Received 1 September 2020 | Revised: 00 Month 0000 | Accepted: 00 Month 0000

DOI: xxx/xxxx

ORIGINAL ARTICLE

First results from three years of GNSS Interference monitoring from Low Earth Orbit

Matthew J. Murrian*[1] | Lakshay Narula[2] | Peter A. Iannucci[1] | Scott Budzien[3] | Brady W. O'Hanlon[4] | Mark L. Psiaki[5] | Todd E. Humphreys[1]

[1]Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, Texas
[2]Department of Electrical and Computer Engineering, The University of Texas at Austin, Austin, Texas
[3]Naval Research Laboratory, Washington, District of Columbia
[4]Cornell University, Ithaca, New York
[5]Department of Aerospace and Ocean Engineering, Virginia Tech, Blacksburg, Virginia

Correspondence
Matthew J. Murrian, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, TX 78712.
Email: matthew.murrian@utexas.edu

Summary

Observation of terrestrial GNSS interference (jamming and spoofing) from low-earth orbit (LEO) is a uniquely effective technique for characterizing the scope, strength, and structure of interference and for estimating transmitter locations. Such details are useful for situational awareness, interference deterrence, and for developing interference-hardened GNSS receivers. This paper explores the performance of LEO interference monitoring and presents the results of a three-year study of global interference, with emphasis on a particularly powerful interference source active in Syria since 2017.

KEYWORDS
GNSS interference; spoofing; emitter localization; Doppler positioning

1 | INTRODUCTION

Terrestrial GNSS interference activity has grown more widespread and sophisticated over recent years. Conspicuous GNSS jamming or spoofing has occurred, or is ongoing, at urban and coastal sites around the globe (Ala'Darabseh & Tedongmo, 2019; Brimelow, 2018; C4ADS, 2019; Sebastian, 2016). Given the dependence of critical infrastructure and safety-of-life systems on GNSS (John A. Volpe National Transportation Systems Center, 2001; Psiaki & Humphreys, 2016a; Shepard, Humphreys, & Fansler, 2012; Wesson & Humphreys, 2013), there is great interest in detecting, characterizing, and localizing sources of interference.

Space-based observation of terrestrial GNSS interference is attractive for several reasons. Most obviously, it offers worldwide coverage: moderately-powerful terrestrial interference sources anywhere on the globe can be detected by low-earth orbit (LEO) satellites multiple times per day, making it possible to maintain a common operating picture of world-wide GNSS interference. Moreover, LEO satellites' stand-off distance from terrestrial interference sources often permits tracking authentic GNSS signals despite the interference, allowing precise estimation of a LEO receiver's position, velocity, and time, which, in turn, supports estimation of interference transmitter locations. A single LEO-based sensor is sufficient to characterize the strength, spectral properties, structural content, and even the location of terrestrial interference sources, provided a Doppler time history can be extracted from a carrier component of the interference signal. For signals from which no carrier can be isolated, multiple synchronized LEO-based sensors can employ time- and frequency-difference-of-arrival (TDOA and FDOA) techniques to infer the source's location (Bhatti, 2015; Bhatti & Humphreys, 2017).

This paper presents the results of a three-year study of terrestrial GNSS interference as observed through a software-defined GNSS receiver operating since February 2017 on the International Space Station (ISS). The FOTON receiver, developed by The University of Texas at Austin (UT) and Cornell University, is part of a larger science

experiment called GPS Radio Occultation and Ultraviolet Photometry—Colocated (GROUP-C), an unclassified experiment aboard the ISS that is part of the Space Test Program—Houston Payload 5 (STP-H5) payload. Serendipitous observations of GNSS interference in the occultation data are an important early result of GROUP-C's scientific objective to characterize GPS signals in the LEO environment. This paper discusses the interference signals detected, their effects, and interference mitigation strategies for receivers deployed in LEO and terrestrial environments.

The FOTON receiver is a science-grade spaceborne dual-frequency (GPS L1 and L2) GNSS receiver (Lightsey et al., 2014). Three levels of FOTON data are available for interference analysis: (1) raw 5.7 Msps intermediate frequency (IF) samples output by the FOTON front-end's analog-to-digital converter, (2) 100-Hz data-modulation-wiped complex correlation products, and (3) 1-Hz standard GNSS observables.

Although spaceborne GNSS sensors have been used for remote sensing via radio occultation (Ao et al., 2009) and reflectometry (Jin & Komjathy, 2010), no prior public literature explores their use for monitoring terrestrial GNSS interference, despite increasing concern over such interference (Humphreys, 2017; Psiaki & Humphreys, 2016a 2016b; Wesson, Gross, Humphreys, & Evans, 2018). Moreover, the recent survey of GNSS interference localization techniques in Dempster & Cetin (2016) makes no mention of single-receiver Doppler-based localization, whether space-based or not. General TDOA and FDOA interference localization has been extensively studied (Amar & Weiss, 2008; Bhatti, 2015; Griffin & Duck, 2002; Ho & Chan, 1997; Pattison & Chou, 2000), and such techniques have been applied for terrestrial interference localization from geostationary orbit (Haworth, Smith, Bardelli, & Clement, 1997; Ho & Chan, 1993; Smith & Steffes, 1989). Application of T/FDOA for localization from LEO can be viewed as an extension of such demonstrations enabling localization of much weaker signals. Interference localization using a single satellite has been explored in Kalantari, Maleki, Chatzinotas, & Ottersten (2016), but only simulation results are presented, and these unrealistically assume perfect-tone interference with a known and constant frequency.

This paper makes three primary contributions. First, it introduces the concept and presents an analysis of expected performance for terrestrial GNSS interference monitoring from LEO. Second, it presents the results of a three-year study of global GNSS interference, with emphasis on a powerful interference source active in Syria since 2017. Via Doppler positioning using the FOTON instrument on the ISS, the Syrian transmitter is located to within less than 1 km, an achievement without precedent in the open literature. Third, this paper explores the implications of interference of the type generated by the Syrian source for GNSS receiver operation and design.

A preliminary version of this paper was published in Murrian, Narula, & Humphreys (2019). The current version extends the analysis period to June 2020, offers a more detailed analysis of localization accuracy, and includes a new section exploring implications for GNSS receivers.

2 | LEO INTERFERENCE MONITORING PERFORMANCE

This section explores the potential performance of LEO-based GNSS interference monitoring in terms of sensitivity, visit interval, and source location accuracy.

2.1 | Sensitivity

Interference in the GNSS bands can be detected via *anomaly* detection or *tailored* detection. Anomaly detection seeks to discover interference by noting its effect on measurements made by a GNSS receiver in its usual course of operation. It assumes no particular model for the interference: any signal whose power intersects the band of interest is potentially detectable. Received power monitoring, signal quality monitoring, and carrier-to-noise ratio monitoring are examples of anomaly detection strategies (Broumandan, Kennedy, & Schleppe, 2020).

Tailored detection assumes that the interference has a low-dimensional (sparse) representation in some basis (Baraniuk, 2007). For example, continuous-wave (tone) interference is sparse in a basis of complex exponentials, and spread-spectrum interference may be sparse in a basis of spreading codes. A standard GNSS receiver can only perform tailored detection on GNSS-like interference signals, or perhaps on signals that are sparse in the frequency domain (if it performs routine spectral monitoring). By contrast, a software-defined GNSS receiver can be configured for tailored detection of any number of sparse signals. In any case, the detection process is identical: the receiver performs matched filtering against an assumed signal basis to compress the interference signal's energy into a small number of coefficients, which in turn are distilled into a single detection statistic.

When a sparse basis can be found for an interference signal, tailored detection can be much more sensitive than anomaly detection (Van Trees, 2001). But any increase in the number of unknown signal parameters (e.g, carrier phase, code delay, Doppler frequency, data modulation, spreading code, etc.), reduces the signal's sparsity from the receiver's viewpoint, resuling in depressed detection sensitivity. Thus, the sensitivity of tailored detection depends on how much the receiver

knows *a priori* about the interference signal it is trying to detect.

What follows is a sensitivity analysis of four example detection techniques applicable to LEO interference monitoring. The first two are anomaly detectors whereas the latter two are tailored detectors.

2.1.1 | Detection via C/N_0 Monitoring

A simple and effective anomaly detection test can be formulated from the standard carrier-to-noise ratio observable, C/N_0, produced by a GNSS receiver. In the presence of interference, C/N_0 actually measures the carrier-to-interference-and-noise ratio, CINR. Let C be the received authentic signal power for a particular satellite-and-signal combination [e.g., the GPS L1 C/A signal corresponding to pseudo-random number (PRN) code 4], N_0 be the (approximately flat) receiver thermal noise power density near the frequency band of interest, and I_0 be the spectrally-flat-equivalent interference noise power density, whose relationship with the actual interference power spectrum is described in Humphreys (2017). Then CINR is defined as

$$\text{CINR} \triangleq \frac{C}{N_0 + I_0} \tag{1}$$

When compensated for satellite- and receiver-side antenna gain patterns and for path loss along the satellite-to-receiver path, and absent signal blockage, strong scintillation, and "flex-power" satellite power adjustments, CINR variations are primarily driven by multipath, which is characterized by a log-normal distribution (Wesson et al., 2018). Let $\boldsymbol{z}$ be a vector of CINR measurements expressed in dB for a particular frequency band, with predictable variations due to antenna gain pattern and path loss removed. A hypothesis test for interference can be formulated in terms of the common decrease in the elements of $\boldsymbol{z}$ due to an increase in I_0. In particular, the distribution of $\boldsymbol{z}$ under the null (H_0) and alternate (H_1) hypotheses may be modeled as

$$H_0 : \boldsymbol{z} \sim \mathcal{N}(\boldsymbol{\mu}, P) \tag{2a}$$

$$H_1 : \boldsymbol{z} \sim \mathcal{N}(\boldsymbol{\mu} - \delta\mathbf{1}, P) \tag{2b}$$

where $\boldsymbol{\mu} \in \mathbb{R}^{n_z}$ is the mean of $\boldsymbol{z}$ under H_0, $P \in \mathbb{R}^{n_z \times n_z}$ is the covariance of $\boldsymbol{z}$ under either hypothesis, $\mathbf{1}$ denotes an all-ones column vector of the same length as $\boldsymbol{\mu}$, and $\delta > 0$ is the amount in dB by which all CINR values drop due to interference under H_1. The parameters $\boldsymbol{\mu}$ and P can be determined based on analysis of historical CINR data for which H_0 is known to be true, as will be detailed in Sec. 4.

The model in (2) conservatively assumes that $\boldsymbol{z}$'s covariance matrix, P, is identical for H_0 and H_1. In practice, although the receiver's multipath environment remains unchanged from H_0 to H_1, interference sources can cause time variations in I_0 that inflate P in the positive definite sense. But because the magnitude of increase in P is impossible to know *a priori*, the less-sensitive model presented above is assumed.

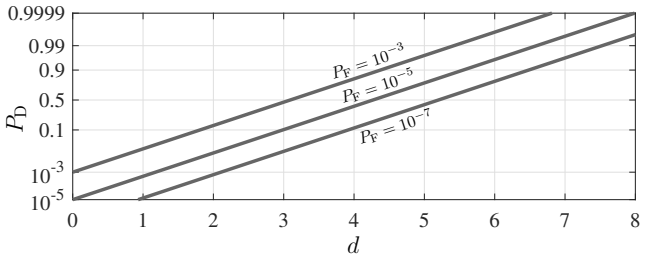

FIGURE 1 Detection probability for the test in (3) as a function of d for three different values of false alarm probability.

The model in (2) is a special case of the general Gaussian problem for which the likelihood ratio test can be reduced to (Van Trees, 2001)

$$l(\boldsymbol{z}) = \mathbf{1}^T P^{-1} \boldsymbol{z} \underset{H_1}{\overset{H_0}{\gtrless}} \nu \tag{3}$$

where $l(\boldsymbol{z})$ is the test's (sufficient) detection statistic and ν is the test threshold, which is chosen such that the test's probability of false alarm is sufficiently small. Note that choice of ν for a given false alarm probability depends on $\boldsymbol{\mu}$ and P.

The test in (3) is optimal despite $\delta > 0$ being unknown *a priori* because $l(\boldsymbol{z})$ is independent of δ (i.e., the test is uniformly most powerful with respect to δ). Note that P may not be diagonal because the elements of $\boldsymbol{z}$ may be correlated through dependence on the spacecraft attitude or because $\boldsymbol{z}$ may contain multiple elements for the same satellite-signal pair taken over a sliding window of time.

As a linear transformation of a Gaussian vector, $l(\boldsymbol{z})$ is itself Gaussian. Hence, the performance of the test in (3) can be completely characterized by the normalized distance between the means of $l(\boldsymbol{z})$ under H_0 and H_1:

$$d \triangleq \frac{\mathbb{E}\left[l|H_0\right] - \mathbb{E}\left[l|H_1\right]}{\sqrt{\text{Var}(l|H_0)}} = \delta\sqrt{\mathbf{1}^T P^{-1}\mathbf{1}} \tag{4}$$

Fig. 1 shows how the performance improves with increasing d.

If the CINR measurements in $\boldsymbol{z}$ are taken at a single epoch of time, and if the effect of multipath on each measurement is only weakly coupled through the spacecraft attitude, then P may be modeled as diagonal. In the simplest case, $P = \sigma_z^2 I$ and d reduces to

$$d = \delta\sqrt{n_z}/\sigma_z \tag{5}$$

For the FOTON receiver on the ISS, the ISS's extended shape and large solar panels create an unfavorable multipath environment, resulting in a relatively high $\sigma_z \approx 1.5$ dB. More compact LEO satellites such as the main sounding rocket payload in Lightsey et al. (2014) enjoy $\sigma_z < 1$ dB.

Approximate LEO interference detection sensitivity in the L1 GNSS band using only CINR measurements can be calculated by assuming $\sigma_z = 1$ dB and $n_z = 15$, which are reasonable parameters for a single-epoch test, a horizontally-oriented hemispherical-gain antenna, and full constellations of GPS, Galileo, and BDS III satellites. From (5) and Fig. 1 , a drop in CINR of $\delta > 1.4$ dB is required at $P_{\mathrm{F}} = 10^{-5}$ to yield $P_{\mathrm{D}} > 0.9$. Conservatively assuming that the interference power is spread evenly across the 4-MHz bandwidth covering the most-widely-used civil L1 GNSS signals, then $I_0 = P_{\mathrm{I}} - 66$ dBW/Hz, where P_{I} is the received interference power in dBW. Assuming $N_0 = -204$ dBW/Hz, a CINR drop by $\delta = 1.4$ dB implies $P_{\mathrm{I}} = -142$ dBW. Denote path loss by L dB, receiver antenna gain by G_r dB, and interference source isotropic radiated power by P_{S} dBW. Then

$$P_{\mathrm{S}} = P_{\mathrm{I}} - G_r + L \tag{6}$$

Path loss at L1 from the surface along the shortest distance to a typical LEO altitude of 400 km is $L = 148.5$ dB. Then, supposing $G_r = 3$ dB, the minimum power of an isotropically-radiating interference source detectable solely from CINR measurements with $P_{\mathrm{F}} \leq 10^{-5}$ and $P_{\mathrm{D}} > 0.9$ is approximately $P_{\mathrm{S}} = 3.5$ dBW.

2.1.2 | Detection via Received Power Monitoring

Like CINR monitoring, received power monitoring is an anomaly detection strategy, but it but avoids the requirement to assemble $\boldsymbol{z}$ only from authentic GNSS signals, which can be difficult under spoofing interference. In fact, received power monitoring requires no tracking of signals at all.

For systems with multi-bit-quantized sampling, total received power P_{T} can be estimated from the dynamic gain setting of an automatic gain control (AGC) unit in the front-end digitizer, or directly from the pre-correlation samples in a constant-gain system, assuming sufficient dynamic range to avoid quantization saturation. The hypothesis test model is identical to (2) with $\boldsymbol{z} = P_{\mathrm{T}} \in \mathbb{R}$ and $n_z = 1$. Its performance is governed by (5), with $\delta < 0$ re-defined as the negative of the increase in P_{T} under H_1, and σ_z^2 as the variance of the unmodelable components of P_{T}.

Received power monitoring for interference detection may be more or less sensitive than CINR monitoring depending on the spectral characteristics of the interference. [MMM: I'd like you to provide here some analysis of the relative sensitivity in three different cases, assuming an equivalent interference power received in the band but three different spectral shapes:

(1) spectrally flat interference across the band. I'm thinking the P_D values will be similar in this case, but I could be wrong.

(2) matched-spectrum interference, in which the source allocates its signal power to match the spectrum of a target authentic GNSS signal, thus maximizing I_0 for a receiver tracking that signal (Humphreys, 2017). It's clear that CINR-based detection will have the advantage here because I_0 will be maximized. But it would be good to show something analytically.

(3) interference with power allocated to the edges of the band. In this case, I expect received power monitoring to do better, because convolution of the GNSS spectrum with the edge-allocated spectrum will leave a depression near the center, which means I_0 will be low.]

2.1.3 | Detection based on Power Spectrum Monitoring

[MMM: Note that wang2017gnss appears flawed in the sense that it doesn't take into account the correlation between different frequency bins under H0, whereas the spectrum monitor you and I designed *does* take such correlation into account. That's why I'm eager for your other paper to be on spectrum monitoring – I'd like for us to set the record straight.]

Power spectrum monitoring is a simple and powerful GNSS interference diagnostic technique, indicating not only the presence but also the nature of interference: whether wideband or narrowband, constant or fleeting (Wang, Cetin, Dempster, Wang, & Wu, 2017). Monitoring is typically effected by generating periodograms (power spectral estimates) from pre-correlation signal samples at a regular cadence. A detection statistic is extracted from each periodogram, based on a likelihood ratio test that is sensitive to the periodogram's departures from a probability distribution learned during interference-free data collection intervals.

The key challenge of interference detection via spectrum monitoring is distinguishing novel spectral variations from background variations due to multipath, changing multi-access interference, and temperature variations. Insofar as the background variations are modelable as stationary random processes, they can be accurately accounted for within the null-hypothesis probability distribution for the power spectrum.

[MMM: Could we offer a "bookend" analysis of power spectrum monitoring sensitivity: one for wideband noise (flat across passband) and one for narrowband noise (CW interference)? Narrowband noise would be much more easily detected. Would be good to offer a number for the power of a terrestrial CW source that could be detected by spectrum monitoring. Then we say that actual detection performance will depend on the type of interference and on the tightness of the H0 probability distribution.]

2.1.4 | Detection via Signal Acquisition

When a matched-spectrum interferer employs a standard GNSS spreading code to achieve the requisite spectrum-matching, it becomes a *matched-code interferer*, which, as will be shown in Section 5, can be effective at denying GNSS service to surrounding receivers on cold-start. However, matched-code interference is itself vulnerable to high-sensitivity detection because a distant receiver can acquire the interference signal just as it does an authentic GNSS signal. Moreover, a receiver in LEO can despread the matched-code interference with the known spreading code, thus extracting a pure carrier tone from whose Doppler time history the source may be geolocated, as will be detailed later on.

Consider the sensitivity of matched-code interference detection via signal acquisition from LEO. Denote the LEO receiver's C/N_0 acquisition threshold by ν_a dB-Hz. Detection via acquisition is possible when $P_{\mathrm{I}} - N_0 > \nu_a$, with P_{I} expressed in dBW and N_0 in dBW/Hz. For the same values of N_0, G_r, and L assumed in Section 2.1.1, and conservatively supposing $\nu_a = 30$ dB-Hz, the minimum detectable EIRP of a terrestrial interference source is approximately $P_{\mathrm{EIRP}} = -28.5$ dBW. Thus, detection of matched-code interference by signal acquisition is more than 1000 times more sensitive than detection of unpredictable wideband interference via C/N_0 monitoring or received power monitoring.

2.2 | Detection Frequency

A terrestrial interference source is potentially detectable by LEO satellite monitoring several times a day. Consider a LEO satellite in a near-ISS orbit: circular, 400-km altitude, and 55° inclination. Assuming detection by C/N_0 monitoring with the parameters given in Section 2.1.1, Fig. 2 shows the average number of times per day that such a satellite could detect an interference source, as a function of P_{EIRP} and latitude. Sources with $P_{\mathrm{EIRP}} = 3.5$ dBW are detectable only when the satellite's ground track crosses directly through the source's location. As P_{EIRP} rises, detection becomes possible even as the satellite ground track passes ever further from the source, increasing detection frequency. This behavior saturates for $P_{\mathrm{EIRP}} \geq 17$ dBW, in which case a minimum of 3 detections occur per day for all latitudes within 75° of the equator.

Besides the average detection frequency shown in Fig. 2 , it is instructive to consider the maximum time between detections for a given P_{EIRP}. Analysis of the ground-track lattice formed by a LEO satellite with the above orbital parameters reveals that every 4 days the lattice is sufficiently dense to guarantee detection of transmitters with $P_{\mathrm{EIRP}} > 6.1$ dBW, and every 17 days detection of transmitters with $P_{\mathrm{EIRP}} > 3.65$ dBW, for all latitudes within 55° of the equator.

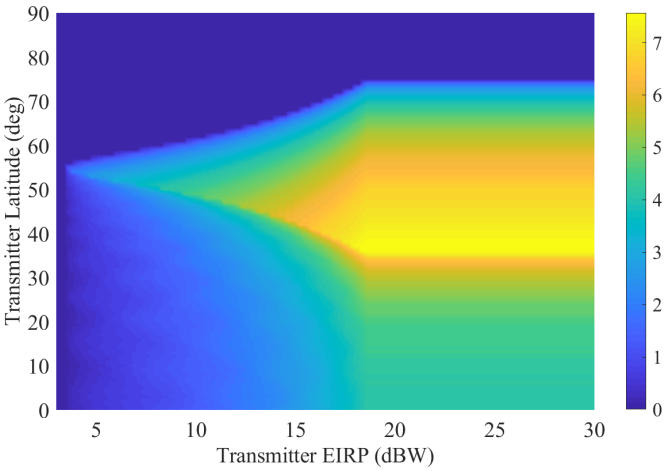

FIGURE 2 Number of times per day that a LEO satellite on a circular, 55° inclination, 400-km altitude orbit could detect a given terrestrial transmitter by C/N_0 monitoring as a function of the transmitter's P_{EIRP} and absolute-value latitude, averaged over a 30-day interval.

2.3 | Geolocation Accuracy

2.3.1 | Multi-Satellite Geolocation

Time- and frequency-difference-of-arrival (T/FDOA) techniques have been explored over the past decades for space-based terrestrial interference localization. These techniques require at least two time-synchronized satellites. Reference (Haworth et al., 1997) studied interference geolocation via EUTELSAT satellites, presenting theoretical models and real-world campaigns assessing the performance of combined FDOA and TDOA techniques. Accuracies from tens to hundreds of km were theorized and demonstrated. The authors identified satellite ephemeris errors as the dominant source of location error.

In Ho & Chan (1993), TDOA-based interference geolocation was analyzed for the scenario of three geostationary satellites able to simultaneously observe the interfering signal. The analysis showed that location accuracy is improved by increased orbital spacing between the observing satellites and by reduced TDOA measurement error. In this scenario, a transmitter at a latitude greater than 40° can be located to 2 km (1-σ) with a satellite spacing of 2° if the TDOA measurements have a standard deviation of less than 3.88 ns. Alternatively, a satellite spacing of 30°. yields the same location precision for TDOA measurements of less than 0.832 µs standard deviation.

Joint TDOA/FDOA geolocation from two satellites has been studied in (Ho & Chan, 1997; Pattison & Chou, 2000), and particularly for LEO satellites in Shilong, Jingqing, & Liangliang (2010). In the latter it was shown that two LEO satellites flying in parallel formation could provide on the

order of 1 km localization from an orbital altitude of 800 km, with a 50-km inter-satellite baseline, TDOA measurement errors of 10 ns, and FDOA measurement errors of 4 Hz for signals centered at 3 GHz. Similar performance could be expected for LEO-based geolocation of interference sources at GNSS frequencies.

2.3.2 | Single-Satellite Geolocation

Assuming a carrier can be extracted from an interference signal, single-satellite-based transmitter geolocation is possible from Doppler measurements alone (Becker, 1992; Ellis, Van Rheeden, & Dowla, 2020). The analysis presented here emphasizes the effect of transmitter clock stability on geolocation accuracy.

Consider a static transmitter emitting a signal at the GPS L1 frequency as observed by a moving receiver. Let λ be the signal wavelength in meters, $\hat{\boldsymbol{r}}$ the unit vector pointing from the transmitter to receiver, expressed in Earth-centered-Earth-fixed (ECEF) coordinates, $\boldsymbol{v}_{\mathrm{R}}$ the receiver velocity with respect to the ECEF frame and expressed in ECEF in m/s, and $\delta\dot{t}_{\mathrm{R}}$ the receiver clock frequency error in s/s, all at the time of signal receipt. Further, let $\delta\dot{t}_{\mathrm{T}}$ be the transmitter clock frequency error in s/s at the time of signal transmission, and w be a zero-mean Gaussian error term that models thermal noise, ionospheric and tropospheric delay rates, and other minor effects, in Hz. Then the observed Doppler frequency in Hz at the receiver can be modeled as

$$f_{\mathrm{D}} = -\hat{\boldsymbol{r}}^{T}\boldsymbol{v}_{\mathrm{R}}/\lambda - c\left[\delta\dot{t}_{\mathrm{R}} - \delta\dot{t}_{\mathrm{T}}\left(1-\delta\dot{t}_{\mathrm{R}}\right)\right]/\lambda + w \quad (7)$$

where c is the speed of light in m/s. It is assumed that $\boldsymbol{v}_{\mathrm{R}}$, $\delta\dot{t}_{\mathrm{R}}$, and the receiver position are known, e.g., via an onboard GNSS receiver. The unknown in (7) is transmitter position, which is embedded in $\hat{\boldsymbol{r}}$, and $\delta\dot{t}_{\mathrm{T}}$. The former is modeled as an unknown constant and the latter as a random walk process that evolves as

$$\dot{\delta t}_{\mathrm{T}}(t_{k+1}) = \dot{\delta t}_{\mathrm{T}}(t_k) + v(t_k) \quad (8)$$

Here, $v(t_k)$ is a discrete-time Gaussian random process with $\mathbb{E}\left[v(t_k)\right] = 0$ and $\mathbb{E}\left[v(t_k)v(t_j)\right] = 2\pi^2 h_{-2}\delta t\delta_{k,j}$, $\forall k, j$, where h_{-2} is the first parameter of the standard clock model based on the fractional frequency error power spectrum, as given in (Brown & Hwang, 2012, Chap. 8); $\delta t = t_{k+1} - t_k$ is the uniform sampling interval; and $\delta_{k,j}$ is the Kronecker delta.

It should be noted that a transmitter could introduce any level of complexity to carrier-phase frequency behavior; e.g., frequency modulation, frequency hopping, etc. Such behaviors, if not discovered and appropriately modeled, would confound geolocation efforts. Here, it is assumed that a nominally-constant carrier frequency is intended by the transmitter and that it is operating in steady-state conditions. In fact, it will be assumed that h_{-2} is sufficiently small that $\dot{\delta t}_{\mathrm{T}}$ can be modeled as constant over a short (e.g., 60-second) data capture interval.

Using the Doppler measurement model from above, a batch maximum likelihood estimator (Crassidis & Junkins, 2011) can be developed to estimate the unknown transmitter position and a constant value for $\dot{\delta t}_{\mathrm{T}}$ from a collection of single-pass Doppler measurements. If Doppler measurements from multiple satellite passes are available, these can be combined for single-batch estimation provided that a new value of $\dot{\delta t}_{\mathrm{T}}$ is estimated for each pass. In other words, $\dot{\delta t}_{\mathrm{T}}$ is viewed as constant over the short capture interval but variable from orbit to orbit.

When $\dot{\delta t}_{\mathrm{T}}$ is modeled as constant over a capture interval, actual transmitter clock instability gives rise to Doppler measurement errors. The impact of such errors on geolocation accuracy was analyzed via Monte Carlo simulation for four levels of transmitter clock quality, from low-quality temperature-compensated crystal oscillator (TCXO) to a laboratory-grade oven-controlled crystal oscillator (OCXO). For each clock quality level, 1000 Monte Carlo simulations were run. Simulation parameters were based on the real-world interference capture discussed in the next section: the transmitter location was 35.4N latitude, 35.95E longitude, 48m altitude; the receiver trajectory was taken from the ISS orbit during the first 60 seconds of the capture interval on day 144 of 2018 (resulting in 441.65 km of total receiver displacement); and the measurement rate was 20 Hz. First, an error-free Doppler time history was generated based on this scenario. Then, for each instance of the Monte Carlo simulation, an independent realization of a Doppler error random process consistent with the clock model being analyzed was generated and added to the error-free Doppler. Doppler error was modeled as a random walk process consistent with (8).

TABLE 1 Single-pass geolocation accuracy as a function of transmitter clock frequency stability parameter h_{-2}. The size of the 95% horizontal geolocation error ellipse, in meters, is characterized by the semi-major (a) and semi-minor (b) axes.

Clock Quality	h_{-2}	a (m)	b (m)
Low-quality TCXO	1×10^{-20}	13027.4	7.8
TCXO	3×10^{-21}	6712.0	2.0
Low-quality OCXO	3×10^{-23}	713.4	2.5×10^{-2}
OCXO	3×10^{-25}	71.2	3.4×10^{-3}

Table 1 shows that transmitter clock frequency stability has a large effect on single-pass geolocation accuracy. Note that the error ellipse is highly eccentric. Its semi-minor axis is oriented in the direction of satellite motion; e.g., if the satellite is moving west to east then transmitter location will be best

resolved in that direction. It follows that additional satellite passes provide the most benefit when, relative to the transmitter location, they are geometrically dissimilar to previous passes.

3 | ANALYSIS OF INTERFERENCE FROM SYRIA

This section presents an in-depth analysis of a particular interference source active on the east coast of the Mediterranean Sea during the period of this paper's study, which spans from March 2017 to June 2020. The analysis illustrates the techniques that can be applied generally to study terrestrial GNSS interference sources using signals collected in LEO.

Recording raw IF data in LEO and relaying these to the ground for processing is an especially flexible approach well suited to studying new or poorly-understood interference. For the case presented here, the FOTON receiver captured 1-minute intervals of raw 5.7-Msps two-bit-quantized IF samples at the GPS L1 (1575.42 MHz) and the GPS L2 (1227.6 MHz) frequencies. These data were packaged and downlinked via NASA's communications backbone. Ground processing using the latest version of UT's software-defined GNSS receiver (Humphreys, Murrian, & Narula, 2020) enabled analysis and tracking of all radio frequency signals near GPS L1 and L2.

The following observations are based on particularly strong interference signals captured on three days in the first half of 2018 along the ground tracks shown in Fig. 3 .

3.1 | Overview

Strong interference is present in both the L1 and L2 bands, but the nature of the interference is markedly different between the two bands. At L2, the interference is narrowband, whereas wideband matched-code interference was discovered at L1. The L1 interference is a composite of individual signals with a common carrier centered near GPS L1 and a unique GPS L1 C/A pseudo-random number (PRN) code. Signals corresponding to almost all PRN codes from 1 to 32 have been detected. When tracked, all false signals exhibit C/N_0 values greater than 40 dB-Hz. No discernible navigation data are modulated on the false GPS L1 signals, rendering them ineffective at spoofing. Moreover, the false signals are not clean simulated GPS L1 C/A signals; they exhibit unexplained fading and spectral characteristics, as if generated from an extremely low-quality GNSS signal simulator. No false Galileo BOC(1,1) signals were detected in the L1 band.

FIGURE 3 Ground tracks for interference-affected captures on days 74, 144, and 151 of 2018. Each capture spans approximately 70 seconds. The estimated transmitter location is marked on the west coast of Syria.

The lack of navigation bit modulation and the coarse nature of the matched-code interference at L1 suggest that its purpose is denial of GPS service (jamming) rather than spoofing. The narrowband interference at L2 also appears intended for jamming. Why different jamming techniques were used at L1 and L2 is unknown.

While some authentic GPS L1 C/A signals in the data are effectively jammed, the majority of authentic signals are still trackable owing to sufficient separation of corresponding false and authentic signals in code-Doppler space. Thus, a correct receiver navigation solution can still be formed despite the interference.

3.2 | Power Spectral Characteristics

Figs. 4 and 5 illustrate the captured signals' spectral characteristics. The spectra of narrowband interference near L2 are simple and remain similar across all three days, but the wideband interference at L1 is more complex and variable. It is clear from the left column of Fig. 4 that the matched-code interference is cluttered by other components. Were it generated by a high-quality signal simulator, L1 interference would tend to be smooth like the authentic signals underlying the spectrum shown in the lower left panel. Instead, it appears to be a strange amalgam of components. Fig. 5 reveals that the rounded prominence in the L1: Day 144 panel exhibits oscillatory behavior with a 5-second period. Whether such variations are deliberate or are caused by transmitter idiosyncrasies is unknown.

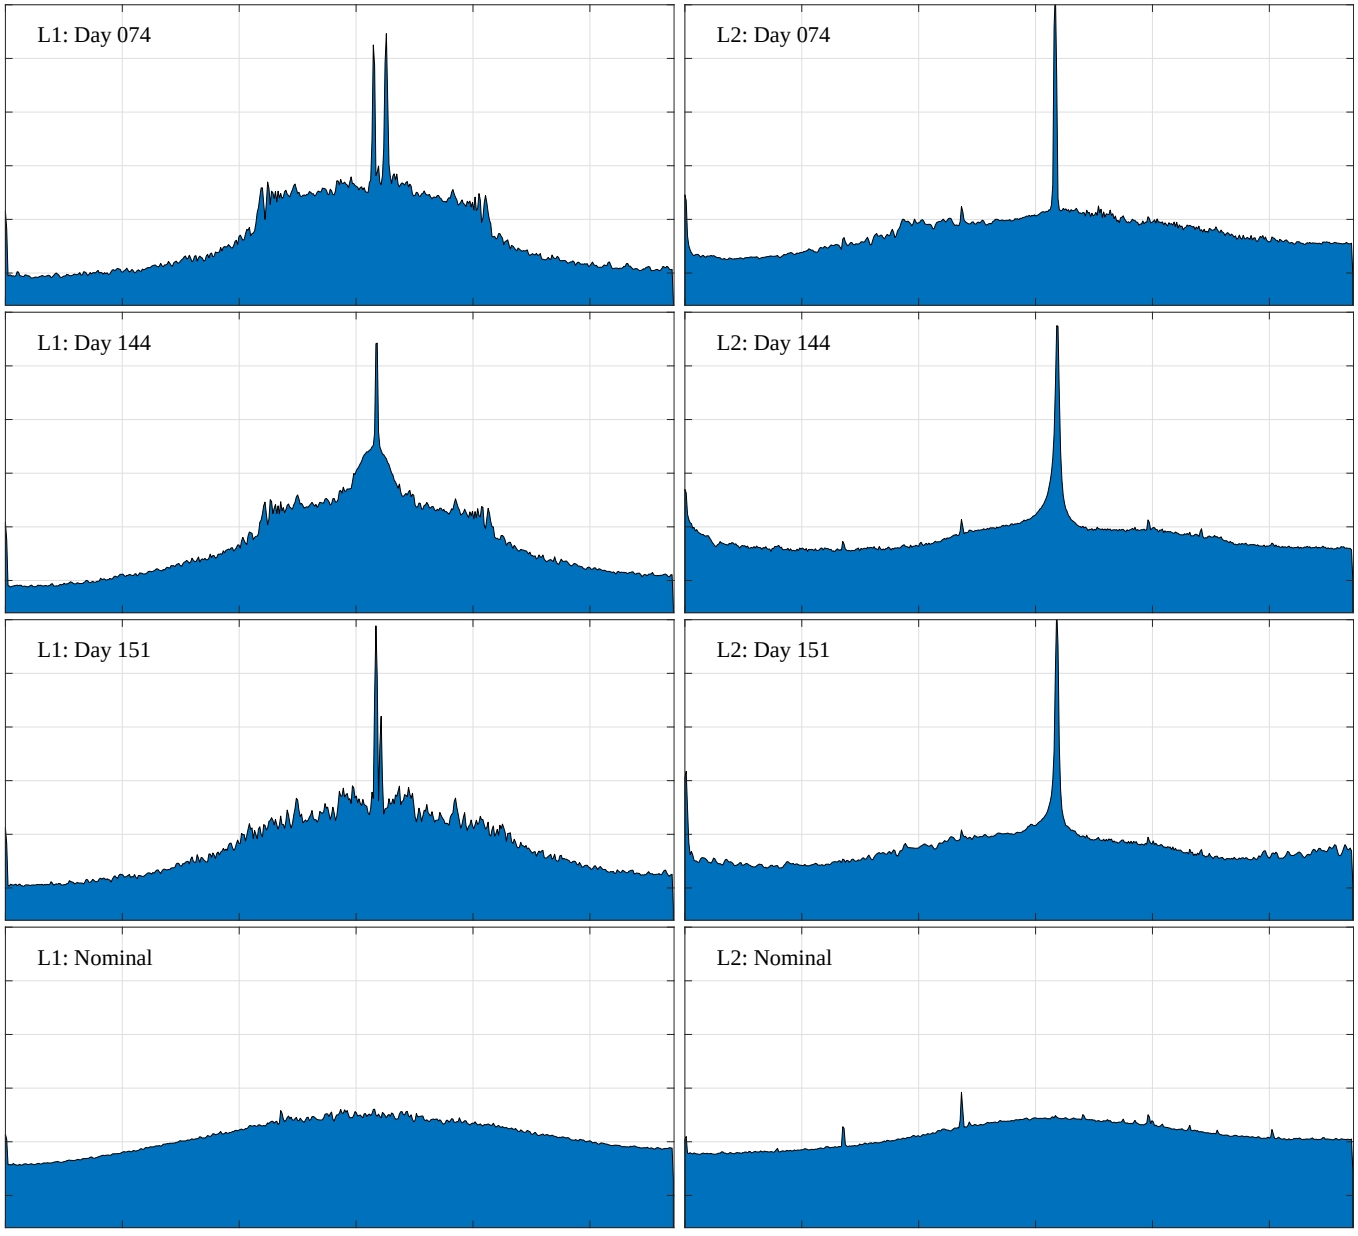

FIGURE 4 Power spectra centered near the GPS L1 (left column) and L2 (right column) frequencies from interference-affected data captured on days 74, 144, and 151 of 2018 (top three rows), and from nominal data captured on day 158 of 2018 (bottom row). The frequency span is approximately 3 MHz wide, scaled linearly with 0.5 MHz divisions. All ordinate axes are in dB and scaled equivalently for ease of comparison. Spectra are estimated by Welch's method (Welch, 1967) from 1-second data intervals with a 5.6-kHz frequency resolution.

3.3 | Baseband Signal Characteristics

Fig. 6 shows time histories of 10-ms-accumulated complex correlation products from both the false (top panel) and authentic (bottom two panels) GPS L1 C/A signals present in the captured L1 band. The false signal's empirical C/N_0 value is 42.5 dB-Hz on average, but the signal is highly irregular, manifesting both gradual and sudden fading. The gradual fading may be a result of scintillation as the signal passes upward through the lower ionosphere (Humphreys, Psiaki, & Kintner, 2010), but the sudden fading, highlighted in the inset of the top panel, is unnatural and likely occurs at the transmitter.

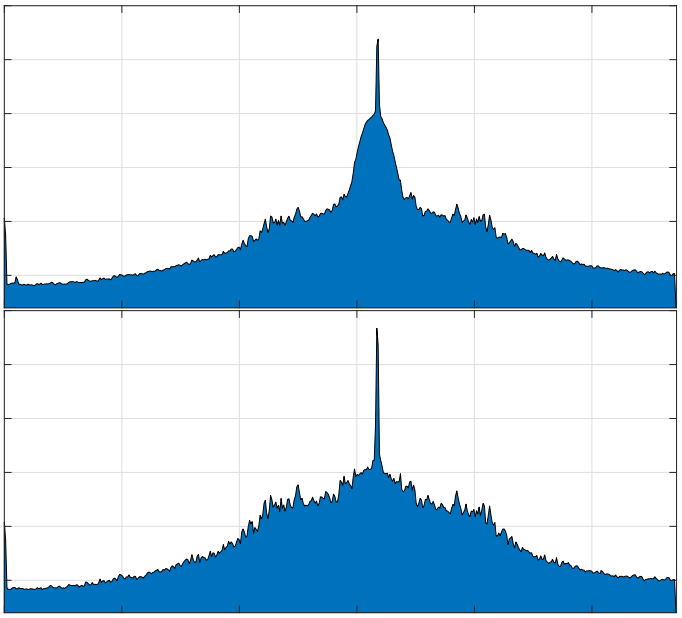

FIGURE 5 Power spectra near L1 for the day 144 capture showing maximum (top) and minimum (bottom) phases of the waxing and waning wideband (∼ 0.25MHz) central interference prominence. The prominence oscillates with a period of approximately 5 seconds. The L1: Day 144 plot in Fig. 4 catches the prominence waning two seconds after the maximum shown in the top plot above.

3.4 | Source Geolocation

The presence of a trackable carrier signal after despreading (cf. top panel of Fig. 6) raised the possibility of geolocating the interference source as described in Section 2.3.2. A receiver navigation solution was first estimated on days 74, 144, and 151 of 2018 using an Extended Kalman Filter (EKF) drawing in pseudorange and Doppler measurements extracted from the authentic GPS L1 C/A, GPS L2C, and Galileo E1 signals. Propagation of the receiver state estimate between measurement updates was based on a nearly-constant acceleration dynamics model. Time histories of the quantities $\boldsymbol{v}_{\mathrm{R}}$, δi_{R}, and the receiver position component of $\hat{\boldsymbol{r}}$ were then extracted from the EKF's state estimate and treated as known for purposes of source geolocation.

A batch estimator for interference source position and clock frequency bias was formulated as described in Section 2.3.2. It was assumed that the interference observed on all three days originated from the same transmitter and that the transmitter was stationary, which allowed multiple days of Doppler measurements, collected on non-repeating ground-tracks, to be combined to form a tightly-constrained estimate. If these assumptions were false, they could be expected to manifest in large post-fit measurement residuals, which was not the case. Consistent with the assumption of a stationary transmitter, transmitter altitude was assumed to be near ground-level and was included as a pseudo-measurement.

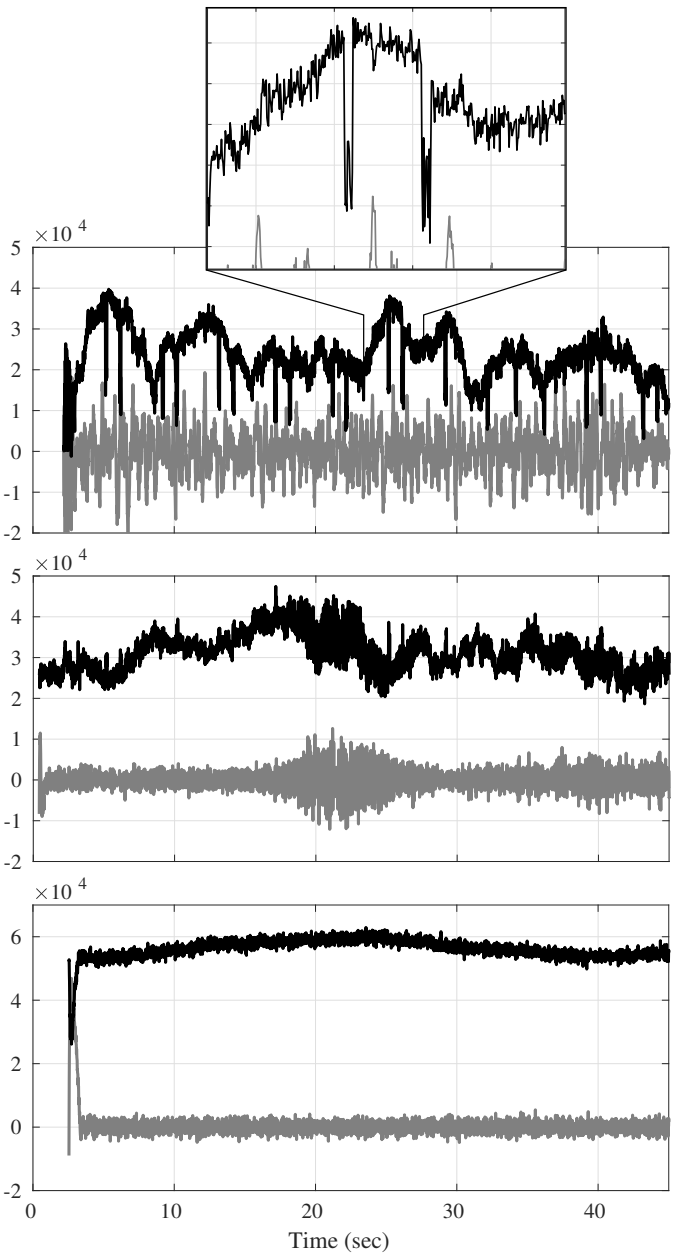

FIGURE 6 In-phase (black) and quadrature (gray) 10-ms accumulation time histories for the strongest false signal from the day 74 capture (top), the strongest authentic signal from the day 74 capture (middle), and the strongest signal from the day 158 nominal capture (bottom). The inset on the top panel shows an amplified view of two sudden amplitude fades in the received false signal. The maximum carrier-to-noise ratio C/N_0 over the intervals shown are, from the top, 42.5, 46.8, and 52.5 dB-Hz.

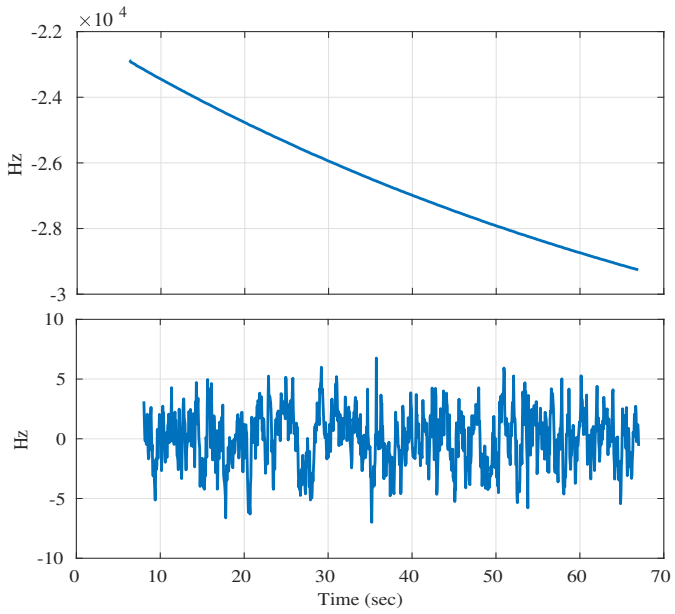

FIGURE 7 Top: Doppler time history corresponding to the false PRN 10 signal from the day 144 capture. Bottom: Post-fit residuals of the Doppler time history assuming the estimated transmitter location and clock rate offset. The standard deviation of the post-fit residuals is 2.3 Hz.

FIGURE 8 Estimated transmitter location overlaid on formal-error 95% and 99% horizontal error ellipses. The location is coincident with an airbase on the coast of Syria.

A constant transmitter clock frequency error $\delta\dot{t}_T$ was assumed to apply during each capture, but a new value of $\delta\dot{t}_T$ was estimated for each of the three captures. Comparing the batch-estimator-produced estimates of $\delta\dot{t}_T$ for days 74 and 144 revealed a two-sample transmitter clock frequency stability of approximately $\sigma_y(2, T, \tau) = 6.85 \times 10^{-9}$ at a sampling interval T of 70 days and an observation time (averaging interval) τ of approximately 50 seconds. The B_2 bias function (Barnes, 1972) was used to convert this two-sample deviation to an Allan deviation, where $B_2(r, \mu) = 1.8144 \times 10^5$ for $r = T/\tau$ and $\mu = 1$, which assumes h_{-2} is the dominant spectral component. This yielded an equivalent Allan deviation for $\tau = 50$ seconds of $\sigma_y(2, \tau, \tau) = 1.6{\times}10^{-11}$, which is consistent with an OCXO.

Thus, given the results of Table 1 , treating $\delta\dot{t}_T$ as constant over each 60-second capture can be expected to introduce 95% errors smaller than 71 meters in single-pass geolocation. A Monte-Carlo simulation like the one that produced the Table 1 data but for the combined three days of collection showed that, assuming independence in the clock frequency errors between passes, this error source can be expected to contribute 95% errors below 24 meters in the combined 3-day solution.

It is worth noting that, because $\delta\dot{t}_R$ and $\delta\dot{t}_T$ enter equivalently into the Doppler measurement model (7), and because no prior knowledge of these parameters is assumed in the batch maximum-likelihood estimator, an error in the EKF's estimate of $\delta\dot{t}_R$ will directly manifest in the batch-estimator-produced estimate of $\delta\dot{t}_T$ for each capture. However, examination of the the EKF's error covariance revealed that its estimate of $\delta\dot{t}_R$ was good to better than 7×10^{-10} (1-σ) for the day 74 and 144 captures. Thus, receiver-side errors are likely small enough that $\sigma_y(2, \tau, \tau) = 1.6 \times 10^{-11}$ remains an accurate assessment of the transmitter clock stability.

Fig. 7 shows time histories of Doppler and post-fit residuals for false PRN 10 collected on day 144. The standard deviation of the post-fit residuals is 2.3 Hz, indicating that the measurement model in (7), and the assumption of a constant $\delta\dot{t}_T$ over each capture, are accurate. Fig. 8 shows the estimated position of the interference source, whose location coincides with an airbase in Syria. The horizontal error ellipses, which indicate a solution better than 220 meters (95%), are formal error ellipses assuming (1) constant $\delta\dot{t}_T$ over each capture, (2) a standard deviation of 5 m for the transmitter altitude constraint, and (3) a standard deviation between 2.3 and 2.5 Hz (depending on the empirical post-fit residuals for each capture) for the measurement error w from (7). Assuming an OCXO-quality clock in the transmitter, the error caused by modeling a constant $\delta\dot{t}_T$ is small compared to these formal error ellipses.

3.5 | Transmitter EIRP

By analyzing the authentic signal CINR values in the captured data one can infer the EIRP of the emitter located in Syria. The average decrease in CINR observed at the ISS when 1340 km from the source was approximately 6 dB. Assume the interference acts as multi-access interference, whose spectral density is $I_0 = (2/3)P_I T_C$ (Humphreys, 2017), where P_I

is the received interference power and $T_C = 1023^{-1}$ ms is the GPS L1 C/A spreading code chip interval. Then, assuming $N_0 = -204$ dBW/Hz, a drop in CINR by 6 dB implies $P_I = -137.4$ dBW. Referring to (6), assume $L = 159$ dB, consistent with a stand-off distance of 1340 km, and $G_r = 3$ dB. It follows that the EIRP of the interference source is $P_{\text{EIRP}} = 18.6$ dBW, which implies a 72-W transmitter.

4 | GLOBAL INTERFERENCE SURVEY VIA RECEIVER-REPORTED CINR

The raw IF data captures from the ISS FOTON receiver enable detailed monitoring of GNSS interference signals and their structure, but such captures are infrequent and limited to short 1-minute intervals. By contrast, the 1-Hz standard GNSS observables and 100-Hz data-wiped complex correlation products have been logged nearly continuously since early 2017. These data facilitate a world-wide survey of strong GNSS interference.

The carrier power C of an authentic signal can be modeled as a function $C(j, f, r_{sr}, z_s, z_r)$, where j is the GNSS satellite identifier (SV ID), f is the frequency band (L1 or L2), r_{sr} is the range between the GNSS satellite antenna and the ISS FOTON antenna, z_s is the angle between the satellite boresight direction and the direction to the ISS antenna (i.e., the satellite antenna zenith angle), and z_r is the angle between the ISS antenna boresight direction and the direction to the satellite (the receiver antenna zenith angle). As discussed in Section 2.1.1, a hypothesis test based on the receiver-reported CINR can be designed to detect whether (H_1) or not (H_0) the receiver is experiencing interference. Under a given P_F, this requires that the statistics $\mathbb{E}[l|H_0]$ and $\text{Var}(l|H_0)$ be known. To obtain these statistics, this section assumes the receiver reports interference-free data (consistent with H_0) when the ISS is over deep ocean bodies.

To isolate the variations in reported CINR due to interference, the data are first pre-processed to eliminate the predictable sources of carrier power variation. First, the dependence of C on r_{sr} is removed by compensating for the free space path loss:

$$\hat{C}(j, f, z_s, z_r) = C(j, f, r_{sr}, z_s, z_r) \times \left(\frac{4\pi r_{sr} f}{c}\right)^2$$

Modeling of interference-free C/N_0 is complicated by the ISS's local multipath environment. The ISS antenna is flanked by solar panels that move with respect to the FOTON antenna, causing a non-stationary signal obstruction and multipath environment. Nevertheless, a zenith angle window $z_r \in [0°, 15°]$ is known to be free of obstructions. Only the signals received in this window are considered for interference detection in this paper's analysis. Confining z_r to this window restricts the geometry between GNSS satellites and the ISS such that $z_s \in [14.2°, 15.2°]$ (see Fig. 9). The GNSS antenna gain pattern can be assumed to be relatively constant over ±0.5°. Thus, $\hat{C}(j, f, z_s, z_r)$ can be assumed independent of z_s.

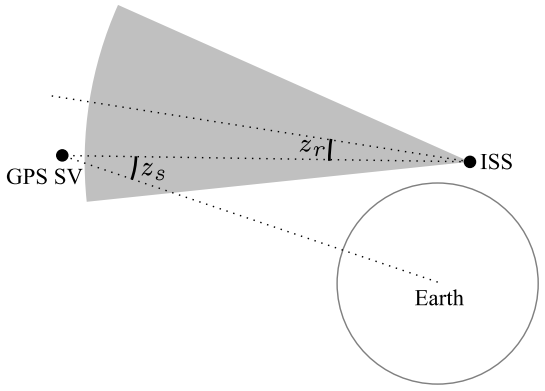

FIGURE 9 For receiver zenith angle $z_r \leq 15°$ (within the gray region), the satellite zenith angle z_s is restricted between $14.2° \leq z_s \leq 15.2°$

The mean and variance of ISS-reported range-compensated-CINR values $\hat{C}/N_0$ collected over deep ocean regions are maintained as control data in a three-dimensional grid of SV ID j, frequency band f, and receiver zenith angle z_r. For a world-wide analysis of GNSS interference events, a hypothesis test is performed on the test statistic derived from $\hat{C}/N_0$ values that fall within $z_r \in [0°, 15°]$. The test is performed separately for the L1 and L2 bands since the interference characteristics are frequency dependent. If the reported test statistics falls below $\mathbb{E}[l|H_0] - 3\sqrt{\text{Var}(l|H_0)}$, the receiver is declared to be under interference. This threshold respects a P_F of approximately 1.35×10^{-3}.

Fig. 10 shows the ratio of the number of potential interference events recorded at L1 (top panel) and L2 (bottom panel) to total number of hypothesis tests performed at each location for the foregoing detection threshold. As expected, a high ratio of potential interference events is reported for both L1 and L2 near Syria (marked with a red dot). Note that the interference "hotspot" appears to the east of the source because the ISS orbit is prograde and the FOTON antenna points in the anti-velocity direction. In other words, the FOTON antenna is exposed to interference only after the ISS passes eastward over an emitter's location.

The high values of the statistic for both L1 and L2 east of Syria indicate that the interference activity in Syria has been persistent over nearly the full interval considered in this paper, from March 2017 to June 2020. A monthly analysis (not shown) revealed that the source has been transmitting at

L2 since no later than March 2017. It began transmitting weak interference at L1 during the second half of 2017, then much stronger interference at L1 during the first quarter of 2018. The interference at L1 and L2 was ongoing in June 2020.

A weaker hotspot is present to the west of the Syrian interference. This hotspot, which emerged in the second half of 2019, is consistent with reports of GNSS interference in the Libyan region (United States Coast Guard, n.d.). The magenta dots in Fig. 10 denote the approximate location of the area in which interference has been documented (33° N, 14° E). Fig. 10 also reveals strong L2 interference over mainland China. This interference has been present since at latest March 2017 and was ongoing in June 2020. The green dot in Fig. 10 , marked at (32° N, 114° E), indicates a hypothesized interference source location based on the shape and location of the observed hotspot.

Note that the above method of counting potential interference events based on CINR degradation ignores cases where interference might lead to complete loss of track of some or all GPS signals. However, the data from the ISS shows that FOTON does not lose track of authentic GNSS signals even when flying by the strong interference source in Syria. In fact, the reported CINR over Syria is well above the weakest signal that FOTON is capable of tracking. As a result, it was concluded that in cases where FOTON seems to track few or no GPS signals, it is likely due to some abnormal behavior of the receiver, and not due to a potential interference event.

In addition to the global average analysis summarized in Fig. 10 , it is instructive to examine the time history of receiver reported CINR as the ISS passes over an interference hotspot. Fig. 11 shows two such histories for signals within the admissible range of z_r as the ISS goes over the strong interference regions in Syria (Fig. 11 (a)) and China (Fig. 11 (b)). Green and blue data points represent range-compensated CINR values for authentic L1 and L2 GNSS signals, respectively, above the applicable threshold, which depends on i, f, and z_r. Light red data points are the same data when below the applicable threshold. Both L1 and L2 signals are declared under interference in Fig. 11 (a), whereas only L2 signals are declared under interference in Fig. 11 (b). The brief dip in Fig. 11 (b) prior to the major dip over China is caused by the Syrian interference. Gaps in the time histories indicate periods with no tracked signals in the admissible zenith angle window.

5 | IMPLICATIONS FOR GNSS RECEIVERS

The matched-code interference captured over Syria is intriguing. So far as this paper's authors are aware, no other GNSS interference captured from an operational (as opposed to experimental) source has exhibited the characteristics observed in the Syrian interference. If the intent behind the signals transmitted at L1 is not spoofing but rather denial of GPS service, as can be inferred from the lack of navigation data bit modulation, then why bother transmitting an ensemble of signals, each one modulated by a separate GPS L1 C/A spreading code? For purposes of maximizing the interference noise power density I_0 in a receiver configured to track GPS L1 C/A signals, the transmitter would do just as well by allocating all its power to a single GPS L1 C/A spreading code, or any code with a similar spectral density (Humphreys, 2017). What motive can be surmised for the additional complication of transmitting a multitude of spreading codes?

The answer appears to be that the Syrian interferer is designed not only to maximize I_0 but also to efficiently disrupt cold-start acquisition of GPS L1 C/A signals, as explained below.

5.1 | Efficient Jamming

The art of jamming is more sophisticated than merely dumping RF energy into a band of interest. An efficient jammer is one that effectively disrupts GNSS service in a given area of operations but does so with as little power as possible. Such frugality extends the life of battery-powered jammers, and makes all jammers less conspicuous. The key to efficient jamming is avoiding wasteful allocation of signal power. Obviously, allocating power outside a target receiver's passband is wasteful because the interference is filtered out by the receiver's RF front-end. Less obviously, narrowband jamming applied directly in the passband is also wasteful. To understand this, consider the vector space of all possible input signals, and a partitioning into a subspace that contains the jamming signal and one that does not. If the jammer-occupied subspace is sparse with respect to the desired signal subspace, and if the receiver's front-end amplification and quantization are not saturated, then a technique can be developed to excise the jammer-occupied subspace with minimal degradation to the desired signals. For a narrowband jammer, the technique is notch filtering; for a pulsed jammer, the technique is pulse blanking (Humphreys, 2017).

An efficient jammer maximizes overlap with the desired-signal subspace for a given power allocation. Jamming that is continuous in the time domain and white (spectrally flat) within the desired signal passband in the frequency domain is fairly efficient because it extensively overlaps the desired signal subspace. But continuous-time matched-spectrum jamming is even more efficient: Instead of spreading the jamming power evenly across the passband, a matched-spectrum jammer shapes it for greater overlap with the desired signal subspace. Consider a random binary spreading code with chip

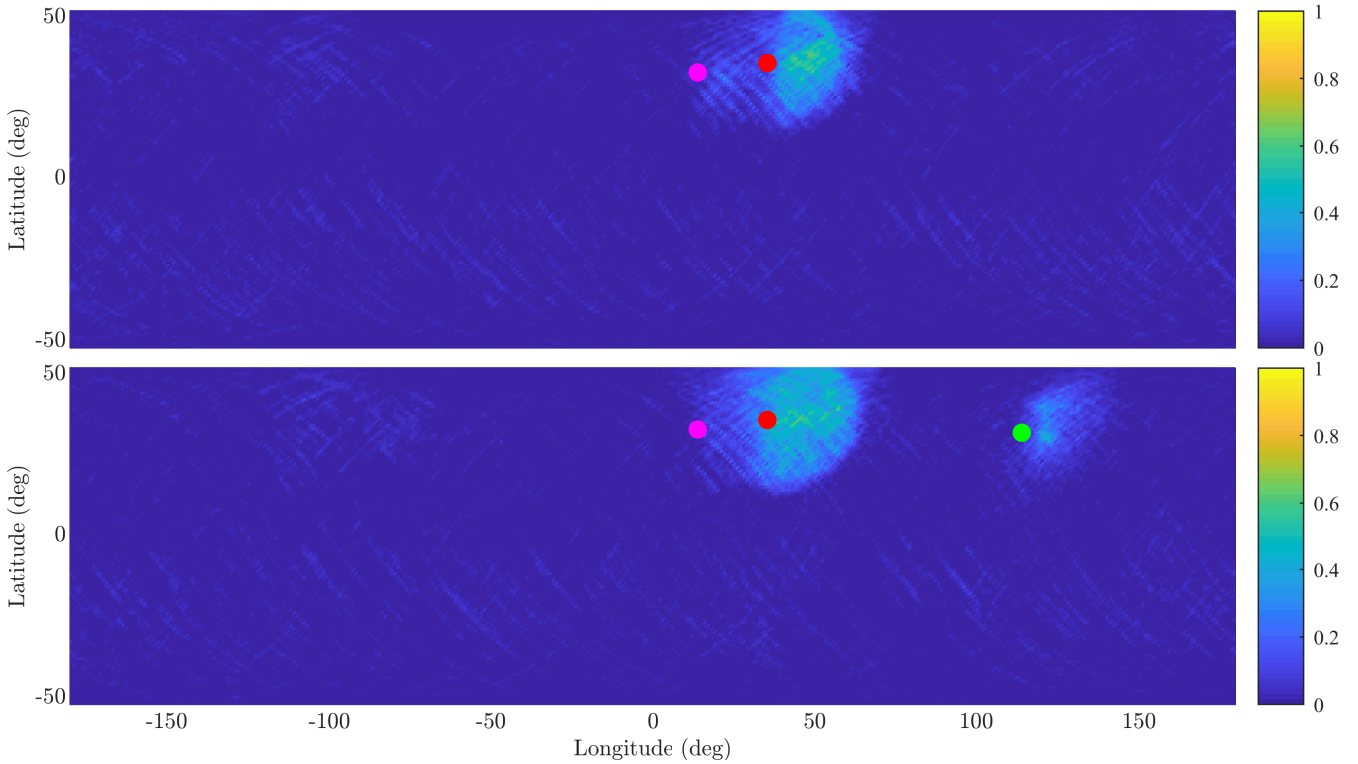

FIGURE 10 Ratio of number of potential GPS L1 (top panel) and L2 (bottom panel) interference events recorded to total number of hypothesis tests performed at each location on the map for the full span of data considered in this paper, from March 2017 to June 2020. The red dots indicate the estimated origin of the Syrian interference based on raw IF recordings. Another hotspot of interference is apparent to the west of the Syrian interference. The magenta dots denote the approximate location of GNSS interference reports in the Libyan region (United States Coast Guard, n.d.). In addition to the interference over the Syrian and Libyan regions, strong L2 interference over mainland China is observed. The green dot at (32° N, 114° E) indicates a hypothesized interference source location based on the shape and location of the observed hotspot.

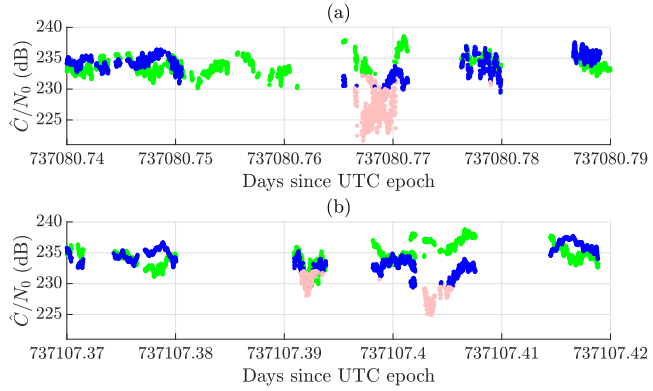

FIGURE 11 Time histories of range-compensated receiver-reported CINR as the ISS flies over potential GPS interference zones over Syria and China.

interval T_C. Suppose a spectrally-flat jammer is designed to cover the spreading code's primary spectral lobe and first two side lobes, for a total frequency span of $4/T_C$ Hz. The noise power density that passes through the receiver's matched filter is $I_0 = P_I T_C/4$, where P_I is the interference power. By contrast, for a matched-spectrum jammer $I_0 = (2/3)P_I T_C$ (Humphreys, 2017). When I_0 is large enough that CINR $\approx C/I_0$, the matched-spectrum jammer is 4.3 dB more potent than the spectrally-flat jammer. What is more, the spectrally-flat jammer spanning $4/T_C$ Hz can be excised by filtering in the frequency domain: even if the main lobe and adjacent two side lobes of the authentic signals are removed along with the jamming, the authentic signals are only attenuated by 13 dB. The spectrally-flat jammer must spread its power even wider to avoid such excision by filtering, resulting in an even less favorable potency compared with matched-spectrum jamming. By contrast, a matched-spectrum jammer cannot be excised by filtering because its spectrum follows the $\text{sinc}^2(fT_C)$ envelope of the authentic binary-code-modulated signals. By generalizing this argument, one can prove that spectrum matching is a necessary condition for efficient jamming.

However, spectrum matching is not a sufficient condition. Consider a jammer emitting a carrier modulated only by a single publicly-known spreading code of arbitrary length. This signal is sparse with respect to the desired signal subspace. It can be excised by the receiver generating a local replica of the interference signal, aligning this replica's code phase, carrier phase, and amplitude with the interference signal, and subtracting the replica from the digitized output of the receiver's RF front end. Assuming sufficient front-end bit depth and amplifier linearity, this procedure can be extended to an arbitrary number of such interference signals, each with a known waveform; the technique is known as successive interference cancellation (SIC) (Madhani, Axelrad, Krumvieda, & Thomas, 2003).

Thus, an effective jammer will not be designed to emit predictable signals; a safer approach to spectrum matching is modulation of the carrier with a non-repeating spectrum-matching spreading code known only to the jammer. But this is only true when the target receiver is capable of SIC. If, for example, the receiver has no way of distinguishing authentic from interference signals, then it cannot apply SIC without also eliminating desired signals.

5.2 | Targeting Cold Start

Under what conditions is a receiver unable to distinguish between authentic and interference signals? When (1) the authentic and interference signals are identical in all aspects of significance (modulation, code phase, carrier phase and frequency, amplitude), or (2) the authentic and interference signals are identical except in ways the target receiver is unable to exploit to distinguish them. In case (1), the interference is hardly a problem: it simply reinforces the authentic signals. Case (2) is more interesting. Let the term *spoofing interference* refer to matched-code interference with all additional modulation requisite to make the interference signal's structure and content identical to an authentic signal's. If a receiver is exposed to spoofing interference while already tracking enough authentic signals to form a navigation solution and when in possession of accurate satellite ephemerides, it can distinguish any authentic and interference signals that differ in code phase, carrier frequency, or amplitude. (It can additionally distinguish by carrier phase if performing precise carrier-based navigation.) Therefore, jamming a navigation-locked receiver with spoofing interference may be ineffective because the target receiver can apply SIC.

However, during a cold start, the target receiver's time and position are uncertain, and it lacks the ephemerides necessary to predict the code phase and Doppler of authentic signals even if its time and position were known. In this case the receiver is highly vulnerable to spoofing interference. Suppose a jammer generates a counterpart power-matched spoofing signal for each authentic GNSS signal available in an area of operations. Suppose further that the ensemble of spoofing signals is self-consistent with a location and time different from the target receiver's true location and time. On cold start, the receiver is jammed not in the traditional sense of being unable to acquire and track the authentic signals, but rather in the sense of being unable to confidently declare which of two plausible-looking navigation solutions is correct. If, under this circumstance, the receiver refuses to provide a navigation solution, the user is effectively denied GNSS service. If instead the receiver mistakenly provides the spoofed solution, the user could be exposed to hazardously misleading information.

Note that this type of spoofing interference is extremely efficient. Suppose the target receiver has a cold-start CINR acquisition threshold of η dB-Hz. Then traditional matched-spectrum jamming would require a jamming-to-authentic power ratio equal to

$$\frac{P_{\mathrm{I}}}{C} = -\left[\eta + 10\log_{10}\left(\frac{2T_{\mathrm{C}}}{3}\right)\right] \tag{9}$$

which, for GPS L1 C/A signals and a typical $\eta = 30$ dB-Hz, amounts to 31.8 dB. By contrast, jamming via single-counterpart power-matched spoofing interference requires only $P_{\mathrm{I}}/C = 0$ dB, which makes it more than 1500 times more efficient for denial of GNSS service at cold start.

5.3 | Discussion

The interference captured over Syria appears to be designed to achieve traditional matched-spectrum jamming at close range, and to disrupt cold-start acquisition far beyond this (along its line-of-sight). Indeed, it would be at least partially effective at preventing FOTON cold start even at the maximum line-of-sight range to the ISS, or approximately 1600 km. However, the interference signals as broadcast have at least four flaws, any one of which could be exploited by receivers to distinguish them from authentic signals: (1) they lack navigation data modulation; (2) they are broadcast on a (nearly) common and constant carrier frequency; (3) they share a common code phase alignment; (4) they include signals for (almost) all GPS PRNs. A receiver built to detect these anomalies could identify the imposter signals and eliminate them via SIC.

However, proper spoofing interference is not so easily distinguished from authentic signals, and is both effective and extremely power-efficient at denying GNSS service on cold start. The best defense against spoofing interference intended to deny GNSS service remains an open problem.

6 | CONCLUSIONS

Low-earth-orbiting instruments capable of receiving signals in GNSS bands are a powerful tool for characterizing GNSS interference emanating from terrestrial sources. Data from one such instrument, the FOTON software-defined GNSS receiver, which has been operational on the International Space Station since February 2017, reveal interesting patterns of GNSS interference from March 2017 to June 2020. A particularly powerful and persistent interference source active in Syria since 2017 was found to generate 72-W (EIRP) transmissions at the GPS L1 frequency containing signals modulated by all 32 GPS L1 C/A spreading codes, but with no data modulation, indicating that the signals' purpose is denial of GNSS service. A global analysis revealed other interference hotspots around the globe in both the GPS L1 and L2 frequency bands. It was argued that matched-spectrum interference is most efficient for jamming signal-locked GNSS receivers, while matched-code and especially spoofing interference are extremely power-efficient for jamming GNSS receivers during cold start.

ACKNOWLEDGMENTS

Work at The University of Texas has been supported by the National Science Foundation under Grant No. 1454474 (CAREER). Work at the Naval Research Laboratory was supported by the Chief of Naval Research. The STP-H5/GROUP-C experiment was integrated and flown under the direction of the Department of Defense Space Test Program.

References

Ala'Darabseh, E. B., & Tedongmo, B. (2019). Detecting GPS jamming incidents in OpenSky data. In *Proceedings of the 7th opensky workshop* (Vol. 67, pp. 97–108).

Amar, A., & Weiss, A. J. (2008). Localization of narrowband radio emitters based on doppler frequency shifts. *IEEE Transactions on Signal Processing*, *56*(11), 5500–5508.

Ao, C., Hajj, G., Meehan, T., Dong, D., Iijima, B., Mannucci, A., & Kursinski, E. (2009). Rising and setting GPS occultations by use of open-loop tracking. *Journal of Geophysical Research: Atmospheres (1984–2012)*, *114*(D4).

Baraniuk, R. G. (2007). Compressive sensing [lecture notes]. *IEEE signal processing magazine*, *24*(4), 118–121.

Barnes, J. A. (1972). Tables of bias functions, b_1 and b_2, for variances based on finite samples of processes with power law spectral densities. In *Precision measurement and calibration* (Vol. 5, p. 479).

Becker, K. (1992). An efficient method of passive emitter location. *IEEE Transactions on Aerospace and Electronic Systems*, *28*(4), 1091–1104.

Bhatti, J. (2015). *Sensor deception detection and radio-frequency emitter localization* (Unpublished doctoral dissertation). The University of Texas at Austin.

Bhatti, J., & Humphreys, T. E. (2017). Hostile control of ships via false GPS signals: Demonstration and detection. *Navigation*, *64*(1), 51–66.

Brimelow, B. (2018, April). *General reveals that US aircraft are being 'disabled' in Syria — the 'most aggressive' electronic warfare environment on Earth.* `https://goo.gl/9B2Nf4`

Broumandan, A., Kennedy, S., & Schleppe, J. (2020). Demonstration of a multi-layer spoofing detection implemented in a high precision gnss receiver. In *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)* (pp. 538–547).

Brown, R. G., & Hwang, P. Y. (2012). *Introduction to random signals and applied kalman filtering*. Wiley.

C4ADS. (2019, April). *Above us only stars: Exposing GPS spoofing in Russia and Syria.* `https://c4ads.org/reports`

Crassidis, J. L., & Junkins, J. L. (2011). *Optimal estimation of dynamic systems*. Chapman and Hall/CRC.

Dempster, A. G., & Cetin, E. (2016). Interference localization for satellite navigation systems. *Proceedings of the IEEE*, *104*(6), 1318–1326.

Ellis, P., Van Rheeden, D., & Dowla, F. (2020). Use of doppler and doppler rate for rf geolocation using a single leo satellite. *IEEE Access*, *8*, 12907–12920.

Griffin, C., & Duck, S. (2002). Interferometric radio-frequency emitter location. *IEE Proceedings-Radar, Sonar and Navigation*, *149*(3), 153–160.

Haworth, D., Smith, N., Bardelli, R., & Clement, T. (1997). Interference localization for EUTELSAT satellites–the first European transmitter location system. *International journal of satellite communications*, *15*(4), 155–183.

Ho, K., & Chan, Y. (1993). Solution and performance analysis of geolocation by tdoa. *IEEE Transactions on Aerospace and Electronic Systems*, *29*(4), 1311–1322.

Ho, K., & Chan, Y. (1997). Geolocation of a known altitude object from TDOA and FDOA measurements. *IEEE transactions on aerospace and electronic systems*, *33*(3), 770–783.

Humphreys, T. E. (2017). Springer handbook of global navigation satellite systems. In (pp. 469–504). Springer.

Humphreys, T. E., Murrian, M. J., & Narula, L. (2020). Deep urban unaided precise Global Navigation Satellite System vehicle positioning. *IEEE Intelligent Transportation Systems Magazine*. doi:

Humphreys, T. E., Psiaki, M. L., & Kintner, P. M., Jr. (2010, Oct.). Modeling the effects of ionospheric scintillation on GPS carrier phase tracking. *IEEE Transactions on Aerospace and Electronic Systems*, *46*(4), 1624–1637.

Jin, S., & Komjathy, A. (2010). GNSS reflectometry and remote sensing: New objectives and results. *Advances in Space Research*, *46*(2), 111–117.

John A. Volpe National Transportation Systems Center. (2001). *Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System.*

Kalantari, A., Maleki, S., Chatzinotas, S., & Ottersten, B. (2016). Frequency of arrival-based interference localization using a single satellite. In *2016 8th advanced satellite multimedia systems conference and the 14th signal processing for space communications workshop (asms/spsc)* (pp. 1–6).

Lightsey, E. G., Humphreys, T. E., Bhatti, J. A., Joplin, A. J., O'Hanlon, B. W., & Powell, S. P. (2014). Demonstration of a space capable miniature dual frequency GNSS receiver. *Navigation, Journal of the Institute of Navigation*, *61*(1), 53–64.

Madhani, P., Axelrad, P., Krumvieda, K., & Thomas, J. (2003, April). Application of successive interference cancellation to the GPS pseudolite near-far problem. *IEEE Transactions on Aerospace and Electronic Systems*, *39*(2), 481–488. doi:

Murrian, M. J., Narula, L., & Humphreys, T. E. (2019). Characterizing terrestrial GNSS interference from low earth orbit. In *Proceedings of the ion gnss+ meeting.* Miami, FL.

Pattison, T., & Chou, S. (2000). Sensitivity analysis of dual-satellite geolocation. *IEEE Transactions on Aerospace and Electronic Systems*, *36*(1), 56–71.

Psiaki, M. L., & Humphreys, T. E. (2016a, August). Attackers can spoof navigation signals without our knowledge. here's how to fight back GPS lies. *IEEE Spectrum*, *53*(8), 26-53. doi:

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

Sebastian, C. (2016, December 2). *Getting lost near the Kremlin? Russia could be GPS spoofing.* CNN.

Shepard, D. P., Humphreys, T. E., & Fansler, A. A. (2012). Evaluation of the vulnerability of phasor measurement units to GPS spoofing attacks. *International Journal of Critical Infrastructure Protection*, *5*(3-4), 146–153.

Shilong, W., Jingqing, L., & Liangliang, G. (2010). Joint FDOA and TDOA location algorithm and performance analysis of dual-satellite formations. In *2010 2nd international conference on signal processing systems* (Vol. 2, pp. V2–339).

Smith, W. W., & Steffes, P. G. (1989). Time delay techniques for satellite interference location system. *IEEE Transactions on Aerospace and Electronic Systems*, *25*(2), 224–231.

United States Coast Guard. (n.d.). *GPS problem reports status.* `https://navcen.uscg.gov/?Do=gpsreportstatus`. Accessed 2020-08-31

Van Trees, H. L. (2001). *Detection, estimation, and modulation theory*. Wiley.

Wang, P., Cetin, E., Dempster, A. G., Wang, Y., & Wu, S. (2017). Gnss interference detection using statistical analysis in the time-frequency domain. *IEEE Transactions on Aerospace and Electronic Systems*, *54*(1), 416–428.

Welch, P. (1967). The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms. *IEEE Transactions on audio and electroacoustics*, *15*(2), 70–73.

Wesson, K. D., Gross, J. N., Humphreys, T. E., & Evans, B. L. (2018, April). GNSS signal authentication via power and distortion monitoring. *IEEE Transactions on Aerospace and Electronic Systems*, *54*(2), 739-754. doi:

Wesson, K. D., & Humphreys, T. E. (2013). Hacking drones. *Scientific American*, *309*(5), 54–59.

How to cite this article: Matthew J. Murrian, Lakshay Narula, Peter A. Iannucci, Scott Budzien, Brady W. O'Hanlon, Todd E. Humphreys, GNSS Interference Monitoring from Low Earth Orbit, *NAVIGATION, Journal of the Institute of Navigation*.