ORIGINAL ARTICLE

WILEY ✦ION

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

Abstract

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 the development of interference-hardened GNSS receivers. This paper presents the results of a three-year study of global interference, with emphasis on a particularly powerful interference source active in Syria since 2017. It then explores the implications of such interference for GNSS receiver operation and design.

KEYWORDS
Doppler positioning, emitter localization, GNSS interference, spoofing

1 | INTRODUCTION

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 fast, orbital, TEC, observables, and navigation (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 (pseudorange, carrier phase, and carrier-to-noise ratio C/N_0).

Although spaceborne GNSS sensors have been used for remote sensing via radio occultation (Ao et al., 2009) and reflectometry (Jin & Komjathy, 2010), there is little public

literature exploring their use for monitoring terrestrial GNSS interference. Isoz et al. (2014) characterized interference observed at a LEO satellite, and approximated the location for one source, but was mainly concerned with determining whether the interference had a detrimental impact on GPS-RO (radio occultation) meteorological products. The more recent survey of GNSS interference localization techniques in Dempster and Cetin (2016) makes no mention of single-receiver Doppler-based localization, whether space-based or not.

General time difference of arrival (TDOA) and frequency difference of arrival (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 et al., 1997; Ho & Chan, 1993; Smith & Steffes, 1989). Application of TDOA or FDOA for localization from LEO can be viewed as an extension of such demonstrations, with the lower-altitude orbits enabling the localization of much weaker signals. Interference localization using a single satellite has been explored in Kalantari et al. (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 improves on the global survey technique in Isoz et al. (2014) by compensating for predictable C/N_0 variations in the detection test. 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, an estimate of the transmitter's location is obtained whose horizontal errors are less than 1 km with 99% confidence based on reasonable clock and noise models. Such an accurate localization of a GNSS interference source from LEO is without precedent in the open literature. Third, this paper explores the implications of interference of the type generated by the source active in Syria for GNSS receiver operation and design.

A preliminary version of this paper was published in Murrian et al. (2019). The current version focuses on the observed interference, extends the analysis period through June 2020, offers a more detailed analysis of localization accuracy, and includes a new section exploring implications for GNSS receivers.

2 | SINGLE-SATELLITE TERRESTRIAL SOURCE GEOLOCATION

As a prelude to the presentation of results from the observation campaign, this section introduces and analyzes the

Doppler-based technique employed to estimate the location of the interference source operating in Syria.

Assuming a carrier can be extracted from an interference signal, single-satellite-based transmitter geolocation is possible from Doppler measurements alone (Becker, 1992; Ellis et al., 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, $\hat{r}$ as the unit vector pointing from the transmitter to receiver, expressed in Earth-centered, Earth-fixed (ECEF) coordinates, v_R as the receiver velocity with respect to the ECEF frame and expressed in ECEF in m/s, and $\delta \dot{t}_R$ as the receiver clock frequency error in s/s, all at the time of signal receipt. Further, let $\delta \dot{t}_T$ 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:

$$f_D = -\hat{r}^T v_R / \lambda - c\left[\delta \dot{t}_R - \delta \dot{t}_T\left(1 - \delta \dot{t}_R\right)\right]/\lambda + w \quad (1)$$

where c is the speed of light in m/s. It is assumed that v_R, $\delta \dot{t}_R$, and the receiver position are known (e.g., via a GNSS receiver co-located with the transmitted signal receiver). The unknowns in Equation (1) largely stem from transmitter position, which is embedded in $\hat{r}$ and $\delta \dot{t}_T$. The former is modeled as an unknown constant and the latter as a random walk process that evolves as:

$$\dot{\delta t}_T(t_{k+1}) = \dot{\delta t}_T(t_k) + v(t_k) \quad (2)$$

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

A transmitter could introduce any level of complexity to carrier-phase frequency behavior (e.g., frequency modulation, frequency hopping). Such behaviors, if not discovered and appropriately modeled, would confound single-pass 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 enough that $\dot{\delta t}_T$ can be modeled as constant over a short (e.g., 60-second) data capture interval.

Based on the above Doppler measurement model, a batch maximum likelihood estimator (Crassidis &

Junkins, 2011) can be developed to estimate the unknown transmitter position and a constant value for $\dot{\delta t}_T$ from a collection of single-pass Doppler measurements. If Doppler measurements from multiple satellite passes are available, they can be combined for single-batch estimation provided that a new value of $\dot{\delta t}_T$ is estimated for each pass. In other words, $\dot{\delta t}_T$ is viewed as constant over each short capture interval but variable from capture to capture.

When $\dot{\delta t}_T$ 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 has been analyzed via the Monte Carlo simulation for three levels of transmitter clock quality, from a temperature-compensated crystal oscillator (TCXO) to a laboratory-grade oven-controlled crystal oscillator (OCXO). Simulation parameters were based on the real-world interference capture discussed in the next section: the true transmitter location was simulated to be 35.4 N latitude, 35.95 E longitude, 48 m 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 Equation (2). These models assume a smooth compensation for temperature control, such as is common for TCXOs used in GNSS receivers. Additionally, h_{-2} is assumed to dominate frequency stability over each short capture interval.

1,000 Monte Carlo trials were conducted for each of the three clock quality levels. Transmitter horizontal location estimation errors were observed to be zero-mean and apparently Gaussian, and they were consistent with the formal error ellipses of the associated linear least-squares estimator. To determine whether 1,000 trials were sufficient for a confident error analysis, subgroups of 250 trials were randomly selected from the 1,000 trials and each of their geolocation error ellipses were calculated.

Those subgroup samples were observed to deviate less than $\pm 10\%$ from the population mean with 99% empirical confidence. For example, 99% of the subgroups for the TCXO simulation had geolocation error ellipse estimates between $6{,}900 \pm 660$ meters for the semi-major axis, and 690 ± 71 meters for the semi-minor axis. Out of the 1×10^5 subgroup samples drawn for each clock quality level, none were observed that deviated more than 17% from the population mean. Thus, 1,000-trial-based error ellipses for each clock quality level given in Table 1 can be assumed to be

TABLE 1 Marginal contribution of transmitter frequency instability to a single-pass geolocation error ellipse; the size of the 95% horizontal geolocation error ellipse, in meters, is characterized by the (*a*) semi-major and (*b*) semi-minor axes

Clock Quality	h_{-2}	a (m)	b (m)
TCXO	3×10^{-21}	6900	690
Low-quality OCXO	3×10^{-23}	720	72
OCXO	3×10^{-25}	67	7.4

no more than 15% smaller, on either axis, than the error ellipses that would be produced in the limit of an infinite number of trials.

Table 1 shows that the marginal contribution of transmitter frequency instability to single-pass geolocation error grows precipitously with reduced transmitter clock quality. These results suggest that single-pass geolocation of a TCXO-based transmitter is marginal at best, and could be even worse if the h_{-2} values for TCXOs in Table 1 are optimistic. On the other hand, if the transmitter is driven by an OCXO-quality clock, then clock instability contributes less than 720 meters (low-quality OCXO) or 67 meters (standard-quality OCXO).

The error ellipse characterized by a and b in Table 1 is highly eccentric, with the semi-minor axis oriented in the direction of satellite motion (e.g., if the satellite is moving west to east then transmitter location would 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 techniques that can be applied generally to study terrestrial GNSS interference sources using signals collected in LEO.

Recording raw intermediate frequency (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 GPS L1 (1,575.42 MHz) and GPS L2 (1,227.6 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

FIGURE 1 Ground tracks for interference-affected captures on days 74, 144, and 151 of 2018; each capture spans approximately 60 seconds; and the estimated transmitter location is marked on the west coast of Syria [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

(Humphreys et al., 2020) enabled analysis and tracking of all radio frequency signals near GPS L1 and L2.

The following observations are based on signals captured on three days in the first half of 2018 along the ground tracks shown in Figure 1.

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 at L1, it is a wideband spread-spectrum signal. The L1 interference is a composite of individual signals with a common carrier centered near GPS L1 but each having a unique GPS L1 C/A pseudorandom number (PRN) spreading code. Such interference can be categorized as matched-code GNSS interference (Humphreys, 2017; Psiaki & Humphreys, 2020). Signals corresponding to almost all GPS L1 C/A PRN codes from 1 to 32 have been detected. When tracked by the UT software-defined GNSS receiver, 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. Moreover, the false signals are not clean simulated GPS L1 C/A signals; they exhibit unexplained fading and spectral characteristics. No false Galileo BOC(1,1) signals were detected in the L1 band.

The lack of navigation bit modulation renders the signals ineffective at spoofing, but matched-code interference is a particularly potent form of jamming (Humphreys, 2017). Why different 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

Figures 2 and 3 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 Figure 2 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 of Figure 2. Instead, it appears to be an amalgam of components. Figure 3 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 caused by transmitter idiosyncrasies is unknown.

3.3 | Baseband signal characteristics

Figure 4 shows time histories of 10-ms accumulated complex correlation products from a false (top panel) GPS L1 C/A signal and two 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
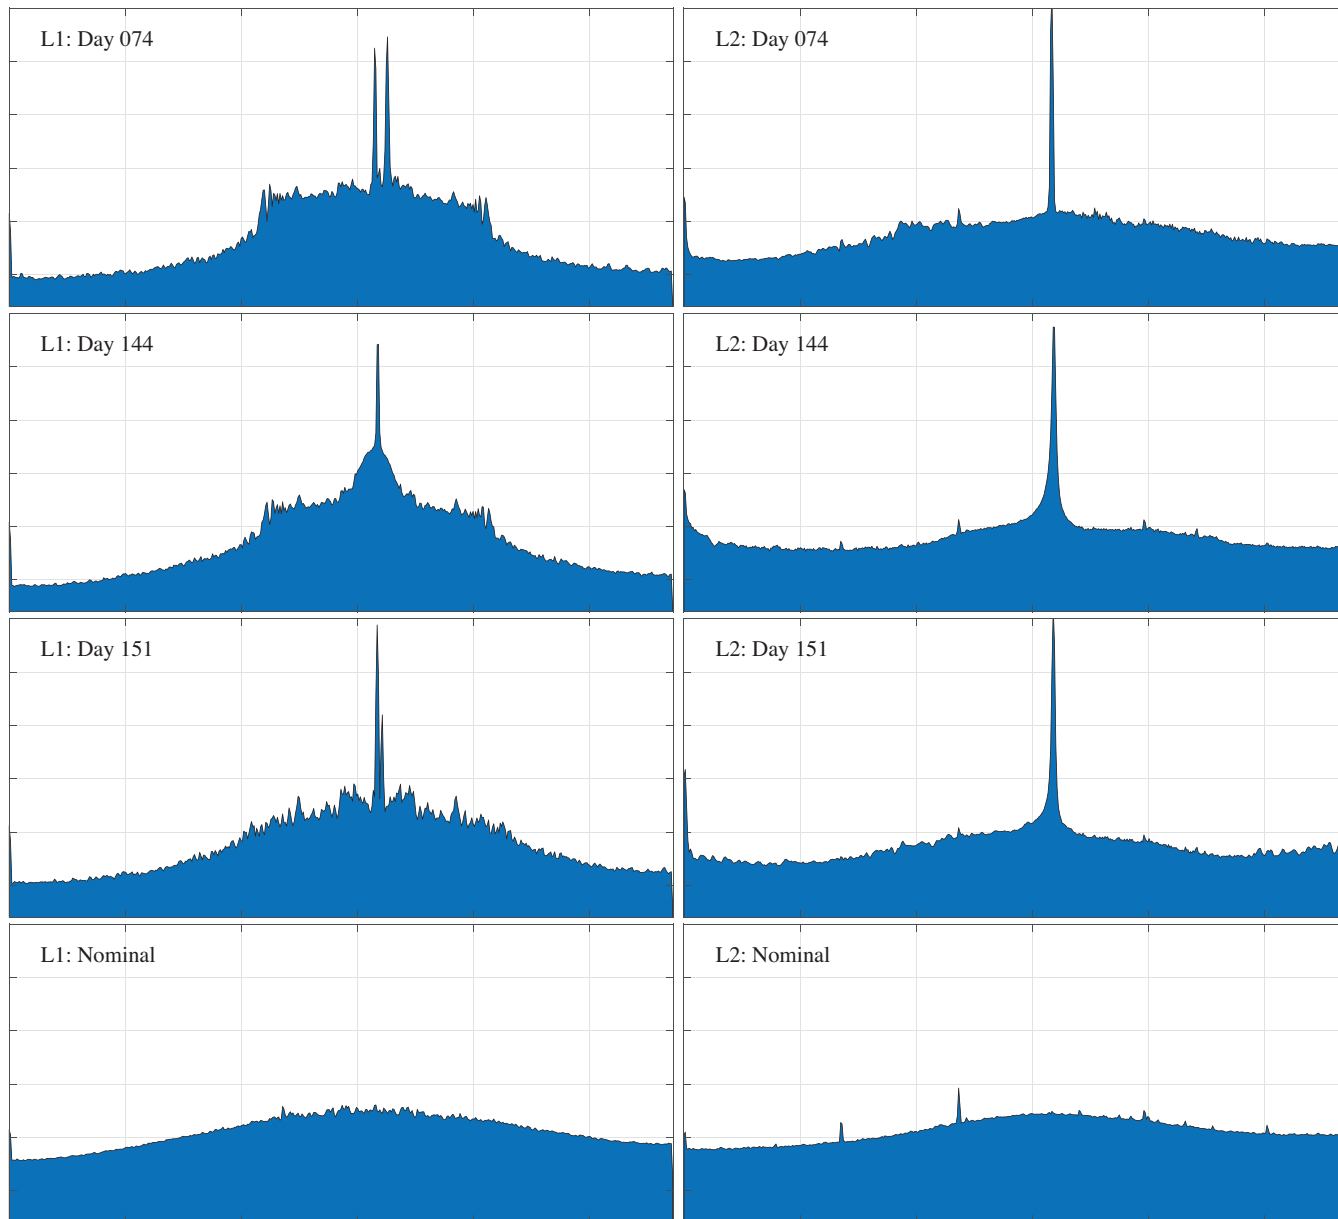
FIGURE 2 Power spectra centered near the GPS L1 (left column) and L2 (right column) frequencies from interference-affected data captured on days 74, 144, and 151 of 2018 (top three rows), and from nominal data captured on day 158 of 2018 (bottom row); the frequency span is approximately 3 MHz wide, scaled linearly with 0.5 MHz divisions; all ordinate axes are in dB and scaled equivalently for ease of comparison; and spectra are estimated by Welch's method (Welch, 1967) from 1-second data intervals with a 5.6-kHz frequency resolution [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

scintillation as the signal passes upward through the lower ionosphere (Humphreys et al., 2010), but the sudden fading, highlighted in the inset of the top panel, is unnatural and likely originates at the transmitter.

3.4 | Source geolocation

The presence of a trackable carrier signal after despreading (compare to the top panel of Figure 4) enables geolocation of the interference source as described in Section 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 v_R, $\delta \dot{t}_R$, and the receiver position component of $\hat{r}$ were then extracted from the EKF's state estimate and treated as known for purposes of source geolocation.
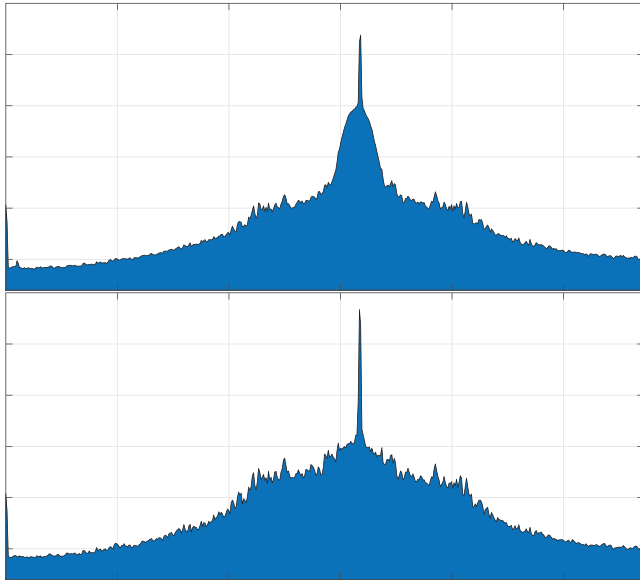
FIGURE 3 Power spectra near L1 for the day 144 capture showing maximum (top) and minimum (bottom) phases of the waxing and waning wideband (~ 0.25MHz) central interference prominence; the prominence oscillates with a period of approximately 5 seconds. The L1: Day 144 plot in Figure 2 catches the prominence waning 2 seconds after the maximum shown in the top plot [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

A batch estimator for interference source position and clock frequency bias was formulated as described in Section 2. It was assumed that the interference observed on all three days originated from the same stationary transmitter, 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, large post-fit measurement residuals could be expected to manifest, although this 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.

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, 1969) 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 $\sigma_y(2, \tau, \tau) =$
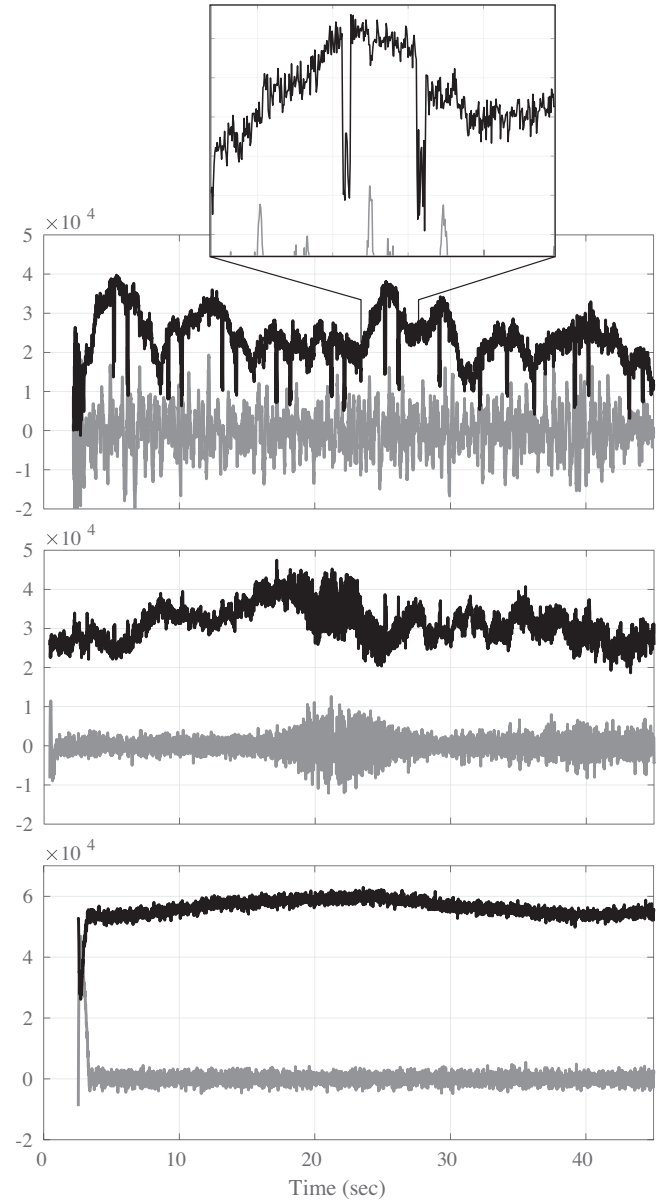
FIGURE 4 In-phase (black) and quadrature (gray) 10-ms accumulation time histories for the strongest false signal from the day 74 capture (top), the strongest authentic signal from the day 74 capture (middle), and the strongest signal from day 158 nominal capture (bottom); data wipe-off was implemented in the middle and bottom figures. 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, respectively

1.6×10^{-11}, which is consistent with a standard-quality OCXO (Bagala et al., 2016).

Thus, given the results of Table 1, treating $\delta \dot{t}_T$ as constant over each 60-second capture can be conservatively expected to introduce 95% errors smaller than 720 meters (that corresponding to a low-quality OCXO) in single-pass geolocation. A Monte-Carlo simulation like the one that
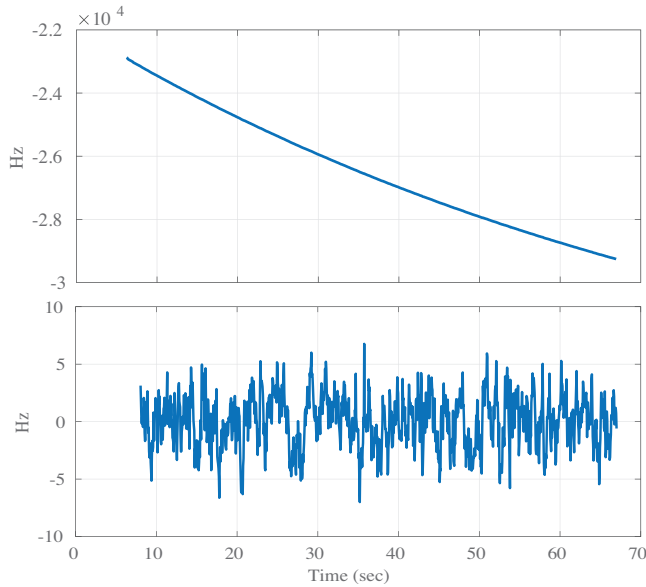
FIGURE 5 Top: Doppler time history corresponding to the false PRN 10 signal from the day 144 capture; Bottom: Post-fit residuals of the Doppler time history assuming the estimated transmitter location and clock rate offset (the standard deviation of the post-fit residuals is 2.3 Hz) [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

FIGURE 6 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. The semi-major axis of the 95% ellipse is 220 meters [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

produced the Table 1 data but for the combined three days of collection showed that, assuming independence in the clock frequency errors between passes, and conservatively assuming a low-quality OCXO, this error source can be expected to contribute 95% errors below 230 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 (1), 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.

Figure 5 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 Equation (1) and the assumption of a constant $\delta\dot{t}_T$ over each capture are reasonably accurate. Figure 6 shows the estimated position of the interference source. 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 (1). 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. While the true location is not known, the geolocation solution based on the model plausibly coincides with a Russian-operated air base in Syria.

3.5 | Transmitter power

In the presence of interference, C/N_0 actually measures the carrier-to-interference-and-noise ratio (CINR). By analyzing the authentic signal CINR values in the captured data, one can infer the transmitted power in the direction toward the ISS of the emitter located in Syria. The data presented here are for the day 74 capture. The average decrease in the CINR values observed at the ISS when 1,340 km from the source was approximately 6 dB.

One may 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 $P_I = -137$ dBW. Let $P_S = P_I - G_r + L$, and assume path loss $L = 159$ dB, consistent with a stand-off distance of 1,340 km, and receiver antenna gain $G_r = 3$ dB.

It follows that the transmitter power of the interference source in the direction toward the ISS during the day 74 capture is $P_S = 19$ dBW, or 79 W. If the transmitter is focused on ground-based targets, then it is possible that the

gain pattern is toroidal. The elevation angle of the ISS as seen from the transmitter is low during this period (varying between 8 and 13.5 degrees) and may have been near the maximum of a toroidal gain pattern.

4 | GLOBAL INTERFERENCE SURVEY VIA RECEIVER-REPORTED CINR

The raw IF data captures from the ISS FOTON receiver enabled 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 worldwide survey of strong GNSS interference.

4.1 | Calculation of receiver-reported CINR

Receiver-reported CINR is calculated as:

$$\text{CINR} = \left(\frac{\mathbb{E}[I^2 + Q^2]}{2\sigma_{IQ}^2} - 1 \right) \frac{1}{T_a} \qquad (3)$$

where the expectation $\mathbb{E}[I^2 + Q^2]$ is estimated by moving average using a Euler approximation to a standard low-pass filter:

$$\mathbb{E}[I^2 + Q^2]_k = \mathbb{E}[I^2 + Q^2]_{k-1}$$
$$+ K\left(I_k^2 + Q_k^2 - \mathbb{E}[I^2 + Q^2]_{k-1}\right) \qquad (4)$$

with subscripts k and $k-1$ indicating the current and previous accumulation interval. The gain parameter $K = \frac{T_a}{\tau}$ with accumulation interval $T_a = 10$ msec and filter time constant $\tau = 0.5$. I_k and Q_k are the in-phase and quadrature prompt correlation products for the current accumulation interval.

The receiver noise floor, $2\sigma_{IQ}^2$, can be derived analytically for a 2-bit quantizing RF front-end and a software-defined GNSS receiver based on the quantization models of both the RF front-end and receiver local carrier replica. It can be shown that:

$$2\sigma_{IQ}^2 = 2N\left(a_0^2 p_{a0} + a_1^2 p_{a1}\right)\left(b_0^2 p_{b0} + b_1^2 p_{b1}\right) \qquad (5)$$

where N is the number of samples per accumulation interval, a_0, a_1, b_0, and b_1 are the low and high quantization values of the RF front-end and local carrier replica, respectively, and p_{a0}, p_{a1}, p_{b0}, and p_{b1} are their associated probabilities, respectively.

In practice, p_{a0} and p_{a1} depend on the implementation of the automatic gain control (AGC) in the RF front-end, and p_{b0} and p_{b1} are selected by the receiver designer (e.g., to minimize quantization distortion). The following values are applicable to the FOTON receiver:

$$a_0 = 1, a_1 = 3, b_0 = 1, b_1 = 3$$

$$p_{a0} = 0.68269, p_{a1} = 0.31731, p_{b0} = 0.38418, p_{b1} = 0.61582$$

$$N = 5714.286$$

which yield a noise floor of $2\sigma_{IQ}^2 = 239669$ front-end units.

4.2 | Methodology

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 off-boresight angle); and z_r is the angle between the ISS antenna boresight direction and the direction to the satellite (receiver antenna off-boresight angle).

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, an off-boresight 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 $z_s \in$
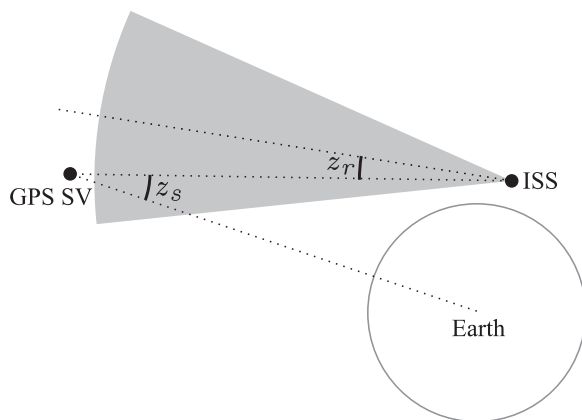
FIGURE 7 For receiver off-boresight angle $z_r \leq 15°$ (within the gray region), the satellite off-boresight angle z_s is restricted between $14.2° \leq z_s \leq 15.2°$

$[14.2°, 15.2°]$ (see Figure 7). The GNSS antenna gain pattern can be assumed to be relatively constant over $\pm 0.5°$. Thus, $\hat{C}(j, f, z_s, z_r)$ can be assumed independent of z_s.

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 off-boresight angle z_r. For a worldwide 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 fall below $\mathbb{E}[l|H_0] - 3\sqrt{\mathrm{Var}(l|H_0)}$, the receiver is declared to be under interference. This threshold respects a P_F of approximately 1.35×10^{-3}.

4.3 | Discussion of results

Figure 8 shows the ratio of the number of potential interference events recorded at L1 (top panel) and L2 (bottom panel) to the 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
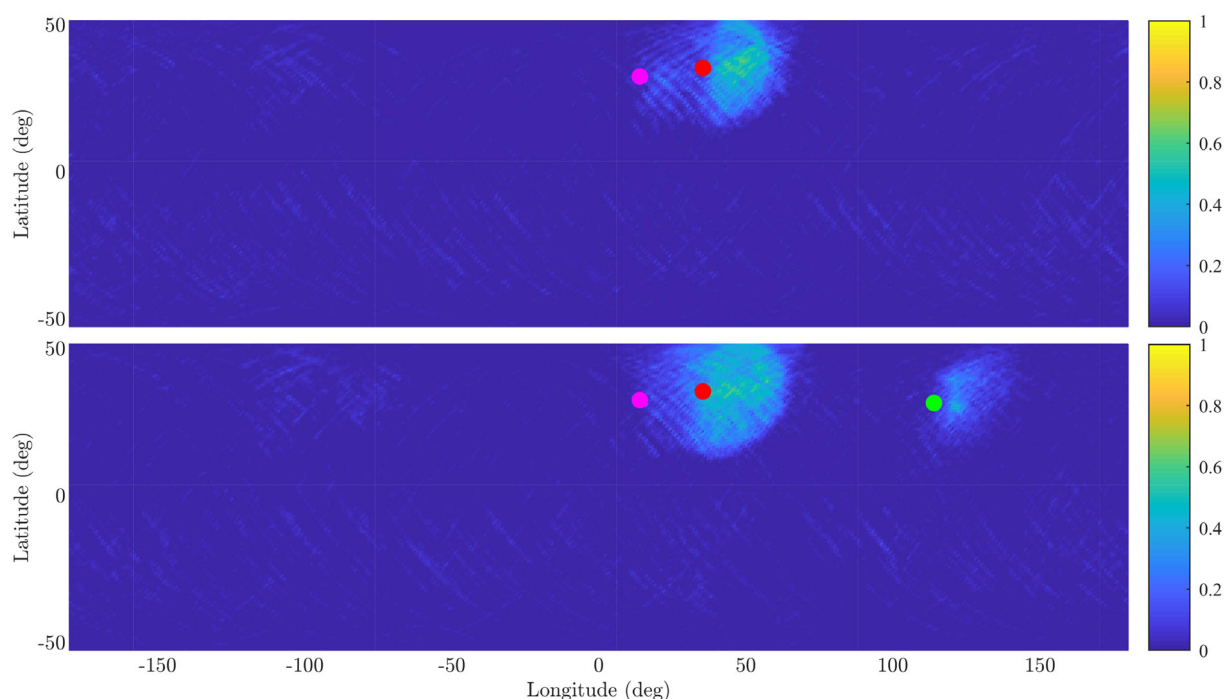
FIGURE 8 Ratio of number of potential GPS L1 (top panel) and L2 (bottom panel) interference events recorded to the total number of hypothesis tests performed at each location on the map for the full span of data considered in this paper, from March 2017 to June 2020. The red dots indicate the estimated origin of the interference from Syria based on raw IF recordings. Another hotspot of interference is apparent to the west of the Syrian location. The magenta dots denote the approximate location of GNSS interference reports in the Libyan region (United States Coast Guard). 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 [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

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 as of June 2020.

An additional hotspot is present to the west of the Syrian location. 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 Figure 8 denote the approximate location of the area in which interference has been documented (33° N, 14° E). Figure 8 also reveals strong L2 interference over mainland China. This interference has been present since March 2017 at the latest and was ongoing in June 2020. The green dot in Figure 8, 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 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 Figure 8, it is instructive to examine the time history of receiver-reported CINR as the ISS passes over an interference hotspot. Figure 9 shows two such histories for signals within the admissible range of z_r as the ISS goes over the strong interference regions in Syria [Figure 9(a)] and China [Figure 9(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 Figure 9(a), whereas only L2 signals are declared under interference in Figure 9(b). The brief dip in Figure 9(b) prior to the major dip over China is caused by the interference originating in Syria. Gaps in the time histories indi-
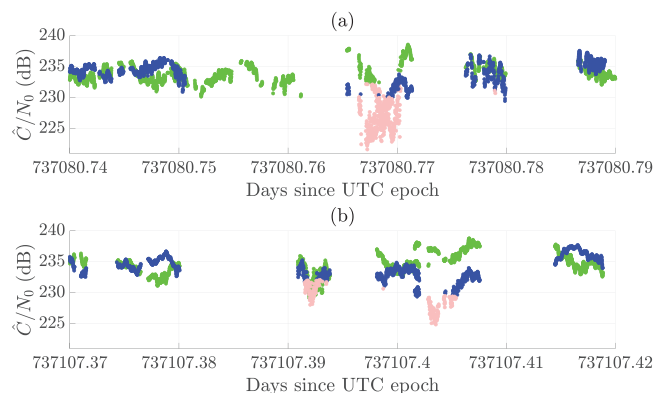
FIGURE 9 Time histories of range-compensated receiver-reported CINR as the ISS flies over potential GPS interference zones over Syria and China [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com and www.ion.org]

cate periods with no tracked signals in the admissible off-boresight 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 interference emanating from Syria. If the intent behind the signals transmitted at L1 is not spoofing, but rather denial of GPS service, as might be inferred from the lack of navigation data bit modulation, then it is unclear why an ensemble of signals, each one modulated by a separate GPS L1 C/A spreading code, was transmitted. The transmitter could just as well allocate its power to a single GPS L1 C/A spreading code, or any code with a similar spectral density. However, transmitting a multitude of spreading codes can be effective at disrupting 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 emitting RF energy into a target band. 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 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 reasonably efficient because it extensively overlaps the desired signal subspace. 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 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 to matched-spectrum jamming. By contrast, a matched-spectrum jammer cannot be excised by filtering because its spectrum follows the $\mathrm{sinc}^2(fT_C)$ envelope of the authentic binary-code-modulated signals. Thus, spectrum matching is a necessary condition for efficient jamming.

However, spectrum matching is not a sufficient condition for effective jamming. 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 et al., 2003).

Thus, an effective jammer will avoid predictable signals: a more sophisticated 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 necessary when the target receiver is capable of SIC. If, for example, the receiver has no way of distinguishing authentic signals from interference signals, then it cannot apply SIC without also eliminating desired signals.

5.2 | The cold start vulnerability

It is instructive to consider the conditions under which a GNSS receiver is unable to distinguish between authentic and interference signals: (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 a 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 hazardous, misleading information.

This type of spoofing interference is highly 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:

$$\frac{P_I}{C} = -\left[\eta + 10 \log_{10}\left(\frac{2T_C}{3}\right)\right] \qquad (6)$$

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 $P_I/C = 0$ dB, which makes it more than 1,500 times more efficient for denial of GNSS service at cold start.

5.3 | Discussion

The interference captured over Syria causes traditional matched-spectrum jamming at close range, and is capable of disrupting 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 1,600 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 common and constant carrier frequency; (3) they share a common code phase alignment; and (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, in general, spoofing interference is not so easily distinguished from authentic signals, and can be both effective and 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. Signals from a particularly powerful and persistent interference source active in Syria since 2017 were captured and characterized, and the source was geolocated to better than 1 km. A global analysis revealed other interference hotspots around the globe in both the GPS L1 and L2 frequency bands. Matched-code interference such as emitted at the GPS L1 frequency by the jammer in Syria is power-efficient for jamming signal-locked GNSS receivers. GNSS receivers are particularly vulnerable to such interference during cold start.

ORCID

Matthew J. Murrian ⓘ https://orcid.org/0000-0002-5299-8904
Mark L. Psiaki ⓘ https://orcid.org/0000-0002-3070-1217

REFERENCES

Amar, A., & Weiss, A. (2008). Localization of narrowband radio emitters based on doppler frequency shifts. *IEEE Transactions on Signal Processing*, *56*(11), 5500–5508. https://doi.org/10.1109/TSP.2008.929655

Ao, C. O., Hajj, G. A., Meehan, T. K., Dong, D., Iijima, B. A., Mannucci, A. J., & Kursinski, E. R. (2009). Rising and setting GPS occultations by use of open-loop tracking. *Journal of Geophysical Research*, *114*(D4). https://doi.org/10.1029/2008jd010483

Bagala, T., Fibich, A., Kubinec, P. & Stofanik, V. (2016). Improvement of short-term frequency stability of the chip scale atomic clock. *2016 IEEE International Frequency Control Symposium (IFCS)*. IEEE. https://doi.org/10.1109/fcs.2016.7546746

Barnes, J. A. (1969). *Tables of bias functions, B₁ and B₂, for variances based on finite samples of processes with power law spectral densities*. National Bureau of Standards. https://doi.org/10.6028/nbs.tn.375

Becker, K. (1992). An efficient method of passive emitter location. *IEEE Transactions on Aerospace and Electronic Systems*, *28*(4), 1091–1104. https://doi.org/10.1109/7.165371

Bhatti, J. (2015). *Sensor deception detection and radio-frequency emitter localization* [Ph.D. thesis, The University of Texas at Austin].

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

Crassidis, J. L., & Junkins, J. L. (2011). *Optimal estimation of dynamic systems*. Chapman and Hall/CRC. https://doi.org/10.1201/b11154

Dempster, A. G., & Cetin, E. (2016). Interference localization for satellite navigation systems. *Proceedings of the IEEE*, *104*(6), 1318–1326. https://doi.org/10.1109/jproc.2016.2530814

Ellis, P., Rheeden, D. V., & Dowla, F. (2020). Use of Doppler and Doppler rate for RF geolocation using a single LEO satellite. *IEEE Access*, *8*, 12907–12920. https://doi.org/10.1109/access.2020.2965931

Griffin, C., & Duck, S. (2002). Interferometric radio-frequency emitter location. *IEE Proceedings - Radar, Sonar and Navigation*, *149*(3), 153. https://doi.org/10.1049/ip-rsn:20020278

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. https://doi.org/10.1002/(SICI)1099-1247(199707/08)15:4⟨155::AID-SAT577⟩3.0.CO;2-U

Ho, K., & Chan, Y. (1993). Solution and performance analysis of geolocation by TDOA. *IEEE Transactions on Aerospace and Electronic Systems*, *29*(4), 1311–1322. https://doi.org/10.1109/7.259534

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. https://doi.org/10.1109/7.599239

Humphreys, T. E. (2017). Interference. In Peter J.G., Teunissen & Oliver M. (Eds.), *Springer Handbook of Global Navigation Satellite Systems* (pp. 469–503). Springer International Publishing.

Humphreys, T. E., Murrian, M. J., & Narula, L. (2020). Deep-urban unaided precise global navigation satellite system vehicle positioning. *IEEE Intelligent Transportation Systems Magazine*, *12*(3), 109–122. https://doi.org/10.1109/MITS.2020.2994121

Humphreys, T. E., Psiaki, M. L., & Kintner, P. M. (2010). Modeling the effects of ionospheric scintillation on GPS carrier phase tracking. *IEEE Transactions on Aerospace and Electronic Systems*, *46*(4), 1624–1637. https://doi.org/10.1109/taes.2010.5595583

Isoz, O., Buehler, S. A., Kinch, K., Bonnedal, M., & Akos, D. M. (2014). Interference from terrestrial sources and its impact on the GRAS GPS radio occultation receiver. *Radio Science*, *49*(1), 1–6. https://doi.org/10.1002/2013rs005243

Jin, S., & Komjathy, A. (2010). GNSS reflectometry and remote sensing: New objectives and results. *Advances in Space Research*, *46*(2), 111–117. https://doi.org/10.1016/j.asr.2010.01.014

Kalantari, A., Maleki, S., Chatzinotas, S. & Ottersten, B. (2016). Frequency of arrival-based interference localization using a single satellite. *2016 8th Advanced Satellite Multimedia Systems Conference and the 14th Signal Processing for Space Communications Workshop (ASMS/SPSC)*. https://doi.org/10.1109/ASMS-SPSC.2016.7601472

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*, *61*(1), 53–64. https://doi.org/10.1002/navi.52

Madhani, P., Axelrad, P., Krumvieda, K., & Thomas, J. (2003). Application of successive interference cancellation to the GPS pseudolite near-far problem. *IEEE Transactions on Aerospace and Electronic Systems*, *39*(2), 481–488. https://doi.org/10.1109/taes.2003.1207260

Murrian, M. J., Narula, L. & Humphreys, T. E. (2019). Characterizing terrestrial GNSS interference from low earth orbit. *Proc. of the 32nd International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2019)*, Miami, FL, 3239–3253. https://doi.org/10.33012/2019.17065

Pattison, T., & Chou, S. (2000). Sensitivity analysis of dual-satellite geolocation. *IEEE Transactions on Aerospace and Electronic Systems*, *36*(1), 56–71. https://doi.org/10.1109/7.826312

Psiaki, M. L. & Humphreys, T. E. (2020). Civilian GNSS spoofing, detection, and recovery. In Y. T. Jade Morton, Frank van Diggelen, James J. Spilker Jr., Bradford W. Parkinson, Sherman Lo, & Grace Gao's (Eds.), *Position, navigation, and timing technologies in the 21st century: Integrated satellite navigation, sensor systems, and civil applications, volume 1*. The Institute of Electrical and Electronics Engineers, Inc. https://doi.org/10.1002/9781119458449.ch25

Smith, W., & Steffes, P. (1989). Time delay techniques for satellite interference location system. *IEEE Transactions on Aerospace and Electronic Systems*, *25*(2), 224–231. https://doi.org/10.1109/7.18683

United States Coast Guard (n.d). GPS problem reports status. https://navcen.uscg.gov/?Do=gpsreportstatus

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. https://doi.org/10.1109/tau.1967.1161901