1k}/X_{m1k}$ for all $k \in \mathcal{K}_\text{l}$. To mitigate the effects of noise, one may apply a filter (e.g., a Savitzky-Golay filter) or a spline fit to the $\breve{H}_k$ for all $k \in \mathcal{K}_\text{l}$ to produce a smoothed function $\tilde{H}_k$. One obtains the value $\tilde{H}_0$ from the smoothed function and uses this to compute a normalized smoothed estimate of the transfer function, $\hat{H}_k = \tilde{H}_k/\tilde{H}_0$. This technique can be used for multiple successive frames of a given receiver/satellite/band combination to refine the channel transfer function estimate.

Fig. 4 shows an example estimated transfer function. Note how the amplitude falls off going from the low-frequency section to the high-frequency section. It also falls off rapidly near the two band edges. The phase stays relatively flat through most of the band, but changes rapidly near the band edges at the same frequencies that experience sharp amplitude roll-off.

The frame-specific constant Z_m is estimated as $\hat{Z}_m = \sqrt{\hat{g}} \exp(j\hat{\theta}_m)$, where $\hat{g}$ is the estimated channel gain and $\hat{\theta}_m$ is the estimated initial carrier phase:

$$\hat{g} = \frac{1}{|\mathcal{K}_\text{l}|} \sum_{k \in \mathcal{K}_\text{l}} |\bar{Y}_{m1k}/\hat{H}_k|^2, \quad \hat{\theta}_m = \angle \sum_{k \in \mathcal{K}_\text{l}} (\bar{Y}_{m1k}/\hat{H}_k) X^*_{m1k} \quad (42)$$

Normalization by $\hat{H}_k$ and $\hat{Z}_m$ equalizes the received symbols for all $k \in \mathcal{K}_\text{l}$ and for all $i \in \mathcal{I}_1$:

$$\tilde{Y}_{mik} = \bar{Y}_{mik}/(\hat{Z}_m \hat{H}_k) \quad (43)$$

The equalized symbols can be modeled by the following simplification of (36)

$$\tilde{Y}_{mik} = \exp(-j2\pi d[k] F \tau_{mi} + j\phi_{mi}) X_{mik} + \tilde{V}_{mik} \quad (44)$$

with $\tilde{V}_{mik}$ being complex AWGN and τ_{mi} and ϕ_{mi} given as in (38) and (39).

Residual Synchronization Parameter Estimation: As is clear from (44), each $\tilde{\boldsymbol{Y}}_{mi}$, $i \in \mathcal{I}_1$ remains distorted by residual synchronization errors and phase misalignment. We remedy this via the multi-stage process described in [34] and summarized below.

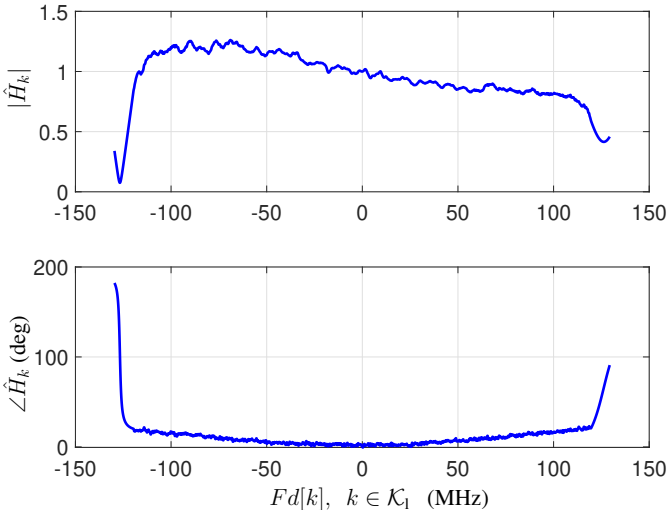


Fig. 4: Example estimated Starlink channel transfer function $\hat{H}_k$ as a function of offset frequency $Fd[k]$ for $k \in \mathcal{K}_\text{l}$. This transfer function is for a data recording system that uses an Ettus X410 USRP with native sampling at 250 MHz and a 1.2-m steered-dish antenna.

First, each constellation $\mathcal{C}_{mi}$ must be identified for $i \in \mathcal{I}_1$. Starlink employs several different quadrature amplitude modulation (QAM) schemes to encode a sequence of digital data onto the N_s complex information symbols that modulate the subcarriers of each OFDM symbol. We distinguish 4QAM from QPSK constellations according to the convention

$$\mathcal{C}^{\text{QPSK}} = \{1, j, -1, -j\}, \quad \mathcal{C}^{\text{4QAM}} = \mathcal{C}^{\text{QPSK}} \exp(j\pi/4) \quad (45)$$

Thus, $\mathcal{C}^{\text{4QAM}}$ is equivalent to $\mathcal{C}^{\text{QPSK}}$ modulo a 45° rotation. We exclude the edge pilots (formally introduced in "Edge Pilots") from the determination of $\mathcal{C}_{mi}$ because edge pilot symbols are uniformly 4QAM. Let $\mathcal{K}_\text{p}$ be the subcarrier indices corresponding to the edge pilots, and let $\mathcal{K}_\text{lnp} \triangleq \mathcal{K}_\text{l} \setminus \mathcal{K}_\text{p}$ be the set of loaded non-pilot subcarrier indices. Only subcarriers $k \in \mathcal{K}_\text{lnp}$ are involved in constellation identification.

We identify $\mathcal{C}_{mi}$ by estimating its cardinality $|\mathcal{C}_{mi}|$ from $\tilde{\boldsymbol{Y}}_{mi}$ via k-means clustering. If the estimated cardinality is 16 or 32 we select $\hat{\mathcal{C}}_{mi} = \mathcal{C}^{\text{16QAM}}$ or $\hat{\mathcal{C}}_{mi} = \mathcal{C}^{\text{32QAM}}$ as appropriate; for cardinality 4, we initially assume $\hat{\mathcal{C}}_{mi} = \mathcal{C}^{\text{QPSK}}$ and defer the QPSK/4QAM disambiguation to after residual synchronization compensation. For rare cases of other estimated cardinality (e.g., 8), we declare the OFDM symbol to have composite modulation within $\mathcal{K}_\text{lnp}$ and discard it. Let $\tilde{\mathcal{I}}_m \subseteq \mathcal{I}_1$ be the set of all OFDM symbols retained after this process as applied to the mth frame.

Given $\hat{\mathcal{C}}_{mi}$, we estimate the residual delay and phase parameters τ_{mi} and ϕ_{mi} from (44) using a symbol-marginalized maximum-likelihood (ML) formulation. Since channel equalization has already normalized the symbols to approximately unit power, the ML formulation focuses solely on phase and delay estimation. Let $\boldsymbol{\theta}_{mi} \triangleq [\tau_{mi}, \phi_{mi}]$ and let

$$\Lambda_{\tilde{\boldsymbol{Y}}_{mi}}(\boldsymbol{\theta}_{mi} \mid \hat{\mathcal{C}}_{mi}) = \prod_{k \in \mathcal{K}_\text{l}} p(\tilde{Y}_{mik} \mid \boldsymbol{\theta}_{mi}, \hat{\mathcal{C}}_{mi}) \quad (46)$$

be the likelihood function of the vector measurement $\tilde{\boldsymbol{Y}}_{mi}$, where $p(\tilde{Y}_{mik} \mid \boldsymbol{\theta}_{mi}, \hat{\mathcal{C}}_{mi})$ is the likelihood of $\tilde{Y}_{mik}$ for a given $\boldsymbol{\theta}_{mi}$ assuming $\mathcal{C}_{mi} = \hat{\mathcal{C}}_{mi}$. For data-bearing information symbols, the likelihood function is obtained by marginalizing over the possible values for the unknown symbol X_{mik}, assuming complex Gaussian noise under a uniform prior on $\hat{\mathcal{C}}_{mi}$:

$$p(\tilde{Y}_{mik} \mid \boldsymbol{\theta}_{mi}, \hat{\mathcal{C}}_{mi}) = \frac{1}{|\hat{\mathcal{C}}_{mi}|} \sum_{c \in \hat{\mathcal{C}}_{mi}} p(\tilde{Y}_{mik} \mid X_{mik} = c, \boldsymbol{\theta}_{mi}) \tag{47}$$

Pilot information symbols are handled as the special case where X_{mik} is known and no marginalization is required. The ML estimate of $\boldsymbol{\theta}_{mi}$ may be found by

$$\hat{\boldsymbol{\theta}}_{mi} = \underset{\boldsymbol{\theta}_{mi}}{\operatorname{argmax}} \log \Lambda_{\tilde{\boldsymbol{Y}}_{mi}}(\boldsymbol{\theta}_{mi} \mid \hat{\mathcal{C}}_{mi}) \tag{48}$$

This process yields a sequence of per-OFDM-symbol delay and phase estimates $\{\hat{\tau}_{mi}, \hat{\phi}_{mi} \mid i \in \tilde{\mathcal{I}}_m\}$. While these estimates are statistically efficient on a per-symbol basis, they are noisy and do not obey the temporal structure imposed by (38) and (39) wherein τ_{mi} and ϕ_{mi} vary linearly with OFDM symbol index i. To exploit this structure and obtain a unified set of frame-level residual synchronization parameters, the per-symbol ML estimates are combined by a joint ML fit as follows. Let $\hat{\boldsymbol{\phi}}_m = [\hat{\phi}_{mi} \mid i \in \tilde{\mathcal{I}}_m]$ and $\hat{\boldsymbol{\tau}}_m = [\hat{\tau}_{mi} \mid i \in \tilde{\mathcal{I}}_m]$ respectively denote column vectors of per-OFDM-symbol phase and delay estimates obtained from the preceding ML stage for frame m. Likewise, let the vector $\boldsymbol{i} = [i \mid i \in \tilde{\mathcal{I}}_m]$ contain the relevant symbol indices, and $\mathbf{1}$ be the all-ones vector of matching dimension. Let $\boldsymbol{\psi}_m \triangleq [\phi_{m0}, \delta\beta_{\mathrm{c}}, \tau_{m0}, \delta\beta_{\mathrm{s}}]$ contain the parameters to be estimated, and define $\boldsymbol{J}(\boldsymbol{\psi}_m)$ as

$$\boldsymbol{J}(\boldsymbol{\psi}_m) \triangleq \begin{bmatrix} \hat{\boldsymbol{\phi}}_m & - & \left(\phi_{m0}\mathbf{1} - 2\pi T_{\mathrm{sym}} F_{\mathrm{c}} \delta\beta_{\mathrm{c}}\, \boldsymbol{i}\right) \\ \hat{\boldsymbol{\tau}}_m & - & \left(\tau_{m0}\mathbf{1} - T_{\mathrm{sym}}\, \delta\beta_{\mathrm{s}}\, \boldsymbol{i}\right) \end{bmatrix} \tag{49}$$

Minimizing the squared magnitude of this function yields the ML estimate for $\boldsymbol{\psi}_m$:

$$\hat{\boldsymbol{\psi}}_m = \underset{\boldsymbol{\psi}_m}{\operatorname{argmin}} \ \left\| \boldsymbol{J}(\boldsymbol{\psi}_m) \right\|_2^2 \tag{50}$$

Compensation of the residual synchronization errors yields the phase-aligned symbols for all $k \in \mathcal{K}_{\mathrm{l}}$ and $i \in \tilde{\mathcal{I}}_m$:

$$Y_{mik} = \tilde{Y}_{mik} \exp\left(-j\hat{\phi}_{mi} + j2\pi d[k] F \hat{\tau}_{mi}\right) \tag{51}$$

Information Symbol Detection: With the phase-aligned information symbol vector $\boldsymbol{Y}_{mi}$ available for all $i \in \tilde{\mathcal{I}}_m$, the QPSK/4QAM ambiguity for cardinality-4 symbols is resolved: For each $\boldsymbol{Y}_{mi}$, we select $\hat{\mathcal{C}}_{mi} = \mathcal{C}^{\mathrm{QPSK}}$ or $\hat{\mathcal{C}}_{mi} = \mathcal{C}^{\mathrm{4QAM}}$ depending on the relative alignment against $\boldsymbol{Y}_{m(i-1)}$, with $\boldsymbol{Y}_{m1}$, the phase-aligned SSS, providing an absolute reference.

Each hard-decoded symbol estimate $\hat{X}_{mik}$ is then obtained by mapping Y_{mik} to its nearest constellation point in $\hat{\mathcal{C}}_{mi}$:

$$\hat{X}_{mik} = \underset{c \in \hat{\mathcal{C}}_{mi}}{\operatorname{argmin}} \ |Y_{mik} - c| \tag{52}$$

This process produces the frame-level information symbol matrix

$$\hat{\boldsymbol{X}}_m \triangleq \left[\hat{\boldsymbol{X}}_{mi} \mid i \in \tilde{\mathcal{I}}_m\right] \in \mathbb{C}^{|\mathcal{K}| \times |\tilde{\mathcal{I}}_m|} \tag{53}$$

whose columns correspond to hard-decoded frequency-domain OFDM symbols. Fig. 5 illustrates the estimated modulation scheme for each OFDM symbol across the exemplar frames.

Alternate Residual Synchronization Parameter Estimation

A more nearly optimal method was developed to remove the residual carrier Doppler shift error $\delta\beta_{\mathrm{c}}$ and phase ϕ_{m0} when determining the edge pilot symbols. In addition to marginalizing over $\mathcal{C}_{mi}$, as in (47), this method marginalized the probability distribution for $\tilde{\boldsymbol{Y}}_{mi}$ over the 4 possibilities for the ith symbol interval, namely QPSK, 4QAM, 16QAM, and 32QAM. This marginalization procedure assumed that each of these 4 possibilities is equally likely, despite evidence to the contrary. The marginalized probability density function for $\tilde{\boldsymbol{Y}}_{mi}$ was conditioned on the residual phase angles of its $\tilde{Y}_{mik}$ components

$$\phi_{mik} = \phi_{m0} - 2\pi(\delta\beta_{\mathrm{c}} F_{\mathrm{c}+} \delta\beta_{\mathrm{s}} F d[k])(T_{\mathrm{sym}}/[1 - \hat{\beta}_m])i \tag{54}$$

The negative-log-likelihood functions for all of the $\tilde{\boldsymbol{Y}}_{mi}$ symbol vectors were summed for $i \in \mathcal{I}_2$, where $\mathcal{I}_2 \triangleq \{2, 3, \ldots, 301\}$ denotes OFDM symbol indices excluding the PSS ($i = 0$) and SSS ($i = 1$). Two additional terms were added to penalize non-zero $\delta\beta_{\mathrm{c}}$ and the difference between ϕ_{m0} and its estimate based purely on the PSS and SSS symbols. The resulting negative-log-likelihood cost function was minimized to obtain estimates of $\delta\beta_{\mathrm{c}}$ and ϕ_{m0}. This minimization applied the approximation $\delta\beta_{\mathrm{s}} = \delta\beta_{\mathrm{c}}$ in the ϕ_{mik} formula above. It started with a brute-force search over an extended $\delta\beta_{\mathrm{c}}$ range in order to get near the global minimum instead of a local minimum. Afterwards, it applied Newton's method to find the optimal estimates of $\delta\beta_{\mathrm{c}}$ and ϕ_{m0}. These estimates were then used to compute the phase-aligned, demodulated symbols $\{Y_{mik} \mid i \in \mathcal{I}_2, k \in \mathcal{K}_{\mathrm{l}}\}$.

All loaded subcarriers—including those for the edge pilots—were included in this data fitting procedure so that the resulting Y_{mik} values could be used generally; e.g., to determine whether edge pilots and other predictable symbols exist.

RESULTS

Edge Pilots

We developed a tailored averaging procedure to search for information symbols that repeat exactly from frame to frame in the Starlink downlink. The procedure computed $|\mathcal{I}_2| \times |\mathcal{K}_{\mathrm{l}}| = 306000$ averages of the received, equalized, and phase-aligned symbols Y_{mik}—one average per unique combination of $i \in \mathcal{I}_2$ and $k \in \mathcal{K}_{\mathrm{l}}$. The averaging was performed across frames (i.e., over multiple m values) from multiple time intervals, multiple Starlink frequency bands, and multiple Starlink satellites (see Supplementary Note S2 for details). If the same coefficient X_{mik} is present for all frames, then the resulting average approaches X_{mik}. If, instead, X_{mik} varies with m, then the average approaches zero. Combinations (i, k) that produced large absolute values of the corresponding averages were examined to ensure that the individual Y_{mik} values used to produce the given average appeared to be noisy versions of the same frame-to-frame-repeated symbol X_{ik}.

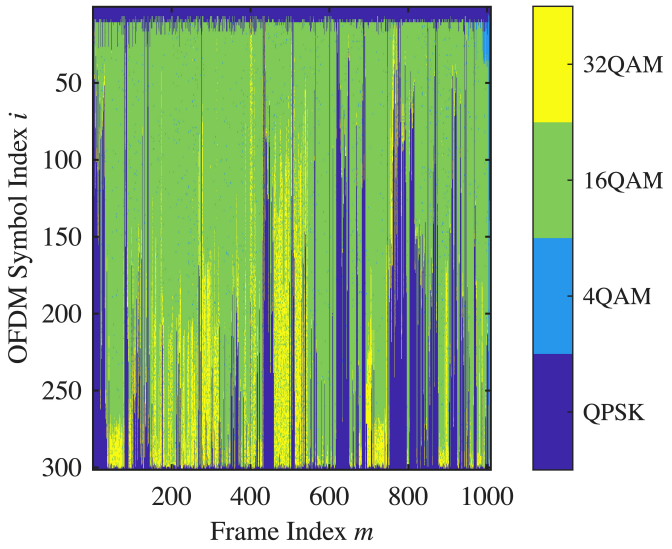


Fig. 5: Modulation scheme estimates for the set of 1009 exemplar frames, with each frame comprising $|\mathcal{I}_1| = 301$ OFDM symbols. Frames and symbols are ordered along the horizontal and vertical axes, respectively. The color of each cell indicates the estimated $\mathcal{C}_{mi}$ type for $m \in \{1, \ldots, 1009\}$ and $i \in \mathcal{I}_1$. Note how the first few symbols (starting with the SSS) of every frame are QPSK-modulated. These symbols constitute the variable-length frame header.

Application of this averaging procedure revealed two bands of 4QAM-modulated pilot symbols, one near each edge of a Starlink channel's allocated bandwidth. The existence of these edge pilots was reported in [10], [28], but not their symbol values, which are published here. Each of the two edge pilot bands occupies 8 subcarriers. The relevant indices of both bands are $\mathcal{K}_\mathrm{p} \triangleq \{488, 489, ..., 495, 528, 529, ..., 535\}$. A 4QAM pilot coefficient is modulated onto each $k \in \mathcal{K}_\mathrm{p}$ for each $i \in \mathcal{I}_2$.

Notably, we found that the edge pilots appear to be identical, not only across all frames from the same Starlink SV, but also across all SVs in the constellation. The pilot coefficient X_{mik} takes the following form for OFDM symbol $i \in \mathcal{I}_2$ and subcarrier $k \in \mathcal{K}_\mathrm{p}$:

$$X_{mik} = \exp\left[j\frac{\pi}{2}\left(s_{\mathrm{p}ik} + \frac{1}{2}\right)\right] \tag{55}$$

$$s_{\mathrm{p}ik} = \left\lfloor \frac{q_{\mathrm{p}k}}{4^{(301-i)}} \right\rfloor \bmod 4 \tag{56}$$

The quantity $s_{\mathrm{p}ik} \in \{0, 1, 2, 3\}$ for $i \in \mathcal{I}_2$ and $k \in \mathcal{K}_\mathrm{p}$ is a function of $q_{\mathrm{p}k}$, a 150-digit hexadecimal number given in Supplementary Note S3 for each $k \in \mathcal{K}_\mathrm{p}$. The foregoing results are identical to those in [35] except for slight notational adjustments and for the basis of representation (base-4 vs. hexadecimal).

Low-Entropy QPSK Symbols

This section reveals the surprising fact that nearly all QPSK-modulated information symbols in the Starlink Ku-band downlink are not high-entropy user data. Rather, outside a short header region, these symbols exhibit a highly repetitive structure across the full capture corpus: deviations relative to a reference set of information symbols called the *reference template* are BPSK-valued and follow a simple periodicity with deterministic shifts across OFDM symbols. This structural repetition enables construction of long, near noise-free local replicas for matched filtering, boosting processing gain.

Reference Template: Our analysis of low-entropy structure (aside from the edge pilots) in decoded frame data is based on the exemplar frames mentioned in "Data Corpus", with discovered patterns verified against data from the larger corpus. The exemplar frames, denoted $\{\hat{\boldsymbol{X}}_m \mid m \in \mathcal{M}_\mathrm{e}\}$ with $\mathcal{M}_\mathrm{e} = \{1, 2, \ldots, 1009\}$, all come from a single capture of data from one SV. They are ordered chronologically, with no frames from occupied slots discarded. The pre-correlation SNR for each exemplar frame is greater than 9 dB and the average over the exemplar set is 13.8 dB.

To identify portions of the Starlink frame that repeat from frame to frame, we compute pairwise cross-correlations. We begin by defining several useful sets and measures. Let the set of non-SSS QPSK-modulated symbol indices for frame m be

$$\mathcal{I}_{\mathrm{Q}m} \triangleq \{\, i \in \mathcal{I}_2 \mid \hat{\mathcal{C}}_{mi} = \mathcal{C}^{\mathrm{QPSK}} \,\} \tag{57}$$

and the corresponding QPSK ratio be $\rho_m \triangleq |\mathcal{I}_{\mathrm{Q}m}|/|\mathcal{I}_2|$. For a pair of frames $\{m, n\}$, define the indices of symbols that are QPSK in both frames as

$$\mathcal{I}_{\mathrm{Q}mn} \triangleq \{\, i \in \mathcal{I}_2 \mid \hat{\mathcal{C}}_{mi} = \hat{\mathcal{C}}_{ni} = \mathcal{C}^{\mathrm{QPSK}} \,\} \tag{58}$$

The corresponding shared QPSK ratio is $\rho_{mn} \triangleq |\mathcal{I}_{\mathrm{Q}mn}|/|\mathcal{I}_2|$. Define the cross-correlation between frames m and n, restricted to their shared QPSK OFDM symbols and loaded non-pilot subcarriers, as

$$R_{mn} \triangleq \sum_{i \in \mathcal{I}_{\mathrm{Q}mn}} \sum_{k \in \mathcal{K}_\mathrm{lnp}} \hat{X}^*_{mik} \hat{X}_{nik} \tag{59}$$

If QPSK-modulated OFDM symbols carried high-entropy user data from frame to frame, then $|R_{mn}|$ would concentrate near zero regardless of ρ_{mn}. An empirical analysis of $|R_{mn}|$ instead shows a clear dependence on ρ_{mn}, as illustrated in Fig. 6: the data form a structured fan with vertically quantized bands. This behavior indicates that QPSK-modulated OFDM symbols exhibit various degrees of correlation across frames and therefore warrant focused analysis.

Further study revealed that pure-QPSK frames, those for which $\rho_m = 1$, exhibit striking consistency: Nearly all information symbols agree across pure-QPSK frames in the exemplar set. Exceptions are a mix of rare hard-decoding errors and some genuine frame-to-frame variation within the frame header. Let the indices of pure-QPSK frames be denoted $\mathcal{M}_\mathrm{ep} \triangleq \{m \in \mathcal{M}_\mathrm{e} \mid \rho_m = 1\}$. We define the *reference template* $\boldsymbol{T} \in \mathbb{C}^{|\mathcal{K}_\mathrm{lnp}| \times |\mathcal{I}_2|}$ as the element-wise mode of the hard-decoded symbols $\{\hat{\boldsymbol{X}}_m \mid m \in \mathcal{M}_\mathrm{ep}\}$. In other words, the (i, k)th element of $\boldsymbol{T}$ for $i \in \mathcal{I}_2$ and $k \in \mathcal{K}_\mathrm{lnp}$ is defined as

$$T_{ik} \triangleq \underset{c \in \mathcal{C}^{\mathrm{QPSK}}}{\operatorname{argmax}} \{\, m \in \mathcal{M}_\mathrm{ep} \mid \hat{X}_{mik} = c \,\} \tag{60}$$

We provide the full reference template $\boldsymbol{T}$ in the Supplementary Material.

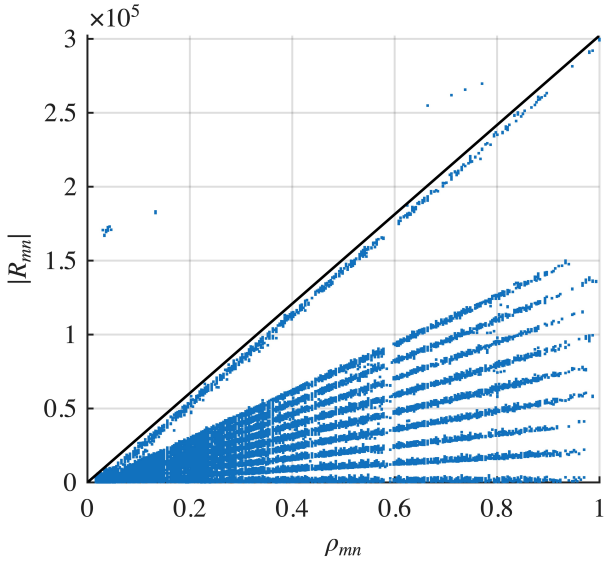


Fig. 6: $|R_{mn}|$ vs. ρ_{mn} for $m, n \in \mathcal{M}_{\mathrm{e}}$, with $n < m$ to avoid duplicate pairwise comparisons. The black line shows the upper limit of $|R_{mn}|$, which obtains when all QPSK-modulated OFDM symbols carry precisely the same information. A small number of points lie above this theoretical limit. Investigation revealed that these correspond to consecutive frames (i.e., $\{m, m+1\}$) where the SV appears to have retransmitted duplicate data.

Tessellation Codes: This section leverages the reference template $\boldsymbol{T}$ to expose the highly regular structure of QPSK-modulated OFDM symbols, then introduces the repeating bit patterns that underlie this structure.

To quantify deviations from the reference template, we restrict $\{\hat{\boldsymbol{X}}_m \mid m \in \mathcal{M}_{\mathrm{e}}\}$ to QPSK-modulated symbols, then compute the element-wise product between the decoded information symbol and the conjugate of the corresponding reference template element. Let D_{mik} denote the (i,k)th element of the deviation, with corresponding OFDM-symbol-level vector $\boldsymbol{D}_{mi} \in \mathbb{C}^{|\mathcal{K}_{\mathrm{lnp}}|}$ and frame-level matrix $\boldsymbol{D}_m \in \mathbb{C}^{|\mathcal{K}_{\mathrm{lnp}}| \times |\mathcal{I}_{\mathrm{Q}m}|}$. The element is defined for $m \in \mathcal{M}_{\mathrm{e}}$, $i \in \mathcal{I}_{\mathrm{Q}m}$, and $k \in \mathcal{K}_{\mathrm{lnp}}$ as

$$D_{mik} \triangleq \hat{X}_{mik} T^*_{ik} \tag{61}$$

Inspection of $\boldsymbol{D}_m$ reveals regular patterns, as follows. BPSK Deviations: A histogram analysis of $\boldsymbol{T}$, and of the set of all QPSK-valued OFDM symbols in the exemplar frames $\{\hat{\boldsymbol{X}}_{mi} \mid m \in \mathcal{M}_{\mathrm{e}}, i \in \mathcal{I}_{\mathrm{Q}m}\}$, and of the set of all deviation matrices $\{\boldsymbol{D}_m \mid m \in \mathcal{M}_{\mathrm{e}}\}$ revealed that, whereas the former two are uniformly distributed on $\mathcal{C}^{\mathrm{QPSK}}$, the deviation matrices are in fact BPSK-modulated; in particular, $D_{mik} \in \{1, -1\}$. Thus, QPSK-modulated OFDM symbols can be expressed as the product of a BPSK-valued sequence and the corresponding column of $\boldsymbol{T}$. Analysis of the sequence of deviation vectors $\{\boldsymbol{D}_{mi} \mid m \in \mathcal{M}_{\mathrm{e}}, i \in \mathcal{I}_{\mathrm{Q}m}\}$ reveals that the QPSK-modulated portion of each frame comprises two distinct temporal regions: a header region and a T-code region.

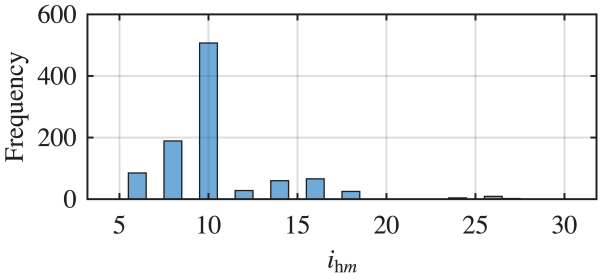


Fig. 7: Distribution of the header length $i_{\mathrm{h}m}$ across the exemplar frames.

Header: After the PSS and SSS, each frame begins with a sequence of QPSK-modulated OFDM symbols that form a header of variable length whose information content differs from the highly regular structure observed later in the frame. In this paper we do not attempt to interpret the header's content semantically; however, preliminary analysis shows multi-modal frame-to-frame repetition.

Let $\mathcal{I}_{\mathrm{H}m} \triangleq \{i \in \mathcal{I}_2 \mid i \leq i_{\mathrm{h}m}\}$ denote the header-symbol indices for the mth frame, where $i_{\mathrm{h}m}$ is the index of the final header symbol. Fig. 7 shows the distribution of $i_{\mathrm{h}m}$ across the exemplar frames. Within the header region, individual OFDM symbols often manifest long sequences of value -1. These can be seen as vertical blue strips on close inspection of Fig. 8. Perhaps certain header regions are reserved for data only transmitted intermittently and are otherwise set to -1. In any case, these runs of constant symbols and other obvious header structure clearly indicate that the header data are not encrypted.

T-Codes: Let $\mathcal{I}_{\mathrm{T}m} \triangleq \{i \in \mathcal{I}_{\mathrm{Q}m} \mid i > i_{\mathrm{h}m}\}$ denote the indices of post-header QPSK-modulated OFDM symbols. Inspection of $\{\boldsymbol{D}_{mi} \mid i \in \mathcal{I}_{\mathrm{T}m}\}$ reveals a tessellated pattern, an example of which is observable in Fig. 8. Between successive OFDM symbols, the deviation's BPSK sequence is circularly shifted by 16 subcarriers from each OFDM symbol to the next. Within each OFDM symbol, the deviation's sequence repeats at 60-subcarrier intervals across the elements of $\mathcal{K}_{\mathrm{lnp}}$, as if gutter and edge-pilot subcarriers were absent. We refer to the 60-bit BPSK sequence as a tessellation code, or *T-code*. Whereas the tessellation periodicity is consistent across frames, the T-codes themselves vary from frame to frame. If non-QPSK OFDM symbols interrupt a continuous sequence of QPSK-modulated OFDM symbols, the pattern continues as if there were no interruption. Continuity is verified by observing that the circular-shift phase accumulates correctly across such gaps. Let $N_{\mathrm{T}} = 60$ be the T-code length and $d_{\mathrm{T}} = 16$ the circular shift period. Then, let $(k_0, k_1, \ldots, k_{1003})$ denote the elements of $\mathcal{K}_{\mathrm{lnp}}$ ordered by ascending physical frequency (i.e., $d[k_0] < d[k_1] < \cdots < d[k_{1003}]$). Consecutive k_n and k_{n+1} are not in general adjacent subcarrier indices, as $\mathcal{K}_{\mathrm{lnp}}$ excludes edge-pilot and gutter subcarriers. Then, for all $i \in \mathcal{I}_{\mathrm{T}m}$ and

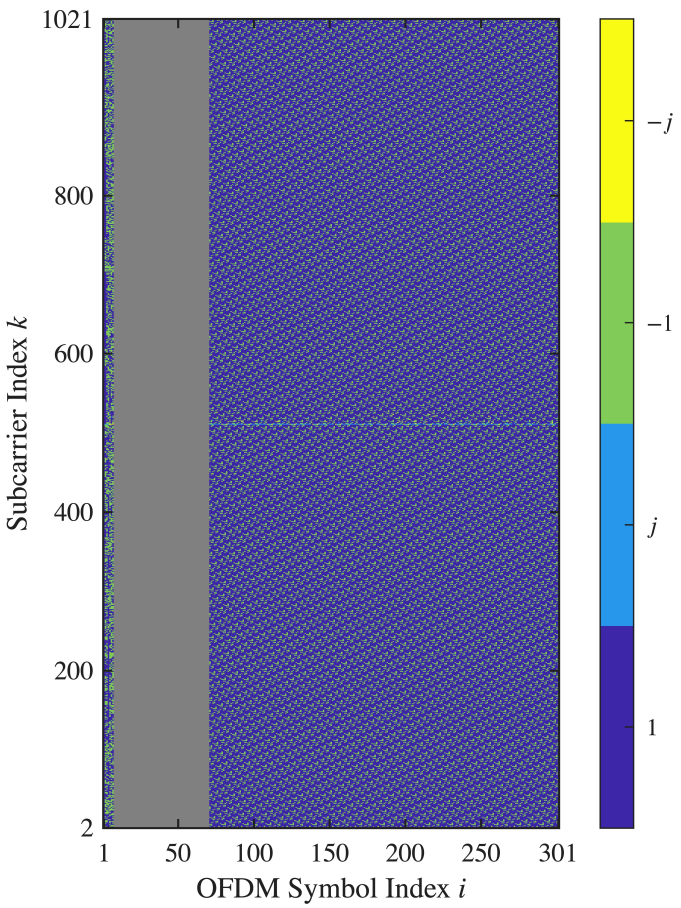


Fig. 8: A color-coded plot of $\boldsymbol{D}_m$ for $m = 814 \in \mathcal{M}_{\text{e}}$. Readers are encouraged to magnify the plot to appreciate its fine structure. The horizontal and vertical axes are ordered respectively by OFDM symbol index $i \in \mathcal{I}_1$ and subcarrier index $k \in \mathcal{K}_{\text{lnp}}$. The narrow vertical strip to the far left is the deviation within the frame's header, $\{\boldsymbol{D}_{mi} \mid i \in \mathcal{I}_{\text{H}m}\}$. The gray region spans the non-QPSK symbols for this frame, with indices $i \in \{\mathcal{I}_2 \backslash \mathcal{I}_{\text{Q}m}\}$. The frame's T-code tiles the remainder of the frame with a tessellated pattern. The thin horizontal line and speckles near the center ($k \approx 511$) are decoding artifacts caused by increased phase noise at the far edges of the Starlink channel. Note how the T-code and the header deviation exhibit only green and blue colors, consistent with $D_{mik} \in \{1, -1\}$.

$0 \leq n < |\mathcal{K}_{\text{lnp}}| - N_{\text{T}}$, we may write

$$D_{mi,k_n} = D_{mi,\,k_{n+N_{\text{T}}}} \tag{62}$$

$$\boldsymbol{D}_{m(i+\Delta i)} = \texttt{circshift}(\boldsymbol{D}_{mi},\ \Delta k) \tag{63}$$

$$\Delta k = (d_{\text{T}} \cdot \Delta i) \bmod N_{\text{T}} \tag{64}$$

where Δi is the OFDM symbol index difference and $\texttt{circshift}(\boldsymbol{x}, d)$ is a function that circularly shifts elements of the vector $\boldsymbol{x}$ by d positions. Because $|\mathcal{K}_{\text{lnp}}|$ is not an integer multiple of N_{T}, the circular shift size must be reduced modulo N_{T} to preserve T-code coherence across all symbol indices. We provide all unique T-codes within the exemplar frames in the Supplementary Material.

T-Codes: Significance, Persistence, and Diversity

Significance: The repeating T-code structure implies that the information content of large blocks within Starlink frames is only $N_{\text{T}} = 60$ bits. If these bits can be determined (e.g., by circularly shifting, stacking, and averaging), then the QPSK-modulated data within each frame can be coherently accumulated. As shown in "Processing Gain", such extended coherent processing directly translates into increased processing gain. Note that T-code symbols appear to be fairly common: across the exemplar frames, 21.5% of all post-header OFDM symbols were T-code symbols.

Persistence: Whether any given signal feature is useful for opportunistic PNT depends on its reliability and persistence. We therefore offer several observations and conjectures regarding the purpose and durability of the T-codes characterized above. SpaceX's 2024 patent [28] indicates that the Starlink SV-side downlink signal processing chain includes both a single-tap whitening scrambler and a low-density parity-check (LDPC) encoder. Whitening scramblers disperse signal energy to ensure a uniform power spectral density, thereby maintaining DC balance and reducing the peak-to-average power ratio of the transmitted OFDM waveform. The use of a single-tap architecture is significant: by applying a deterministic complex-domain phase rotation specific to each subcarrier independently, the transmitter preserves inter-carrier orthogonality and enables computationally efficient, memoryless descrambling at the receiver.

The presence of a constant reference template is consistent with the output of a single-tap scrambler, while the observed tessellation pattern is consistent with the numerological interaction between a 16-wide vectorized hardware processing architecture and a 60-bit periodicity introduced by structured channel coding (e.g., an LDPC encoder). The observed frame-to-frame T-code variations suggest that the scrambler state is re-seeded between frames.

We conjecture that the tessellation pattern shown in Fig. 8, which is present in some form in the majority of exemplar frames, is the result of a single-tap whitening scrambler and an LDPC encoder operating on *unladen* payload allocations—OFDM symbols that are not bearing any user data. Such empty resource blocks arise due to time-varying user demand and unavoidable inefficiency in resource scheduling. Starlink SVs likely populate empty payload allocations with a deterministic fill pattern (e.g., an all-ones vector) prior to scrambling and LDPC encoding. The resulting QPSK symbols thus appear to be user data upon demodulation, yet are in fact low-entropy (60-bit) transformations of the same reference template.

These observations suggest that T-codes are neither arbitrary artifacts nor transient byproducts of a particular software release, but are instead tied to the transmitter's physical-layer architecture. As such, they are likely to remain a durable feature of the Starlink downlink, even if the fraction of OFDM symbols exhibiting T-code structure varies with per-beam user demand. Should a future design alter the LEEs disclosed here, however, the methods described above can be applied to quickly decode frames and re-derive the LEE set.

Diversity: Although a 60-bit-sequence set could theoreti-

cally include 2^{60} members, analysis of T-codes extracted from the exemplar frames suggests a much smaller active set. Of the 1009 exemplar frames, 530 contained T-codes, among which there were only 40 unique sequences. Each unique sequence appeared on average 13.25 times, the least frequent 3 times and the most frequent 24 times. These results suggest that the active T-code family comprises far fewer than 2^{60} unique T-codes.

Processing Gain and TOA Bounds

To understand the implications of these discoveries for PNT, we approximate the potential processing gain now available via exploitation of all low-entropy structure in the Starlink frame known to date, relying on empirical data obtained from the exemplar frames. For this analysis, we make several simplifying assumptions. First, we conservatively ignore the CP on all OFDM symbols other than the PSS and the SSS. Second, we treat the header as unknown data. Third, we assume that T-code symbols can be distinguished from non-T-code symbols even at low SNR.

To calculate the frame-level processing gain, let $\bar{N}$ denote the average number of information symbols per frame to which we can assign a non-uniform distribution. Since symbols with a uniform prior do not contribute to the metric L, we may substitute $\bar{N}$ for N in (6). We count in $\bar{N}$ all PSS, SSS, and edge pilot symbols, as these are fully known. For frames containing T-codes, we assume that the corresponding T-code sequence is hard-decoded by stacking, shifting, and combining the frame's T-code symbols. Let $|\mathcal{I}_{\mathrm{T}m}|$ denote the number of T-code OFDM symbols in frame m, and define the stacking factor for frame m as $M_m = |\mathcal{I}_{\mathrm{T}m}||\mathcal{K}_{\mathrm{lnp}}|/N_{\mathrm{T}}$. Define $\bar{M}$ as the average of M_m over the exemplar frames $m \in \mathcal{M}_{\mathrm{e}}$. Thus, T-code information symbols are included in $\bar{N}$ as the hard-decoded bits that permit the assignment of a non-uniform distribution to each symbol.

Next, we estimate the average (across frames) of the correlation coefficient μ from (5). For each subcarrier $k \in \mathcal{K}_{\mathrm{l}}$, we construct the local replica element l_k as follows. For pilot symbols, l_k is set to the known information symbol value, and since the local replica is exact, $|\mu| = 1$. For T-code symbols, l_k is reconstructed from the hard-decoded T-code bits and the known tessellation pattern. Assume that l_k agrees with the true symbol with probability $(1 - p_{\mathrm{e}})$ and disagrees with probability p_{e}. Then $|\mu| = |1-2p_{\mathrm{e}}|$, where bit-error probability for coherent BPSK demodulation is

$$p_{\mathrm{e}} = \frac{1}{2}\,\mathrm{erfc}\left(\sqrt{\bar{M}\,\mathrm{SNR}_{\mathrm{pre}}}\right) \tag{65}$$

with $\mathrm{erfc}(\cdot)$ being the complementary error function. The average correlation coefficient $\bar{\mu}$ is calculated as the weighted average of $|\mu|$ across all exploitable symbols.

For the exemplar frames, whose mean $\mathrm{SNR}_{\mathrm{pre}}$ is 13.8 dB, $\bar{N} \approx 69{,}500$ and $\bar{\mu} \approx 1$. Applying (6) yields an average processing gain $\bar{L} \approx 48.4\,\mathrm{dB}$, which is only 6.6 dB below the theoretical maximum $L \approx 55\,\mathrm{dB}$ calculated in "Processing Gain" for full-frame correlation with completely known symbols.

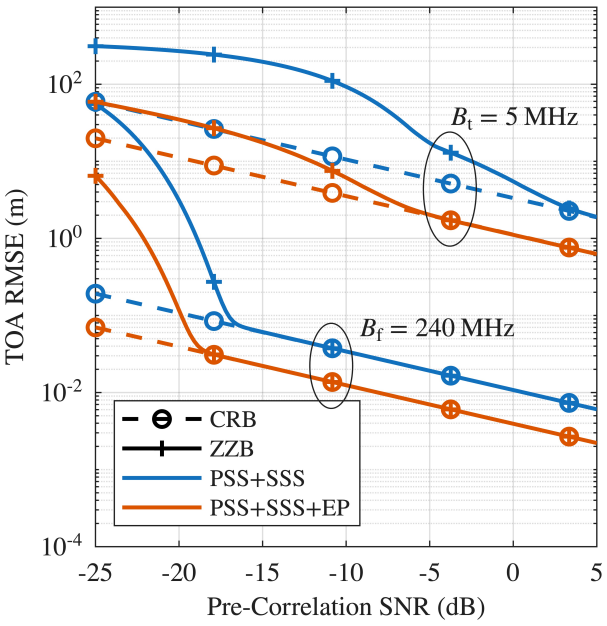


Fig. 9: Lower bounds on single-frame TOA RMSE vs. $\mathrm{SNR}_{\mathrm{pre}}$ for processing with the PSS, SSS, and edge pilots (EP) under two representative capture bandwidths. From top to bottom, bounds are shown for (1) narrowband correlation with a 5 MHz subband of the PSS+SSS combination; (2) narrowband correlation with a coherent combination of the 5 MHz PSS+SSS subband concatenated in time with a single edge pilot band; (3) full-bandwidth correlation over 240 MHz with the PSS+SSS combination, as shown in Fig. 2; and (4) full-bandwidth correlation over 240 MHz with PSS+SSS and both edge pilot bands.

At lower SNR, processing gain degrades due to increased T-code bit estimation errors. However, for the exemplar frames, the large stacking factor $\bar{M} \approx 1044$ provides a boost of approximately 30 dB to the effective SNR for T-code bit decisions. Thus, even at $\mathrm{SNR}_{\mathrm{pre}} = -25\,\mathrm{dB}$, T-code recovery remains nearly perfect, yielding $\bar{\mu} \approx 0.995$ and $\bar{L} \approx 48.4\,\mathrm{dB}$.

The gray traces in Fig. 2 show the resulting single-frame TOA precision bounds when exploiting all known low-entropy elements (LEEs)—the PSS, SSS, edge pilots, and T-codes—assuming T-code occurrence statistics as observed in the exemplar frames. The bounds are computed with the receiver's local oscillator center frequency tuned such that the edge pilots occupy subcarriers at the outer edges of the captured bandwidth, thereby maximizing the mean-square bandwidth and tightening both the ZZB and CRB. Notably, the LEEs ZZB remains coincident with its corresponding CRB across the entire pre-correlation SNR range plotted, matching the behavior of the full-frame bounds. This indicates that the discovered predictable elements are sufficient to suppress ambiguity sidelobes in the local replica correlation function.

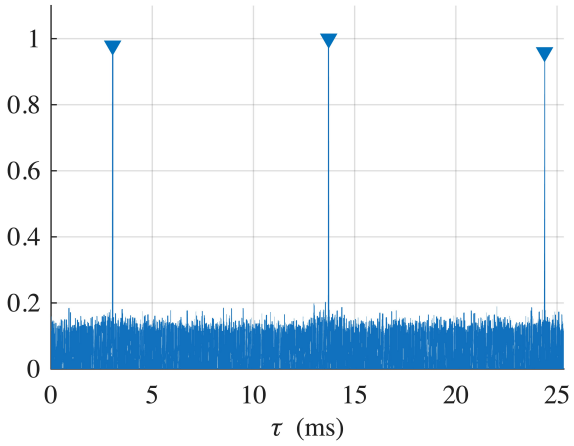


Fig. 10: Normalized matched-filter correlation using a local replica constructed exclusively from the EP, with no PSS or SSS content. Using the same captured signal as in Fig. 1, the EP-only correlation produces detectable peaks at the expected periodicity, aligned with those from the PSS+SSS correlator.

TOA Based on Edge Pilots

A lower-performance but especially simple implementation of Starlink-based opportunistic LEO PNT is possible with only the PSS, SSS, and edge pilots. Fig. 9 shows analytical TOA RMSE bounds as functions of $\mathrm{SNR}_{\mathrm{pre}}$ when the PSS+SSS local replica is augmented with the known edge pilot symbols. For full-bandwidth processing with $B_{\mathrm{f}} = 240\,\mathrm{MHz}$, adding both edge pilot bands yields only a modest change in the ZZB-CRB divergence point, but reduces the TOA RMSE bound by an approximately constant factor of 2.75 across the plotted SNR range due to additional processing gain provided by the edge pilots.

For narrowband processing with a tight $B_{\mathrm{t}} = 5\,\mathrm{MHz}$, augmenting the narrowly filtered PSS+SSS with a single edge pilot provides a larger benefit. Concatenating the edge pilot band with a matched $5\,\mathrm{MHz}$ PSS+SSS subband increases the effective coherent span and significantly suppresses TOA ambiguity, shifting the ZZB-CRB divergence point from approximately 3.3 to $-5\,\mathrm{dB}$.

Fig. 10 provides empirical confirmation that the edge pilots alone provide sufficient processing gain for frame detection.

Other Frame-Invariant Data

Reference [10] reported additional inter-frame-correlated data besides the PSS, SSS, header, and edge pilots, modeling these additional "beacon" data as appearing deterministically in all frames. Fig. 6 does indicate various levels of pairwise correlation between frames, but as we have revealed, this is due to the small-family T-code deviations from the reference template $\boldsymbol{T}$. Since T-codes vary from frame to frame and are approximately zero-mean, there is no guarantee of significant inter-frame correlation of QPSK-modulated data unless one wipes the T-code modulation from the underlying reference template $\boldsymbol{T}$. We carefully searched multiple Starlink satellite signals for the purported additional frame-to-frame repeated data using the averaging technique presented in "Edge Pilots" and did not find them. We conjecture that the additional correlated data found by [10] are not actually invariant across Starlink frames but only appeared so in the cases analyzed therein.

Discussion

We developed a post-processing acquisition and demodulation framework to hard-decode Starlink Ku-band downlink frames and applied it to a corpus of frames captured across multiple satellites and dates. Our analysis revealed key low-entropy structures whose existence or details were previously unknown to the public: edge pilots, the reference template, and T-codes. Exploiting these discoveries yields approximately $48\,\mathrm{dB}$ of processing gain, making precise TOA extraction feasible at minimum SNRs well below those required for communications. These results support robust opportunistic PNT using compact feedhorn receivers and suggest that Starlink side-beam signals may be far more usable for navigation than previously assumed.

Data Availability

The data supporting the findings of this study are available in the Supplementary Material. This repository includes decoded frame data and derived products used for the analyses reported in this paper.

Code Availability

The Supplementary Material includes MATLAB code for generating the reference template and edge pilot replica, correlating signals against the edge pilots, and reproducing the diagnostic plots from the included decoded exemplar frames.

Acknowledgements

Research was supported by the U.S. Department of Transportation (DOT) under Grant 69A3552348327 for the CARMEN+ University Transportation Center, by another U.S. DOT Grant for the CARNATIONS University Transportation Center, by the Department of Defense NDSEG Fellowship program, and by affiliates of the `6G@UT` center within the Wireless Networking and Communications Group at The University of Texas at Austin. The authors alone are responsible for the content of the paper.

Author Contributions

W.Q. led the study, developed the signal processing framework, performed the primary analysis, and wrote the initial manuscript draft. T.E.H. supervised the work and contributed to study design, interpretation of results, and manuscript editing. M.L.P. discovered the edge pilot structure and contributed modeling, analysis, and manuscript text. J.R.B. conducted data collection and pre-processing that supported pilot identification and waveform characterization. All authors reviewed and approved the final manuscript.

Competing Interests

The authors declare no competing financial or non-financial interests.

References

[1] S. McNeil, "NATO says it is working to counter Russia's GPS jamming after interference with EU leader's plane," https://apnews.com/article/europe-security-russia-gps-jamming-nato-rutte-6efed7d030da19a51dfacec39ed77325, Sept. 2025, accessed: 2025-10-24.

[2] European Union Aviation Safety Agency, "Safety Information Bulletin 2022-02R3: Global Navigation Satellite System Outage and Alterations Leading to Communication / Navigation / Surveillance Degradation," July 2024, https://ad.easa.europa.eu/ad/2022-02R3. [Online]. Available: https://ad.easa.europa.eu/ad/2022-02R3

[3] S. Gebrekidan, "Electronic warfare confounds civilian pilots, far from any battlefield," The New York Times, Nov. 2023.

[4] G. S. Workgroup, "GPS spoofing: Final report of the GPS spoofing workgroup," OPSGROUP, Tech. Rep., 2024. [Online]. Available: https://ops.group/blog/gps-spoofing-final-report

[5] H. W. Jones, "The recent large reduction in space launch cost," in *48th International Conference on Environmental Systems (ICES)*. Albuquerque, New Mexico: ICES, July 2018, iCES-2018-81.

[6] P. A. Iannucci and T. E. Humphreys, "Fused low-Earth-orbit GNSS," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 60, no. 4, pp. 3730–3749, Jun. 2022.

[7] Z. Tan, H. Qin, L. Cong, and C. Zhao, "Positioning using IRIDIUM satellite signals of opportunity in weak signal environment," *Electronics*, vol. 9, no. 1, p. 37, 2019.

[8] M. Lei, Y. Liu, M. Gao, Z. Fang, J. Chen, and Y. Xu, "Robust Helmert variance component estimation for positioning with dual-constellation LEO satellites' signals of opportunity," *Electronics*, vol. 14, no. 17, p. 3437, 2025.

[9] Z. M. Kassas, N. Khairallah, and S. Kozhaya, "Ad astra: Simultaneous tracking and navigation with megaconstellation LEO satellites," *IEEE Aerospace and Electronic Systems Magazine*, vol. 39, no. 9, pp. 46–71, Jan. 2024.

[10] S. Kozhaya, J. Saroufim, and Z. M. Kassas, "Unveiling Starlink for PNT," *NAVIGATION*, vol. 72, no. 1, Mar. 2025.

[11] N. Jardak and R. Adam, "Practical use of Starlink downlink tones for positioning," *Sensors*, vol. 23, no. 6, p. 3234, Mar. 2023.

[12] M. Neinavaie and Z. M. Kassas, "Cognitive sensing and navigation with unknown OFDM signals with application to terrestrial 5G and Starlink LEO satellites," *IEEE Journal on Selected Areas in Communications*, vol. 42, no. 1, pp. 146–160, Jan. 2024.

[13] E. Grayver, R. Nelson, E. McDonald, E. Sorensen, and S. Romano, "Position and navigation using Starlink," in *Proceedings of the 2024 IEEE Aerospace Conference*. Big Sky, MT, USA: IEEE, Mar. 2024, pp. 1–12.

[14] S. C. Morgan, Z. M. Komodromos, W. Qin, Z. L. Clements, A. M. Graff, W. J. Morrison, and T. E. Humphreys, "A mock implementation of fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.

[15] Z. M. Komodromos, S. C. Morgan, Z. L. Clements, W. Qin, W. J. Morrison, and T. E. Humphreys, "Network-aided pseudorange-based LEO PNT from OneWeb," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.

[16] SpaceX, "Starlink Mini specifications," 2025, accessed: 2026-05-08. [Online]. Available: https://starlink.com/public-files/specification_sheet_mini.pdf

[17] M. Neinavaie, J. Khalife, and Z. M. Kassas, "Acquisition, Doppler tracking, and positioning with Starlink LEO satellites: First results," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 3, pp. 2606–2610, 2022.

[18] S. E. Kozhaya and Z. M. Kassas, "Positioning with Starlink LEO satellites: A blind Doppler spectral approach," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023, pp. 1–5.

[19] M. L. Psiaki, "Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids," *NAVIGATION*, vol. 68, no. 3, pp. 621–641, 2021.

[20] B. McLemore and M. L. Psiaki, "Navigation using Doppler shift from LEO constellations and INS data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4295–4314, 2022.

[21] A. Baron, P. Gurfil, and H. Rotstein, "Implementation and accuracy of Doppler navigation with LEO satellites," *NAVIGATION*, vol. 71, no. 2, Jun. 2024.

[22] R. C. Johnson, Ed., *Antenna Engineering Handbook*, 3rd ed. New York, NY, USA: McGraw-Hill, 1993.

[23] C. A. Balanis, *Antenna Theory: Analysis and Design*, 3rd ed. Hoboken, NJ, USA: John Wiley & Sons, 2005.

[24] C. Ruf, J. Perez, and S. White, "EVLA memo 222: Assessment of low Earth orbit broadband constellation downlink power flux density at Earth," National Radio Astronomy Observatory (NRAO), Socorro, NM, USA, Tech. Rep., Aug 2023. [Online]. Available: https://library.nrao.edu/public/memos/evla/EVLAM_222.pdf

[25] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 5, pp. 6016–6030, Oct. 2023.

[26] W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, "Timing properties of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 62, pp. 727–744, 2026.

[27] A. M. Graff and T. E. Humphreys, "OFDM-based positioning with unknown data payloads: Bounds and applications to LEO PNT," *IEEE Transactions on Wireless Communications*, vol. 25, pp. 7505–7519, 2026.

[28] M. S. McCormick, T. Larsson, A. Tamrat, J. C. Strait, C. Chen, P. J. Treigherman, S. O. Roche, and R. Gottula, "Configurable orthogonal frequency division multiplexing (OFDM) signal and transmitter and receiver for user terminal to satellite uplink communications," June 4 2024, US Patent 12,003,350 B1.

[29] S. Kozhaya, H. Kanj, and Z. M. Kassas, "Multi-constellation blind beacon estimation, Doppler tracking, and opportunistic positioning with OneWeb, Starlink, Iridium NEXT, and Orbcomm LEO satellites," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2023, pp. 1184–1195.

[30] A. M. Graff and T. E. Humphreys, "Ziv-Zakai-optimal OFDM resource alocation for time-of-arrival estimation," *IEEE Transactions on Wireless Communications*, vol. 24, no. 8, pp. 6886–6901, 2025.

[31] SpaceX Services, Inc., "Earth station authorization, file no. SES-MOD-20260227-00740," Federal Communications Commission, April 2026, call Sign E230035, granted Apr. 15, 2026. [Online]. Available: https://fccprod.servicenowservices.com/icfs?id=ibfs_application_summary&number=SES-MOD-20260227-00740

[32] Federal Communications Commission, "Order and authorization, DA 26-36 (Space Exploration Holdings, LLC)," FCC Space Bureau, Jan. 2026, released Jan. 9, 2026. [Online]. Available: https://docs.fcc.gov/public/attachments/DA-26-36A1.pdf

[33] D. Rife and R. Boorstyn, "Single tone parameter estimation from discrete-time observations," *IEEE Transactions on information theory*, vol. 20, no. 5, pp. 591–598, 1974.

[34] W. Qin, Z. M. Komodromos, S. C. Morgan, and T. E. Humphreys, "Maximum likelihood time of arrival and Doppler estimation for precise Starlink-based PNT," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.

[35] M. L. Psiaki, "The why and the how of research into LEO-based signals of opportunity," in *Proceedings of the ION GNSS+ Meeting*, Baltimore, MD, Sept. 2025.

Received: 00 Month 0000 | Revised: 00 Month 0000 | Accepted: 00 Month 0000

DOI: 10.1109/xxx.xx.xxxx

REGULAR PAPERS

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

Zachary L. Clements*[1] | Argyris Kriezis[2] | Todd E. Humphreys[1]

[1]Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Texas, USA

[2]Department of Aeronautics and Astronautics, Stanford University

Correspondence

*Zachary L. Clements, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, TX 78712
Email: zclements@utexas.edu

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.

Keywords

GNSS interference monitoring, interference detection, interference localization

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 (Humphreys, 2012, 2017; Humphreys et al., 2008; Scott, 2003). The past five years have seen a crescendo of GNSS disruptions in the aviation and maritime sectors (GPS Spoofing Workgroup, 2024; Marcos et al., 2018; Osechas et al., 2022; 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 (Borio et al., 2012; Clements et al., 2022; Psiaki & Humphreys, 2016, 2020). It has been shown that dedicated GNSS interference monitoring systems provide valuable spectrum situational awareness to enhance navigation security (Gunawardena & Van Graas, 2011; Gunawardena et al., 2009; Kriezis et al., 2024; Morrison et al., 2023; Sokolova et al., 2022; Stader & Gunawardena, 2021), 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 (Clements et al., 2023, 2026; Murrian et al., 2021).

Networks of terrestrial GNSS receivers can also be enlisted for GNSS interference monitoring (Abraha et al., 2024; Borio & Gioia, 2015; Jada et al., 2026), including of space-based interference. York et al. (2014) revealed that GPS PRNs 24 and 27

Licensed under CC-BY 4.0 © 2023 Institute of Navigation

(SVN65 and SVN66) were transmitting leakage tones at $\pm$10 MHz and $\pm$20 MHz from the GPS L5 frequency. The anomaly was subsequently fixed. More recently, Patil et al. (2023, 2024) used a network of terrestrial reference stations to identify a BeiDou satellite (NORAD ID 40749) as a source of interference in the B3I band (1268.52 MHz). The offending satellite continuously transmitted a tone centered at 1268.52 MHz and tones at $\pm$10 MHz and $\pm$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 by 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 by Patil et al. (2023, 2024) and 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; Johnston et al., 2017; Kouba, 2009; Montenbruck 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 (GLRT) 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 by 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

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

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

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

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

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 jth 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

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

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 $I_0^i = 10\log_{10}(\tilde{I}_0^i)$. 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 $S_{\text{I}}^i(f)$ be its power spectral density (PSD, dBW/Hz). Then I_0^i may be defined as

$$I_0^i = S_{\text{I}}^i(0) \qquad \text{(dBW/Hz)} \tag{5}$$

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

$$\tilde{S}_{\text{I}}^i(f) = \tilde{S}_{\text{C}}^i(f) \star \tilde{S}_{r_{\text{I}}}^i(f + \hat{f}_{\text{D}}^i) \tag{6}$$

Similar to (2), let P_{I}^i denote the received interference signal's power at the ith station, given by

$$P_{\text{I}}^i = P_{\text{I}} + G_{\text{I}}\left(\theta_{\text{I}}^i, \phi_{\text{I}}^i\right) + G_{\text{R}}^i\left(\theta_{\text{R}}^{i\text{I}}, \phi_{\text{R}}^{i\text{I}}\right) + L_{i\text{I}} \qquad \text{(dBW)} \tag{7}$$

where P_{I} is the transmitted interference power and G_{I} is the interference source antenna's gain function. The arguments of G_{I} and G_{R}^i, and the quantity $L_{i\text{I}}$ are as before except that they apply to the interference source rather than to the jth GNSS signal.

At the kth time epoch, the receiver at the ith station reports the measurement $z_{ij}[k]$, which is either the measured CNR_{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

$$H_0: \quad z_{ij}[k] = \text{CNR}_{ij}[k] + w_{ij}[k] \tag{8}$$

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

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. CNR-based GNSS interference detectors have been presented by Borio and Gioia (2015) and Jada et al. (2026). Borio and Gioia (2015) present a method based on the correlation introduced on CNR measurements from interference. Jada et al. (2026) developed an autonomous GNSS interference detection algorithm using the time-difference between CNR measurements from continuously operating reference stations (CORS), similar to the detector presented by Clements and Humphreys (2025). The problem in this paper is to identify a *transient* interference event that typically persists for 3 to 5 seconds, thus, the optimal detector requires a different construction.

Detection techniques from the nascent transient change detection literature are well suited to this problem (Egea-Roca et al., 2018; Guépié et al., 2012; Moustakides, 2014). 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 by 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 $l = \lceil T_{\text{max}}/\Delta t \rceil + 1 \in \mathbb{N}$ as the detector's index stride. Then for the kth time-epoch, ith station, and jth tracked GNSS signal, let the signal-specific detection statistic be given by

$$\xi_{ij}[k] = \tfrac{1}{2}\left(z_{ij}[k+l] - 2z_{ij}[k] + z_{ij}[k-l]\right) \tag{10}$$

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

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

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

where the mean μ_i is

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

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 $\xi_{ij}[k]$:

$$\Lambda_i[k] = \frac{1}{|\mathcal{J}_i[k]|}\sum_{j\in\mathcal{J}_i[k]}\left(\xi_{ij}[k]\right. \tag{14}$$

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

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

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

The quantity μ_i amounts to the average drop in observed CNR at the ith 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 ith 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 ith station then becomes

$$\Lambda_i[k] \underset{H_0}{\overset{H_1}{\gtrless}} \nu_i \tag{17}$$

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.

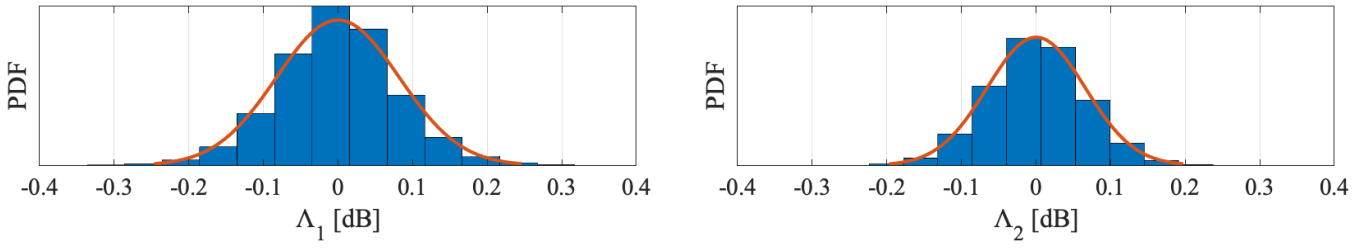


FIGURE 1 The distributions of Λ_i under H_0 (no interference present) for stations METG ($i = 1$) and MATE ($i = 2$).

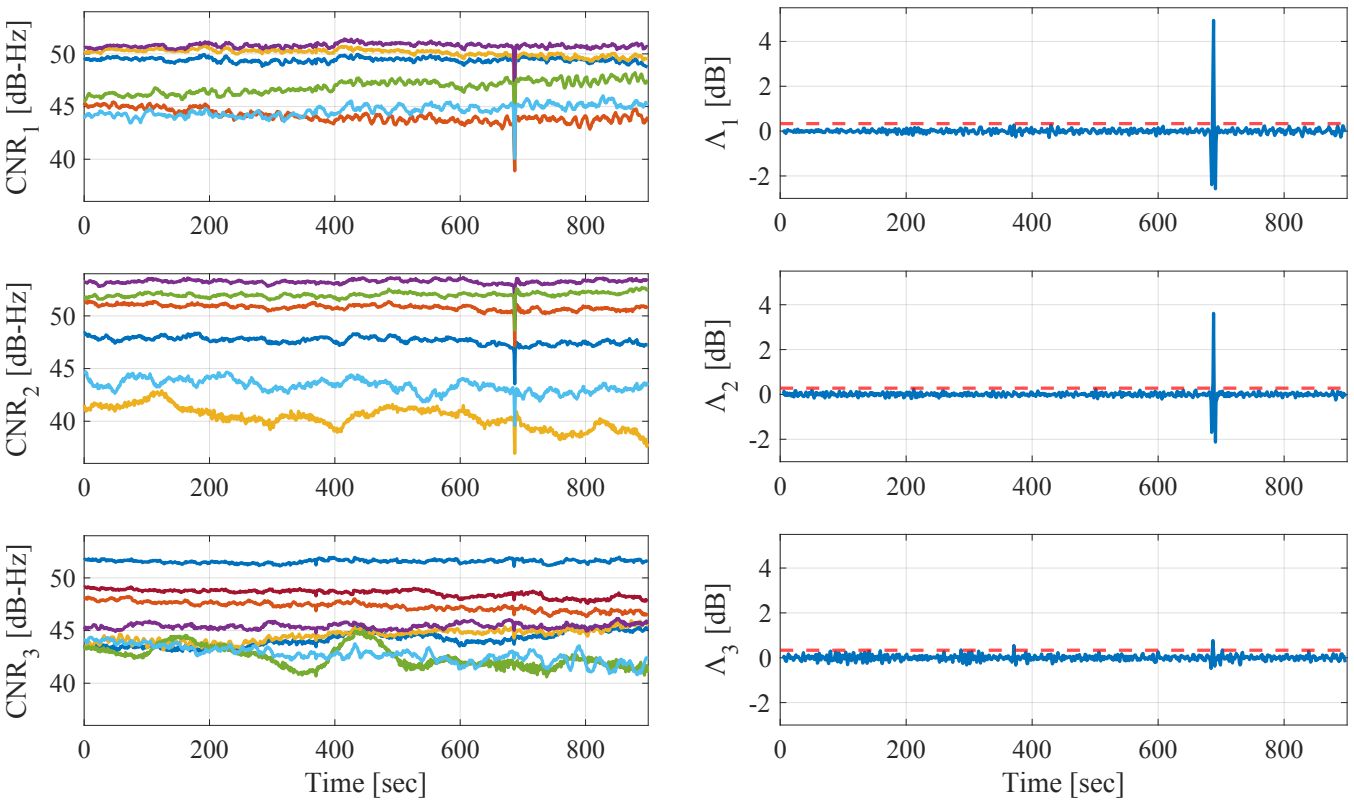


FIGURE 2 Reported CNR (left) and the detection statistic Λ_i (right) for IGS stations METG ($i = 1$, Finland), MATE ($i = 2$, Italy), and THU2 ($i = 3$, Greenland) over a 15-minute interval on day 160 of year 2021 for the received GPS L1 C/A signals. The dashed red line is the detection threshold with 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

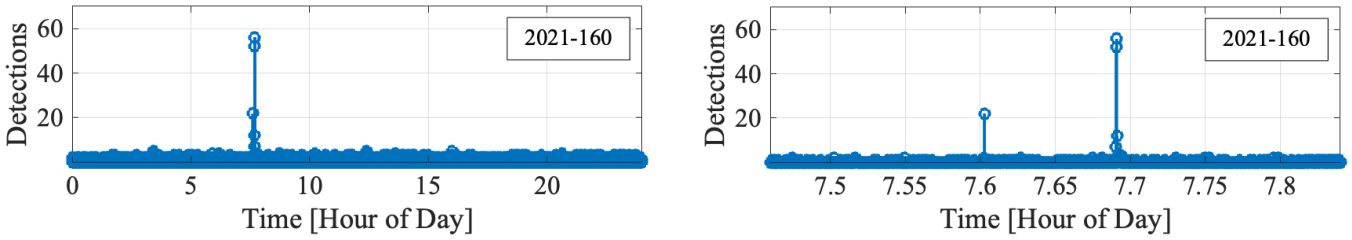


FIGURE 3 The number of stations that detected interference during day 160 of year 2021. The expanded view on the right shows a lower-power event detected by 21 stations, followed by a higher-power event detected by 58 stations.

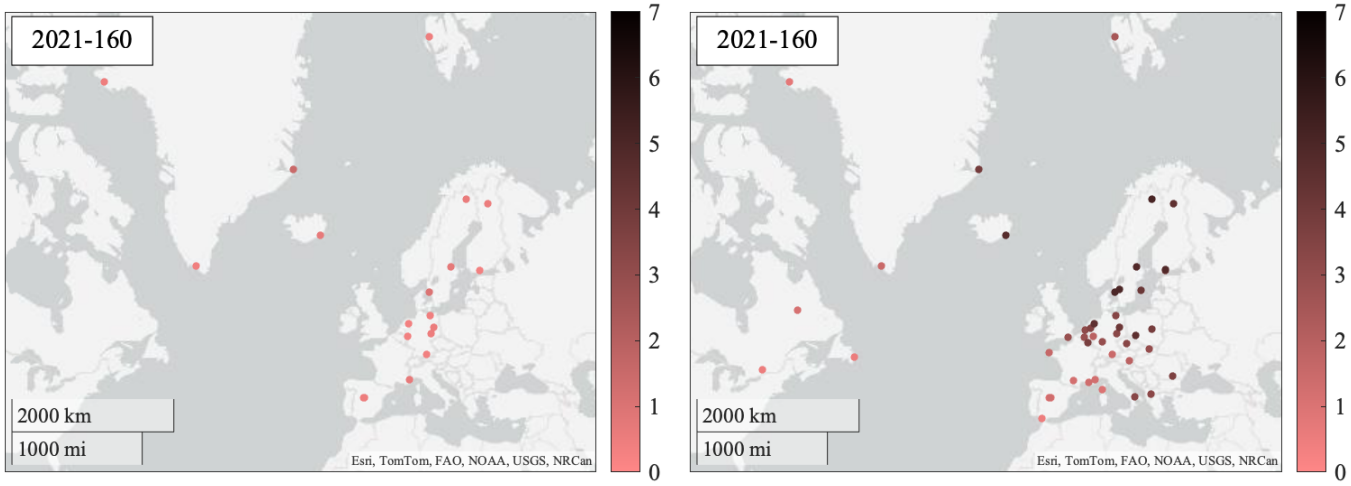


FIGURE 4 Heat map of the test statistic (in dB) at triggering stations for the lower-power event (left) and the higher-power event (right). The drop in GPS L1 C/A CNR during the more powerful interference event was as large as 6 dB, centered near the Baltic region.

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.

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.

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

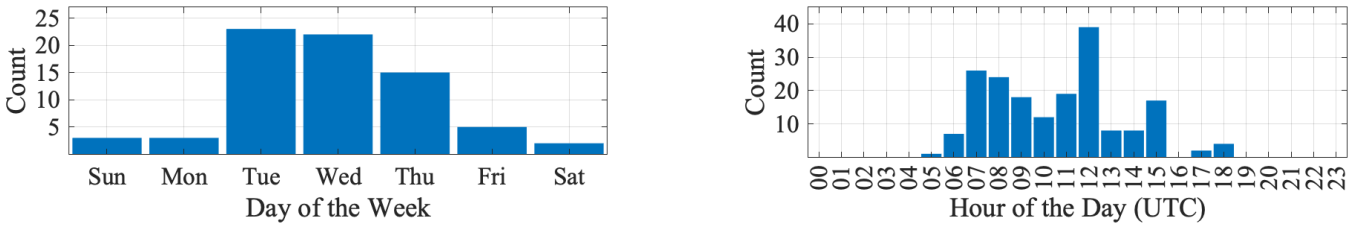


FIGURE 5 Distribution of the day of the week and hour of day (with respect to UTC) during which interference events with at least one station suffering a drop of 5 dB or greater occurred. Clearly, the high-power interference events typically occur during business days and business hours.

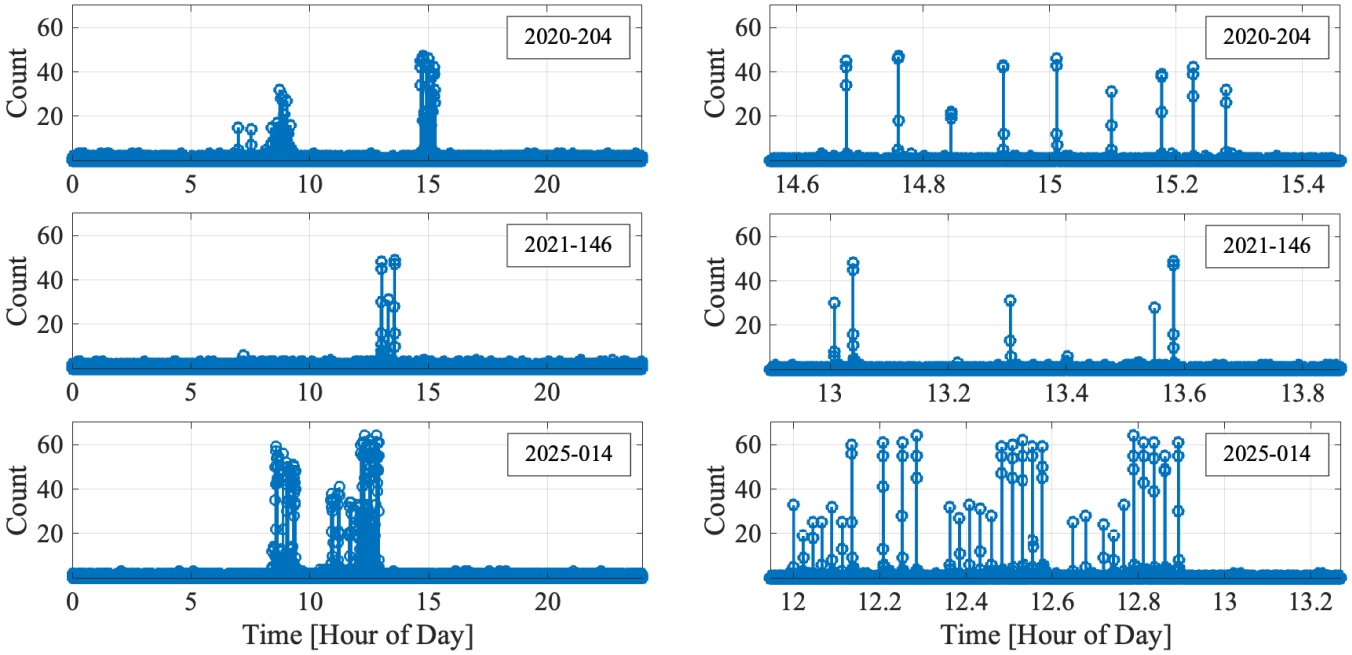


FIGURE 6 The number of stations that detected interference during three different days. Interference events can occur several times a day. The left column shows the entire day, while the right column is an expanded sub-interval.

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.

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

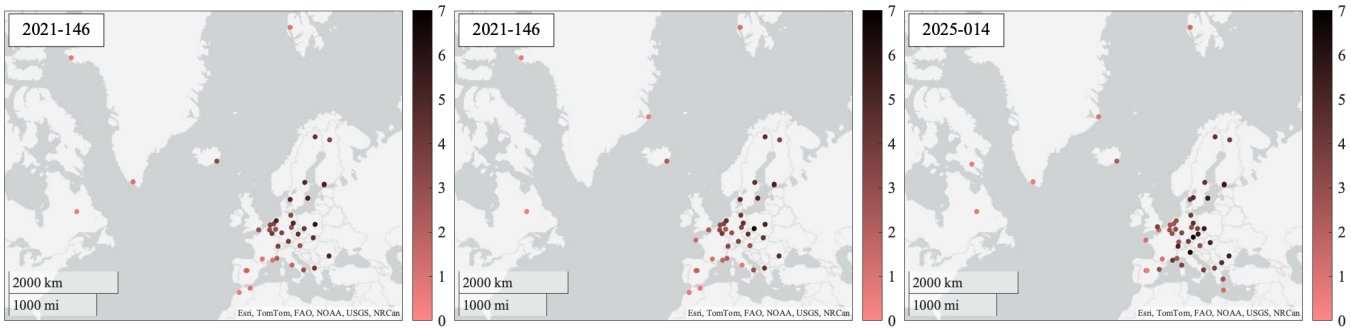


FIGURE 7 Detection statistic heat maps, in dB, for day 146 of year 2021 and day 014 of year 2025. GNSS receivers in the Baltic region are impacted most, which is generally representative of nearly all of the interference events.

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.

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 & 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 $\tilde{S}_{\mathrm{C}}^{i}(f)$ which, as shown in (6), leads to a different interference component PSD $\tilde{S}_{\mathrm{I}}^{i}(f)$ and ultimately to a different interference power density I_0^i.

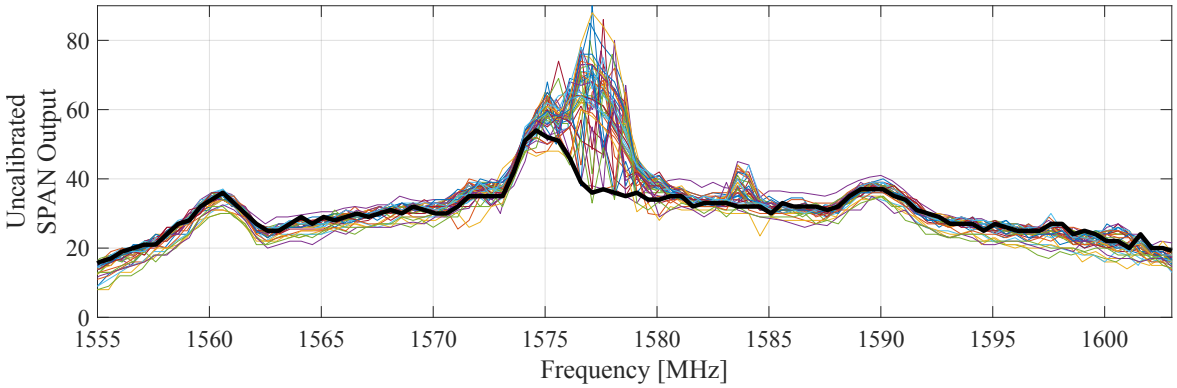


FIGURE 8 Uncalibrated average PSD near the GPS L1 band during nominal operation (thick black line) and 1-Hz PSD estimates during 48 transient interference events, as recorded in Gdynia, Poland.

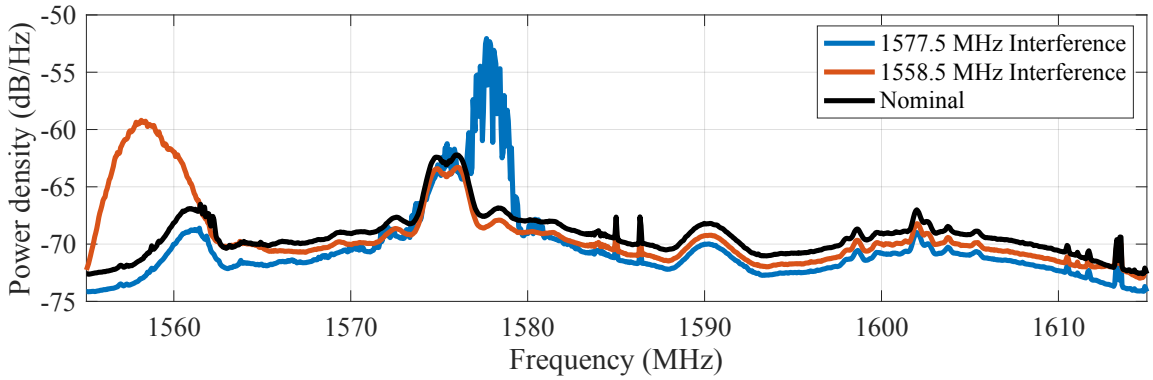


FIGURE 9 Power spectrum derived from raw wideband samples captured in Amsterdam, Netherlands, during an interference event on February 11, 2026. An initial burst of interference at 1577.5 MHz was followed by a burst at 1558.5 MHz.

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 $I_0^i = S_{\mathrm{I}}^i(0) = 0$. 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.

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.

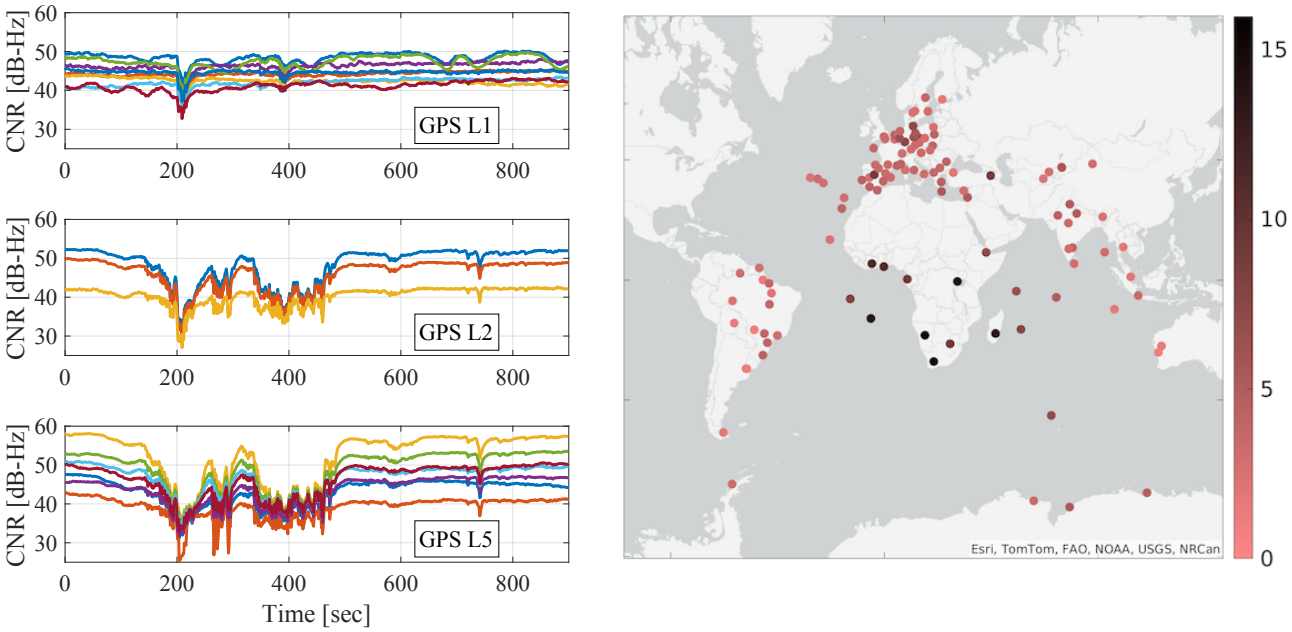


FIGURE 10 Left: CNR time history of tracked GPS L1, L2, and L5 signals produced by IGS station SUTM in South Africa during the solar radio burst on November 11, 2025. Right: Heat map of maximum CNR degradation at GPS L2 during the solar radio burst.

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. Due to the intensity of the solar radio burst, detections were flagged by the developed CNR-based framework despite being tailored to transient interference. 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.

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 $\mathcal{I}_\mathrm{d} \subseteq \mathcal{I}$ 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 $i \in \mathcal{I}_\mathrm{d}$ 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 $i \in \mathcal{I}_\mathrm{d}$.

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

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

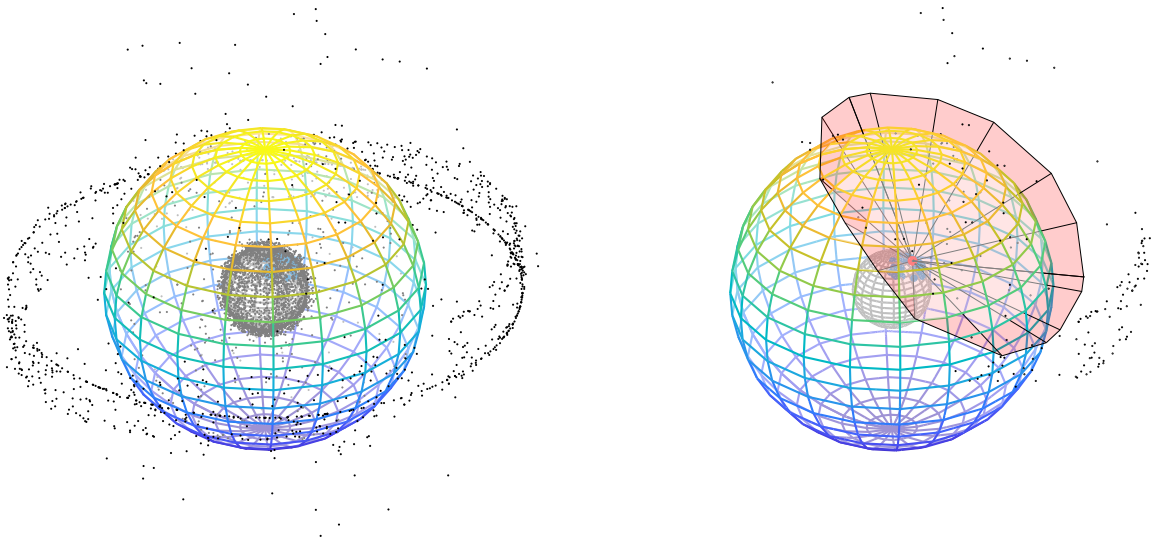


FIGURE 11 Left: The position of all tracked objects during the high-power interference burst on day 160 year 2021. Right: The position of all satellites that satisfy a 0° elevation mask, excluding debris and rocket bodies. The feasible region in which the interference source could have been positioned 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

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 $\zeta^* = \|\boldsymbol{r}^*\| - r_{\mathrm{E}}$, 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 $\boldsymbol{r} \in \mathbb{R}^3$ 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.

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.

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 $\boldsymbol{x}$ is estimated from the measurements $\boldsymbol{z}$ via maximum likelihood estimation under each hypothesis. Second, these estimates are used to evaluate likelihoods as if they were the correct values of $\boldsymbol{x}$ 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 $\boldsymbol{x}_s$ 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 ith station and jth GNSS signal at epoch k, let $z_{ij}[k]$ be an interference-free CNR measurement, modeled as in (8). The measurement noise variance $\sigma^2_{ij}[k] = f(\theta^{ij}_{\mathrm{R}}[k])$ is modeled as a nonlinear function of the received zenith angle, with $\sigma_{ij} = 0.25$ dB at zenith ($\theta^{ij}_{\mathrm{R}}[k] = 0$) and $\sigma_{ij} = 1.25$ dB at the horizon ($\theta^{ij}_{\mathrm{R}}[k] = 90°$). Both receiver and transmitter antenna gain patterns are modeled as azimuthally symmetric—$G^i_{\mathrm{R}}(\theta, \phi) = G^i_{\mathrm{R}}(\theta)$, $G^j_{\mathrm{T}}(\theta, \phi) = G^j_{\mathrm{T}}(\theta)$—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 jth GNSS signal, and $\mathcal{S}_{\mathrm{e}}[k]$ be the set of candidate source satellites that have all passed elevation mask vetting. Let the measurement vector $\boldsymbol{z}_j[k]$ 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 $s \in \mathcal{S}_{\mathrm{e}}[k]$ is consistent with $\boldsymbol{z}_j[k]$ given varying degrees of prior knowledge about the satellite's transmitter. The unknown parameter vector $\boldsymbol{x}_s[k]$ may include any combination of transmit power $P_s[k]$, beamwidth $\beta_s[k]$, and beam vector $\boldsymbol{v}_s[k] = [\theta_s[k],\ \phi_s[k]]^{\mathsf{T}}$, 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, $\boldsymbol{x}_s$ includes all three parameters:

$$\boldsymbol{x}_s[k] = [P_s[k]\,,\ \beta_s[k]\,,\ \boldsymbol{v}^{\mathsf{T}}_s[k]]^{\mathsf{T}} \tag{19}$$

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

$$H_s: \quad \boldsymbol{z}_j[k] \sim \mathcal{N}\left(\bar{\boldsymbol{z}}_j[k, \boldsymbol{x}_s],\ R_j[k, \boldsymbol{x}_s]\right), \quad s \in \mathcal{S}_{\mathrm{e}}[k] \tag{20}$$

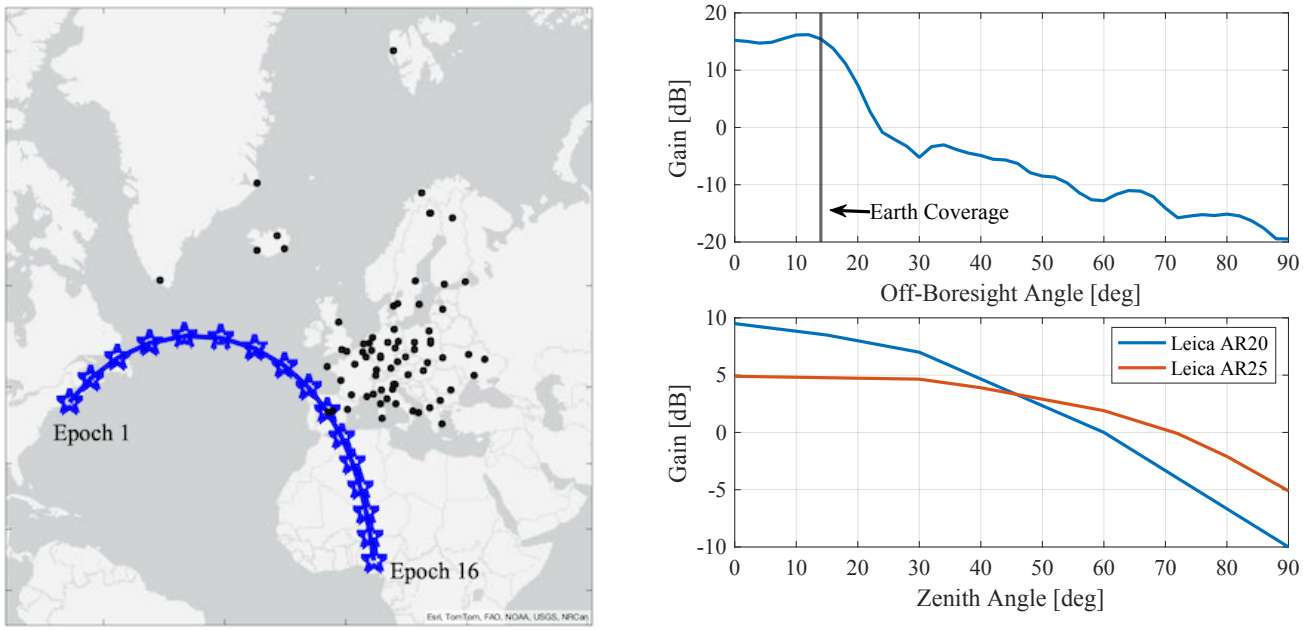


FIGURE 12 Simulation study setup. Left: The geometry of the true source GNSS satellite over 16 epochs at 15-minute intervals, with terrestrial reference stations shown as black dots. Right: Transmitter (top) and receiver (bottom) gain patterns. The receiver gain patterns are from a model provided by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

Here, $\bar{\boldsymbol{z}}_j[k, \boldsymbol{x}_s]$ is the vector of expected CNR measurements for signal $j \in \mathcal{J}$ emanating from satellite $s \in \mathcal{S}_\mathrm{e}[k]$, and $R_j[k, \boldsymbol{x}_s]$ is a square diagonal matrix with diagonal elements $\sigma^2_{ij}[k]$ for $i \in \mathcal{I}_j[k]$, where the angle $\theta^{ij}_\mathrm{R}[k]$, on which $\sigma^2_{ij(s)}[k]$ 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 $\boldsymbol{x}_s$, which, for the model in (20), leads to a cost function of the form

$$J_j[k, \boldsymbol{x}^*_s] = \left(\boldsymbol{z}_j[k] - \bar{\boldsymbol{z}}_j[k, \boldsymbol{x}^*_s]\right)^\mathsf{T} \left(R_j[k, \boldsymbol{x}^*_s]\right)^{-1} \left(\boldsymbol{z}_j[k] - \bar{\boldsymbol{z}}_j[k, \boldsymbol{x}^*_s]\right), \quad s \in \mathcal{S}_\mathrm{e}[k] \tag{21}$$

where $\boldsymbol{x}^*_s$ is the value of $\boldsymbol{x}_s$ that minimizes $J_j[k, \boldsymbol{x}_s]$. Thus, for $a, s \in \mathcal{S}_\mathrm{e}[k]$, if $J_j[k, \boldsymbol{x}^*_a] < J_j[k, \boldsymbol{x}^*_s]$, then H_a is more likely than H_s. Since $\mathcal{S}_\mathrm{e}[k]$ 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)

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

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

$$J_j[k, \boldsymbol{x}^*_u] \sim \chi^2_{|\mathcal{I}_j[k]|} \tag{23}$$

From this one can readily calculate $\nu_j[k]$ to ensure $P\left(J_j[k, \boldsymbol{x}^*_u] > \nu_j[k]\right)$ 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 $\mathcal{S}_\mathrm{e}[k]$ of candidate satellites was generated by applying

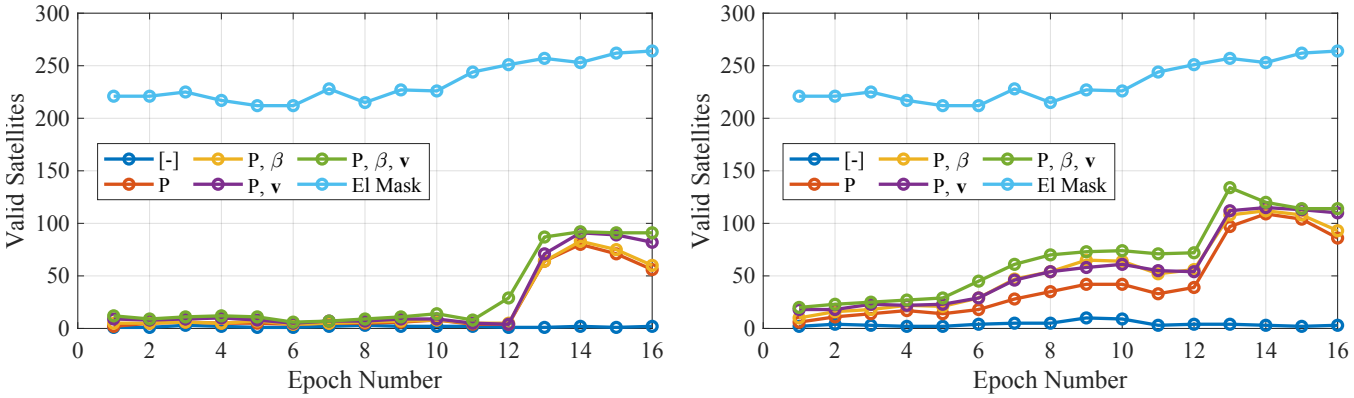


FIGURE 13 The number of satellites that, for at least 5% of the Monte Carlo trials, remained viable candidates after association testing with $P\left(J_j[k, \boldsymbol{x}_s^*] > \nu_j[k]\right) < 10^{-3}$. The legends indicate which parameters are assumed unknown. The number of satellites satisfying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

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 $\boldsymbol{x}_s$ under the model (1); (4) a realization of the measurement $z_{ij}[k]$ was generated according to (8) assuming $\sigma_{ij}^2[k] = f(\theta_{\mathrm{R}}^{ij}[k])$ for each $i \in \mathcal{I}_j[k]$.

Costs $J_j[k, \boldsymbol{x}_s^*]$ were calculated for all $s \in \mathcal{S}_{\mathrm{e}}[k]$, with $\boldsymbol{x}_s$ 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 $\boldsymbol{v}_s$; and (S5) unknown P_s, β_s, and $\boldsymbol{v}_s$. The maximum likelihood (minimum cost) estimates β_s^*, and $\boldsymbol{v}_s^*$ were found via grid search, while P_s^* was taken as the mean of the elements of the difference $\boldsymbol{z}_j[k] - \bar{\boldsymbol{z}}_j[k, \tilde{\boldsymbol{x}}_s^*]$, with $\tilde{\boldsymbol{x}}_s^* = [P_s = 0\,,\, \beta_s^*\,,\, \boldsymbol{v}_s^{*\mathsf{T}}]^{\mathsf{T}}$.

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 $J_j[k, \boldsymbol{x}_s^*]$ below $\nu_j[k]$ when $P\left(J_j[k, \boldsymbol{x}_u^*] > \nu_j[k]\right)$ 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 $\mathcal{J}_{\mathrm{g}} \subset \mathcal{J}$ be the subset of GPS L1 C/A signals. Based on historical data from

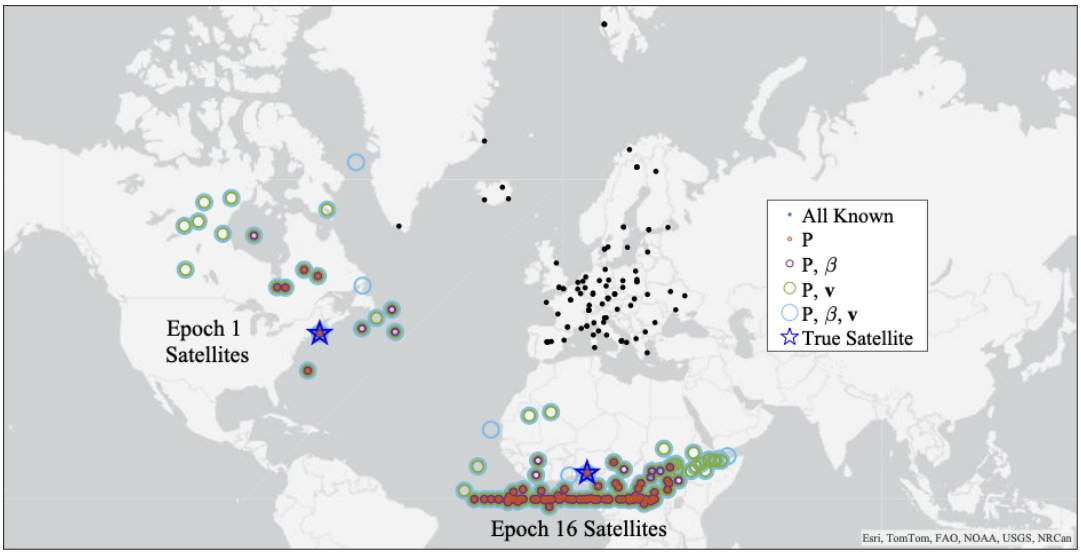


FIGURE 14 Satellites that remained viable candidates on more than 5% of Monte Carlo trials for epoch $k = 1$ and epoch $k = 16$ when the reference stations are modeled as using the Leica AR25 antenna. The legend gives symbols for the true satellite and for valid candidates under scenarios with an increasing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

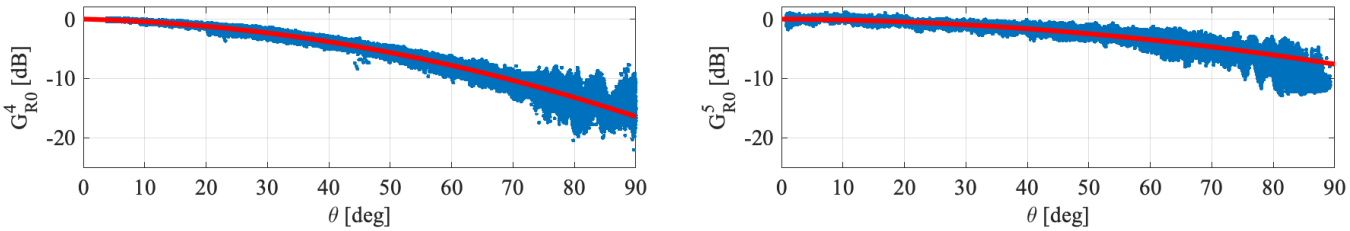


FIGURE 15 Normalized estimates of $G^i_{\text{R0}}(\theta)$ for stations MEDI (Leica AR20, $i = 4$) and NICO (Leica AR25, $i = 5$). The blue dots represent collected values of $\text{CNR}_{ij} - G^j_{\text{T}} - L_{ij}$, while the red line indicates a weighted third-order polynomial fit, which becomes $G^i_{\text{R0}}(\theta)$, all normalized to $G^i_{\text{R0}}(0) = 0$. As expected, measurement noise increases significantly as the elevation angle decreases.

each site $i \in \mathcal{I}$ and on the known azimuthally symmetric antenna gain pattern $G^j_{\text{T}}(\theta)$ and transmit power P^j_{T} for $j \in \mathcal{J}_{\text{g}}$ from a select number of GPS satellites $\mathcal{S}_{\text{g}} \subset \mathcal{S}_{\text{e}}$, empirical estimates of the combined receiver antenna gain pattern and receiver noise floor

$$G^i_{\text{R0}}(\theta) = G^i_{\text{R}}(\theta) - N^i_0 \tag{24}$$

were obtained for all $i \in \mathcal{I}$ across zenith angles $\theta \in [0, 90°]$, assuming azimuthal symmetry. Accurate estimates of $G^j_{\text{T}}(\theta)$ for $j \in \mathcal{J}_{\text{g}}$ and $s \in \mathcal{S}_{\text{g}}$ are available in public documents (Marquis & Reigh, 2015), and corresponding estimates for P^j_{T} were obtained by observation with a calibrated high-gain antenna. Normalized estimates of $G^i_{\text{R0}}(\theta)$ 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 $j \in \mathcal{J}_{\text{g}}$ for $s \in \mathcal{S}_{\text{g}}$, according to (8) with $\sigma^2_{ij} = f_i(\theta^{ij}_{\text{R}})$. 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 $G^i_{\mathrm{R}}(\theta)$ and noise floor N^i_0 must be known independently, rather than combined as in (24), because the model for CINR_{ij} given in (4) involves an isolated value of $\tilde{N}^i_0$. Yet it is not possible to independently estimate $G^i_{\mathrm{R}}(\theta)$ and N^i_0 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^i_0 can be obtained for $\mathcal{I}_{\mathrm{a}} \in \mathcal{I}$, the subset of reference stations for which an accurate antenna gain pattern $G^i_{\mathrm{R}}(\theta)$ is available, e.g., one provided by the manufacturer. But the authors of the current paper found $\mathcal{I}_{\mathrm{a}}$ was limited to stations with Leica AR20 or AR25 antennas. For all other stations (all $i \in \mathcal{I} \setminus \mathcal{I}_{\mathrm{a}}$), either $G^i_{\mathrm{R}}(\theta)$ was not available or the available model did not match the shape of the empirically derived combined model $G^i_{\mathrm{R0}}(\theta)$. (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 $\mathcal{I}_{\mathrm{a}}$ 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 $G^i_{\mathrm{R}}(\theta)$ 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 $G^i_{\mathrm{R}}(\theta)$ is known, and leveraging the known transmit gain pattern $G^j_{\mathrm{T}}(\theta)$ and transmit power P^j_{T} for $j \in \mathcal{J}_{\mathrm{g}}$, N^i_0 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 $\mathcal{I}_{\mathrm{d}}[k] \subset \mathcal{I}$ be the subset of stations that detect the event, and let the vector $\boldsymbol{z}[k]$ contain ordered station-specific detection statistics at epoch k, i.e., values from the set $\{\Lambda_i[k] \mid i \in \mathcal{I}_{\mathrm{d}}[k]\}$. For every satellite $s \in S[k]$ and for an assumed parameter vector $\boldsymbol{x}_s$, one can develop a hypothesis model for $\boldsymbol{z}[k]$ like the one in (20) and find the subset of $S[k]$ consistent with $\boldsymbol{z}[k]$, just as with the GNSS satellite identification problem. Let the resulting probabilities, analogous to $P(s|\boldsymbol{z}_j[k])$ in (22), be denoted $P(s|\boldsymbol{z}[k])$. 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 $\boldsymbol{g}_1$ and $\boldsymbol{g}_2$ in the ECEF reference frame. Let $\mathcal{K} = \{1, \ldots, K\}$ be the set of TDOA measurement epochs over an observation interval, $\mathcal{S}_{\mathrm{e}}$ be the set of candidate satellites that satisfy elevation mask vetting for all $k \in \mathcal{K}$, and $\boldsymbol{r}_s[k]$ be the ECEF position of satellite $s \in \mathcal{S}_{\mathrm{e}}$ at epoch k. For model simplicity, $\boldsymbol{r}_s[k]$ is assumed to remain constant over the signal's time-of-flight for each measurement. Denoting by $\rho_{is}[k] = \|\boldsymbol{g}_i - \boldsymbol{r}_s[k]\|$ the range between receiver $i \in \{1, 2\}$ and satellite $s \in \mathcal{S}_{\mathrm{e}}$, 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 $u \in \mathcal{S}_{\mathrm{e}}$ is the true interference source. Then the actual TDOA measurement is modeled as

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

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 $\boldsymbol{y} = [y[1], \ldots, y[K]]^{\mathsf{T}}$, and modeled values for $s \in \mathcal{S}_{\mathrm{e}}$ as $\bar{\boldsymbol{y}}_s = [\bar{y}_s[1], \ldots, \bar{y}_s[K]]^{\mathsf{T}}$.

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

$$\boldsymbol{r}_s[k] = \tilde{\boldsymbol{r}}_s[k] + \boldsymbol{e}_s, \quad s \in \mathcal{S}_\text{e}, \;\; k \in \mathcal{K} \tag{26}$$

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

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

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

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

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

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 $s \in \mathcal{S}_\text{e}$ is consistent with the measurement $\boldsymbol{y}$. For each s, the elements of $\tilde{\boldsymbol{y}}_s$ can be readily calculated as in (28). Let $\boldsymbol{1}$ be the all-ones vector of appropriate size, and let $\boldsymbol{w} = [w[1], \ldots, w[K]]^\mathsf{T} \sim \mathcal{N}(\boldsymbol{0}, R)$, with $R = \sigma^2 I$. Then the TDOA measurement residual with respect to the TLE-modeled satellite s is

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

Note that if $s = u$, i.e., if s is the true satellite, then $\boldsymbol{\gamma}_s$ is zero mean. A cost can be assigned to each satellite $s \in \mathcal{S}_\text{e}$ 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 $p(\eta_s | \boldsymbol{y})$:

$$J_s(\eta_s) = (\boldsymbol{\gamma}_s - \eta_s \boldsymbol{1})^\mathsf{T} R^{-1} (\boldsymbol{\gamma}_s - \eta_s \boldsymbol{1}) + \left(\frac{\eta_s}{\sigma_s}\right)^2 \tag{30}$$

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

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

This distribution arises because the total number of independent measurements is $K + 1$ (K from $\boldsymbol{y}$ 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 5.4, 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 $P(s|\boldsymbol{z})$ from Section 5.4 as the prior probability for the TDOA-based identification problem. Then, since $\mathcal{S}_\text{e}$ is assumed to include the true interference satellite, and approximating σ_s as equal for all $s \in \mathcal{S}_\text{e}$, the probability that H_s is true is

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

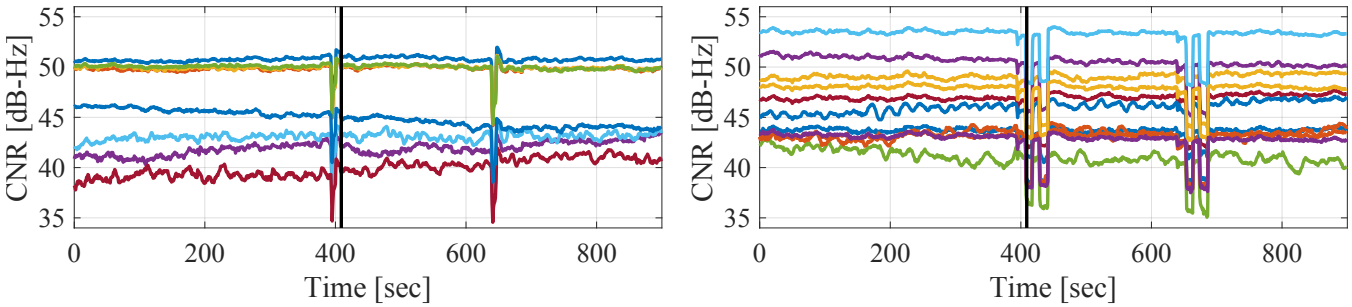


FIGURE 16 CNR time history of tracked GPS L1 C/A (left) and BeiDou B1I (right) signals at the IGS station METG in Finland. The GPS L1 C/A signals are first affected by each interference event, followed by the BeiDou B1I signals. The black line indicates the 2.3-second time-overlapped raw IQ capture interval.

Here, the vector $\boldsymbol{z}$ contains the ordered station-specific detection statistics for those stations $\mathcal{I}_\mathrm{d}$ that detect the interference event associated with the TDOA measurement, i.e., values from the set $\{\Lambda_i \mid i \in \mathcal{I}_\mathrm{d}\}$. As before, one can choose the highest-probability satellite, or the top N most probable, or all those for which $P(s|\boldsymbol{y}, \boldsymbol{z})$ 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 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 (Morrison et al., 2023; Sokolova et al., 2022). 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 us. 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 us, then later at 292.9 us.

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).

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

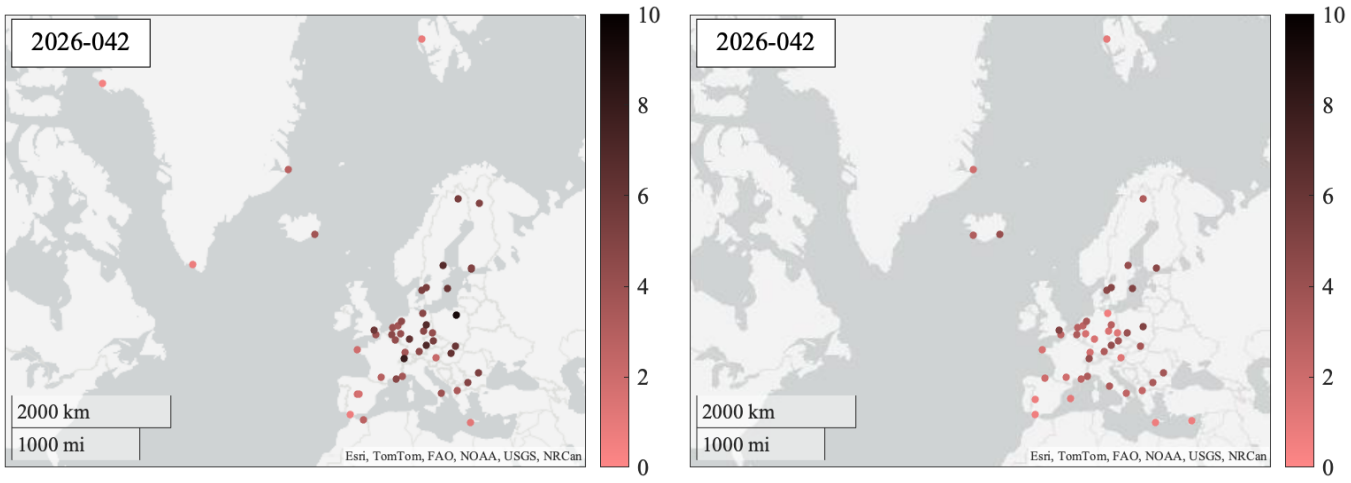


FIGURE 17 Interference heat maps for day 42 of year 2026 from tracked GPS L1 C/A (left) and BeiDou B1I (right) signals.

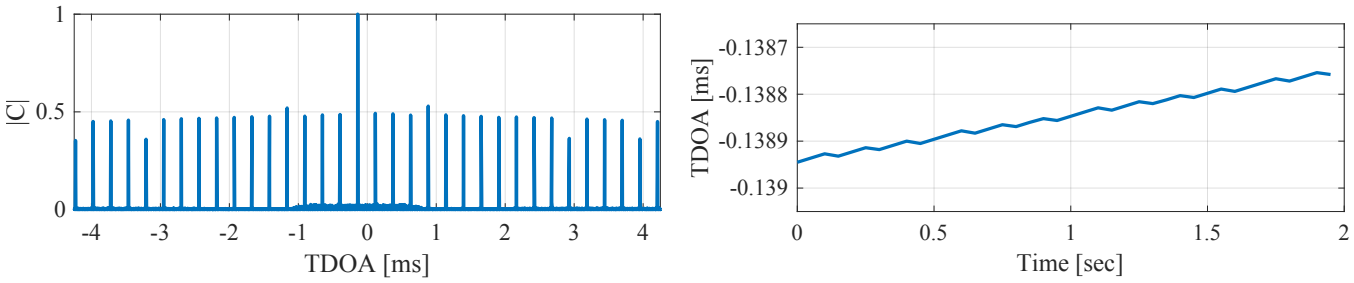


FIGURE 18 Left: Example normalized cross correlation at the optimal FDOA. Right: TDOA measurement time history.

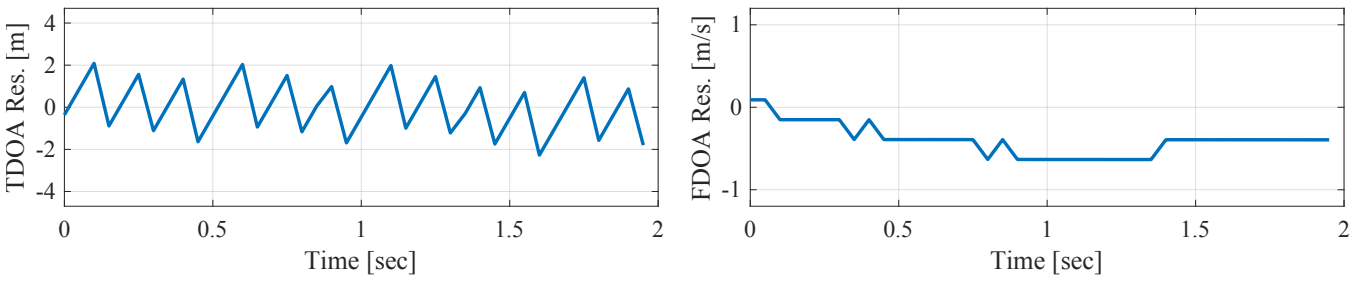


FIGURE 19 TDOA (left) and FDOA (right) residuals for Cosmos 2546 (NORAD ID 45608). The effects of nearest-sample quantization and Doppler search quantization are evident in the respective plots.

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 us 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 $\boldsymbol{y}$ 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 $s \in \mathcal{S}_\text{e}$. 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 $\sigma_\text{n}^2 = 0.76^2$ m^2 for 75-MHz sampling. Alternatively, a conservative choice of σ^2 can be determined by selecting the width of half a sample, $\sigma^2 = 2^2$ 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, $\sigma^2 = 2^2$ m.

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 $s \in \mathcal{S}_\text{e}$, 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 $\mathcal{S}_\text{e}$ was composed of all satellites satisfying elevation angle masking with $\alpha_0 = 0$ for all stations that detected the event.

For each satellite $s \in \mathcal{S}_\text{e}$, 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 $P(s|\boldsymbol{z}) = 1/|\mathcal{S}_\text{e}|$ for $s \in \mathcal{S}_\text{e}$, the probability $P(s|\boldsymbol{y}, \boldsymbol{z})$ calculated according to (32) for Cosmos 2546 was numerically indistinguishable from unity. Cosmos 2546 also satisfied an $\alpha_0 = 35°$ elevation mask with respect to every $i \in \mathcal{I}_\text{d}$. 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 $\mathcal{S}_\text{e}$, 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 $\mathcal{L} = \{1, \ldots, L\}$ be the index set for this TLE archive, and let t_l represent the epoch for the lth TLE. Let $\boldsymbol{r}_k(t_l)$ denote the position of Cosmos 2546 based on the kth TLE but propagated to the epoch of the lth TLE. For all $l \in \mathcal{L}$ and all $k \in \mathcal{L} \setminus l$, one may assume that $\boldsymbol{r}_l(t_l)$ is more accurate than $\boldsymbol{r}_k(t_l)$. Let $\mathcal{L}_l \subset \mathcal{L}$ be the set of TLEs within a 60-hour window of t_l, and let $\mathcal{R}_l = \{\|r_k(t_l) - r_l(t_l)\| \mid k \in \mathcal{L}_l\}$ 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 (Griffiths & Ray, 2009; Hadas & Bosy, 2015).

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. Furthermore, as shown in Fig. **??**, the interference signal's power spectrum over two years remains identical, indicating that similar software/hardware configurations are being used. 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., & 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., & Kirubarajan, T. (2001). *Estimation with applications to tracking and navigation*. John Wiley; Sons. https://doi.org/10.1002/0471221279

Bhatti, J. A., Humphreys, T. E., & Ledvina, B. M. (2012). Development and demonstration of a TDOA-based GNSS interference signal localization system. *Proceedings of the IEEE/ION PLANS Meeting*, 1209–1220.

Borio, D., & Gioia, C. (2015). Real-time jamming detection using the sum-of-squares paradigm. *2015 International Conference on Localization and GNSS (ICL-GNSS)*, 1–6.

Borio, D., O'Driscoll, C., & Fortuny, J. (2012). GNSS jammers: Effects and countermeasures. *2012 6th ESA Workshop on Satellite Navigation Technologies (Navitec 2012) & European Workshop on GNSS Signals and Signal Processing*, 1–7.

Cerruti, A. P., Kintner, P. M., Gary, D. E., Lanzerotti, L. J., de Paula, E. R., & 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., & Coster, A. J. (2008). Effect of intense December 2006 solar radio bursts on GPS receivers. *Space Weather*, *6*(10).

Clements, Z. L., Ellis, P. B., & Humphreys, T. E. (2023). Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit. *Proceedings of the IEEE/ION PLANS Meeting*, 458–469. https://doi.org/10.1109/PLANS53410.2023.10140058

Clements, Z. L., Ellis, P. B., Murrian, M. J., Psiaki, M. L., & Humphreys, T. E. (2026). Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit. *NAVIGATION*, *73*(1). https://doi.org/10.33012/navi.750

Clements, Z. L., & Humphreys, T. E. (2025). Transient space-based GNSS interference: Observations and analysis. *Proceedings of the ION GNSS+ Meeting*, 71–81.

Clements, Z. L., Yoder, J. E., & Humphreys, T. E. (2022). Carrier-phase and IMU based GNSS spoofing detection for ground vehicles. *Proceedings of the ION International Technical Meeting*, 83–95. https://doi.org/10.33012/2022.18252

Dow, J. M., Neilan, R. E., & 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., & 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* (tech. rep.). OPSGROUP. https://ops.group/blog/gps-spoofing-final-report

Griffiths, J., & Ray, J. R. (2009). On the precision and accuracy of IGS orbits. *Journal of Geodesy*, *83*(3), 277–287.

Guépié, B. K., Fillatre, L., & Nikiforov, I. (2012). Sequential detection of transient changes. *Sequential Analysis*, *31*(4), 528–547.

Gunawardena, S., & Van Graas, F. (2011). Multi-channel wideband GPS anomalous event monitor. *Proceedings of the ION GNSS+ Meeting*, 1957–1968.

Gunawardena, S., Zhu, Z., de Haag, M. U., & van Graas, F. (2009). Remote-controlled, continuously operating GPS anomalous event monitor. *NAVIGATION*, *56*(2), 97–113.

Hadas, T., & 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* (pp. 469–503). Springer International Publishing. https://doi.org/10.1007/978-3-319-42928-1_16

Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O'Hanlon, B. W., & Kintner, P. M., Jr. (2008). Assessing the spoofing threat: Development of a portable GPS civilian spoofer. In *Proceedings of the ION GNSS meeting*.

Jada, S., Psiaki, M., & Joerger, M. (2026). Network-based GNSS jamming prediction enabling wideband interference observation. *NAVIGATION*, *73*(1).

Johnston, G., Riddell, A., & Hausler, G. (2017). The International GNSS Service. In *Springer handbook of global navigation satellite systems* (pp. 967–982). Springer International Publishing. https://doi.org/10.1007/978-3-319-42928-1_33

Komodromos, Z. M., Morgan, S. C., Clements, Z. L., Qin, W., Morrison, W. J., & Humphreys, T. E. (2025). Network-aided pseudorange-based LEO PNT from OneWeb. *Proceedings of the IEEE/ION PLANS Meeting*.

Kouba, J. (2009). A guide to using international GNSS service (IGS) products.

Kriezis, A., Chen, Y.-H., Akos, D., Lo, S., & Walter, T. (2024). GNSS RFI detection and impact characterization in various interference environments using low-cost receivers. *Proceedings of the ION GNSS+ Meeting*, 3348–3360.

Kriezis, A., Chen, Y.-H., Akos, D., Lo, S., & Walter, T. (2025). GNSS jamming and spoofing monitoring using low-cost COTS receivers [Submitted for review.]. *NAVIGATION*.

Krzan, G., Dawidowicz, K., & 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., & Meurer, M. (2018). Interference awareness and characterization for GNSS maritime applications. *Proceedings of the IEEE/ION PLANS Meeting*, 908–919. https://doi.org/10.1109/plans.2018.8373469

Marquis, W. A., & 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., & 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., & Humphreys, T. E. (2025). A mock implementation of fused LEO GNSS. *Proceedings of the IEEE/ION PLANS Meeting*.

Morrison, A., Sokolova, N., Gerrard, N., Rødningsby, A., Rost, C., & Ruotsalainen, L. (2023). Radio-frequency interference considerations for utility of the Galileo E6 signal based on long-term monitoring by ARFIDAAS. *NAVIGATION*, *70*(1). https://doi.org/10.33012/navi.560

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., & 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., & Humphreys, T. E. (2022). Software-defined GNSS is ready for launch. *Proceedings of the ION GNSS+ Meeting*.

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., & 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., & Won, J.-H. (2024). GNSS software defined radio: History, current developments, and standardization efforts. *NAVIGATION*, *71*(1). https://doi.org/https://doi.org/10.33012/navi.628

Patil, A., Phelts, R. E., Chen, Y.-H., Lo, S., & Walter, T. (2023). Detecting space based interference on GNSS signals. *Proceedings of the ION GNSS+ Meeting*, 1232–1244.

Patil, A., Phelts, R. E., Walter, T., & Thoelert, S. (2024). Detecting and localizing space based interference on GNSS signals using machine learning. *Proceedings of the ION International Technical Meeting*, 532–545.

Psiaki, M. L., & Humphreys, T. E. (2016). GNSS spoofing and detection. *Proceedings of the IEEE*, *104*(6), 1258–1270.

Psiaki, M. L., & Humphreys, T. E. (2020). Position, navigation, and timing technologies in the 21st century: Integrated satellite navigation, sensor systems, and civil applications. Wiley-IEEE. https://doi.org/10.1002/9781119458449.ch25

RIN Maritime GNSS Interference Working Group. (2026). *The impacts of GNSS interference on maritime saftey report* (tech. rep.). Royal Institute of Navigation. https://rin.org.uk/page/RIN_Maritime_Report

Scott, L. (2003). Anti-spoofing and authenticated signal architectures for civil navigation systems. *Proceedings of the ION GNSS Meeting*, 1542–1552.

Sokolova, N., Morrison, A., & Diez, A. (2022). Characterization of the GNSS RFI threat to DFMC GBAS signal bands. *Sensors*, *22*(22). https://doi.org/10.3390/s22228587

Stader, J., & Gunawardena, S. (2021). Leveraging worldwide, publicly-available data to create an automated satnav interference detection system. *Proceedings of the ION International Technical Meeting*, 69–83.

Stein, S. (1981). Algorithms for ambiguity function processing. *IEEE Transactions on Acoustics, Speech and Signal Processing*, *29*(3), 588–599. https://doi.org/10.1109/TASSP.1981.1163621

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., & 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.