Orbital War Driving: Assessing Transient GPS Interference from LEO

Daniel M. LaChapelle*, Lakshay Narula†, Todd E. Humphreys*
*Radionavigation Laboratory, The University of Texas at Austin
†Amazon Lab126

BIOGRAPHIES

Daniel LaChapelle (BS, Mechanical Engineering, Cornell University; MS, Aerospace Engineering, The University of Texas at Austin) is a PhD student with the Department of Aerospace Engineering and Engineering Mechanics, and a Graduate Research Assistant in the University of Texas at Austin Radionavigation Laboratory.

Lakshay Narula received the B.Tech. degree in electronics engineering from IIT (BHU), Varanasi, India, and the M.S., and Ph.D. degrees in Electrical and Computer Engineering from The University of Texas at Austin. He is currently an Applied Scientist with Amazon Lab126. His research interests include robust state estimation and sensor fusion. He was the recipient of the Qualcomm Innovation Fellowship in 2017 and the IEEE Walter Fried Award in 2020.

Todd Humphreys (BS, MS, Electrical Engineering, Utah State University; PhD, Aerospace Engineering, Cornell University) is a professor in the department of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, where he directs the Radionavigation Laboratory. He specializes in the application of optimal detection and estimation techniques to problems in secure, collaborative, and high-integrity perception, with an emphasis on navigation, collision avoidance, and precise timing. His awards include The University of Texas Regents' Outstanding Teaching Award (2012), the National Science Foundation CAREER Award (2015), the Institute of Navigation Thurlow Award (2015), the Qualcomm Innovation Fellowship (2017), the Walter Fried Award (2012, 2018), and the Presidential Early Career Award for Scientists and Engineers (PECASE, 2019). He is a Fellow of the Institute of Navigation and of the Royal Institute of Navigation.

ABSTRACT

Low Earth orbit provides a unique vantage point for observing GNSS interference: it is close enough to the source of the interference for a single sensor to characterize the strength, power spectra, and signal content of terrestrial jamming and spoofing sources, but far enough that authentic GNSS signals may still be tracked and navigation solutions computed. These observations permit geo-referenced characterization of terrestrial GNSS interference, which is an important step on the way to understanding the extent of the phenomenon and developing mitigation strategies. By working directly on 100 Hz data-wiped complex IQ correlation products instead of lower-frequency receiver products, it is possible to identify interference with greater sensitivity, permitting detection not long after it is first received. Successive ground passes of the Fast, Orbital, TEC, Observables, and Navigation (FOTON) receiver aboard the International Space Station (ISS) form an impressively complete GNSS interference survey of the globe at latitudes below the ISS inclination of 51.6 degrees.

INTRODUCTION

Global Navigation Satellite Systems (GNSS) signals are relied upon for a number of safety critical applications where there is a need for precise localization or clock synchronization. Due to the low strength of the signals at point-of-use, they are easily overwhelmed by RF interference — malicious or unintentional. This interference may simply deny a navigation and timing solution; it may also induce inauthentic solutions unbeknownst to the victim. Identifying when and where a GNSS receiver is affected by interference is an important first step towards locating and mitigating the interference itself. The work presented in this paper takes advantage of the receiver's rather unique platform — the International Space Station (ISS) — to detect GNSS interference as it is occurring. This task is aided by the fact that the ISS is one of the most-observed spacecraft presently in orbit; even if the navigation solution were severely degraded, a position estimate can be obtained from regularly updated public ephemerides.

Preprint of the 2021 ION GNSS+ Conference St. Louis, MO, Sept. 20-24, 2021

The strength of a signal interfering with GNSS is also its Achilles' heel: it is easily heard, provided one is listening. A number of recent efforts have exploited this fact to monitor GNSS interference across the globe. One such approach takes advantage of the public Automatic Dependent Surveillance-Broadcast (ADS-B) used for air traffic control [1]. Interference was recognized as irregularities in ADS-B reports, which are collected by the community receiver network OpenSky. In fact, the authors of [1] were able to provide an estimate of the interference source's location by noting the effects on multiple flight paths and convex optimization techniques. This approach has its limits: the OpenSky Network's receivers are primarily (although not totally) located in populated areas with reliable internet access, with little coverage over oceans. Furthermore, the use of ADS-B means that it has limited coverage over conflict areas where overflights are rare but GNSS interference is likely to be present [2]. On the other hand, this approach is vindicated by the observation that GNSS interference tends to be inordinately powerful, affecting aviation up to 300 km from its estimated origin [3].

Commercial efforts have also entered the business of interference monitoring: hosted payloads managed by Aireon aboard Iridium NEXT satellites monitor ADS-B transmissions from orbit, enabling Aireon to observe areas not covered by OpenSky [4]. Yet another project takes advantage of data from over 500 reference receivers collected by organizations under the aegis of the International GNSS Service (IGS) [5]. It is desirable to search for GNSS interference in a way that is restricted to neither commercial flight paths nor the vicinity of reference stations. HawkEye360 is attempting to do exactly that: as a demonstration mission, three *Pathfinder* spacecraft launched in late 2018 were used to geolocate land-based reference signals [6], [7]. The Pathfinder spacecraft were placed in a Sun-synchronous orbit (SSO) at an altitude of 575 km, spaced apart by 100 to 200 km. Onboard software-defined receivers are capable of 144 MHz to 15 GHz; GNSS signals are well within this range in the L band (1 to 2 GHz). At the time of writing, HawkEye 360 has yet to publish its attempts to geolocate GNSS interference sources.

War driving is a colloquial term for the practice of searching for unsecured wireless access points, which are then mapped for later use. As the name implies, this exercise is facilitated by a car, enabling the war drivers to cover more ground [8]. This work borrows the term as a metaphor for a slightly different application. Historically, war driving used GNSS geolocation to map Wi-Fi networks; here, the GNSS receiver itself is used to identify radio signals of interest. In this case, however, the signals do not originate from harmless, low-power wireless routers; rather, they are broadcast by high-power spoofing and jamming equipment that has the potential to disrupt GNSS-derived positioning, navigation, and timing (PNT) over a large geographic area. The strategies most effective at detecting and mitigating this interference depend on the (generally unknown) strategy used by the spoofer [9]. A simple, widely-implemented strategy requires the GNSS receiver to report when interference is present, prompting the user to discard the navigation solutions as invalid. This approach generally does not permit the user to entirely ignore the interference, as it can easily be strong enough to overwhelm the relatively weak GPS signal (at the Earth's surface). Being an eavesdropper in LEO, the FOTON receiver isn't itself led astray by spoofing attacks, although careful processing of raw captures can elicit the structure of the interference signal [10, §3].

There is also the possibility, not yet publicly demonstrated, that clever terrestrial interference could target an unsuspecting satellite in LEO. Such an attack could be particularly damaging to a satellite that relies on precise localization or timing for its mission, and not require a transmitter significantly stronger than those presently used to interfere with terrestrial receivers. Spacecraft-targeting interference would present an insurmountable challenge for the work presented here; one could then no longer make the assertion that the receiver is tracking authentic signals. However, spaceborne GNSS receivers in LEO are aided by the fact that any interference source on the Earth's surface will quickly fall out of sight as the spacecraft continues in its orbit.

Previous work in this area identified and located persistent sources of GNSS interference that were present over many ground passes [10, §4]. However, much GNSS interference is transient with a duration of only hours or days. Despite its shorter duration, transient interference can compromise the safety of GNSS-dependent systems, with implications for injury and economic damage.

Guided by NOTAMs, ADS-B data, and news reports, among other sources, it is possible to narrow down the search for transient GNSS interference [1]. This paper's technical approach involves tuning the detection tests using these known sources of interference. Armed with a highly sensitive means to detect interference, it will be possible to identify heretofore unknown transient sources of GNSS interference over much of the globe from 2017 to the present. Additionally, this work has the advantage of a data set beyond compare: three years of 100 Hz data-wiped complex IQ correlation products captured by the Fast, Orbital, TEC, Observables, and Navigation (FOTON) receiver aboard the International Space Station (ISS).

There are two basic interference categories that this analysis can expect to find: narrowband and wideband. The latter may be due to spoofing or matched-code interference. Narrowband interference is the simplest and most common form of GNSS interference. It entails broadcasting a narrowband waveform in navigation bands for denial-of-service (DOS) purposes. On the other hand, a transmitter "spoofing" a GNSS signal broadcasts a counterfeit signal intended to deceive the recipient into thinking it is authentic. Somewhere in between these two techniques is matched-code interference, in which a GNSS satellite's pseudorandom ranging codes are broadcast sans navigation message; this is intended to fool receivers into acquiring the signal, but deny them a navigation solution. It is also possible that unintentional interference may be detected; for example, radio signals produced by malfunctioning electrical equipment or natural phenomena like solar radio bursts [11]. However, the former is less likely than intentional interference to be powerful enough to detect from LEO, and the latter is relatively rare. Nevertheless, it is important to rule out unintentional or natural interference to the extent possible before casting aspersions.

A simple technique by which GNSS interference can be detected is that of monitoring the carrier-to-noise ratio, C/N_0, also referred to as the carrier-to-interference-and-noise ratio (or CINR) in the presence of interference. GNSS interference manifests as a decrease in the C/N_0 of an authentic signal by a magnitude unlikely to be caused by multipath, the typical source of C/N_0 variation. Once a likely interference source has been detected, it may even be possible to narrow down its location on the globe by cross-referencing interference episodes with navigation solutions. The estimated carrier-to-noise ratio for each tracked signal is regularly reported by the receiver along with its navigation solutions, but this work has access to lower-level, higher-frequency, more sensitive receiver outputs: the complex IQ correlation products. Making the assumption that phase error remains negligible, the I (in-phase) component may be substituted for the C/N_0 (Section).

This work (i) extends existing methods for GNSS interference detection via C/N_0 monitoring from standard GNSS observables to higher-sampling-frequency complex correlation product data, and (ii) evaluates these methods on likely persistent instances of GNSS interference originating from hotspots identified in previous work.

RELATED WORK

The problem of identifying GNSS interference from data available to the receiver can be thought of as a special case of the more general anomaly detection problem; that is, determining whether or not observations of a stochastic process correspond to an expected (nominal) model, or an alternative (anomalous) model. Anomaly detection has been studied extensively [12]–[14]. More general approaches ([13], [14]) may be appropriate if the data domain is not amenable to analytical modeling. In the case of detecting GNSS interference from a receiver mounted on a spacecraft — especially one as well-studied as the ISS — it is possible for the model to incorporate patterns that may otherwise be classified as anomalies by a more domain-agnostic algorithm. One example of such features in the context of C/N_0 monitoring is the regular occultation of the GPS satellite vehicles (SVs) by the Earth. There are also features of the data that may be identified as anomalies, but are not of interest to this work. For example, rapid signal fading due to ionospheric scintillation [11], [15] is a natural phenomenon that is not as easily modeled. These features must be identified or ruled out in some other way (see Section).

A common limitation imposed on anomaly detection techniques is the restriction to a subset of the data, typically in the context of real-time detection [14]. As this paper studies historical data, the analyses herein are not constrained in this manner; however, the sheer quantity of data — years of 100 Hz data for each tracked signal — necessitates examining the data in segments.

There is strong interest in developing low-cost methods of detecting GNSS interference in order to (i) alert users that the navigation solutions may not be valid and (ii) if possible, recover the authentic solution. As (ii) is generally not a concern for the ISS, only methods to perform (i) are needed. A "low cost" solution is one that leverages quantities that are observable to a typical GNSS receiver installation, in comparison to those available only with specialized hardware [16]. Some common interference-detection metrics available to a typical GNSS receiver are:

1) Carrier-to-Noise (density) ratio, C/N_0 [11], [15], [17]
2) Received power (AGC gain) [11], [15], [17]–[19]
3) Spectral analysis [11], [17]
4) Number of observed signals [17]
5) Correlator output power [18]
6) Correlator output power variance [18]
7) Carrier phase vacillation [18]

8) Pseudorange outliers [20]
9) Signal quality monitoring (SQM) [15], [19]
10) Complex ambiguity function monitoring [21]

While the presence of an anomaly in one of these metrics can suggest interference, some successful detection strategies use more than one in order to improve the probability of detection or discriminate between different interference types [15], [18].

A problem related to anomaly detection is that of *quickest detection*, also referred to as quickest change detection (QCD) [22]. The goal of quickest detection is to identify a sudden statistical change in an observed signal with minimal detection delay. A key assumption made in quickest detection theory is that the duration of this change is effectively infinitely long — a change occurs only once in the signal's time history. More relevant to this work is transient change detection (TCD) theory, which studies to the case in which the change occurs only over an interval. A similar approach to the same underlying problem is termed *offline change point detection* [23]. Recently-published work in TCD has been adapted to fit this problem and forms the core of the statistical framework used to detect GNSS interference events (Section).

TECHNICAL APPROACH

Overview

The problem of finding possible interference events in complex IQ correlation products can be simplified to that of identifying a negative shift in the mean of a random series. The mean-shift problem has a long history; solutions may be found in the literature as early as 1955 [24], [25].The identification may be performed offline, in a retrospective manner. In theory, this would permit analyzing years of data at once; however, practical constraints on computer memory limit inspection to only one day's IQ data at a time.

The total number of interference events — or other mean-shifting anomalies — in a given range of data is *a priori* unknown. However, one can conclude that each interference events will consist of two mean shifts, a negative shift at the start of the event and a positive shift of a similar magnitude at the end. Under the assumption that only a single interference source is present at a time, it is possible to bound the expected length of the shift. It is assumed that interference sources are terrestrial and (approximately) stationary relative to the orbital motion of the ISS: the likelihood of spaceborne GNSS jamming and spoofing is considered to be negligible.

A complicating factor is imposed by the nature of the FOTON receiver's placement: the antenna's boresight is directed along the negative velocity vector of the ISS, in close proximity to solar arrays and radiators that cause time-varying blockage and multipath. Thus, each GPS SV is only visible for a relatively brief window before it appears to set behind the Earth from the perspective of the receiver. Unlike multipath, interference is likely to be highly correlated across SVs. That is, authentic signals subject to multipath or blockage may see decreases in the carrier-to-noise ratio that vary only due to the changing geometry of the ISS structure. Interference, on the other hand, is likely to be present across multiple SVs.

The FOTON Receiver

The FOTON receiver was designed by researchers at The University of Texas at Austin and Cornell University as a low-cost, space-capable, software-defined GNSS receiver [26]. Prior to its installation on the ISS, other FOTON receivers were tested in sounding rockets [26] and integrated with the ARMADILLO CubeSat (launched in 2019) [27]. FOTON is a part of the GROUP-C (GPS Radio Occultation and Ultraviolet Photometry-Colocated) experiment, which, in tandem with a compact ultraviolet photometer, is intended to provide ionospheric electron density profiles, scintillation measurements, and lower atmosphere profiles. The GROUP-C experiment was launched to the ISS in 2017 as part of the STP-H5 payload, under the aegis of the United States Department of Defense (DoD) Space Test Program (STP) [28]. These measurements are made through GPS radio occultation: from the point of view of the aft-mounted GROUP-C antennas, GPS SVs appear in the field of view and then "set" behind the Earth. The antenna's ground plane and the presence of multipath- or blockage-inducing equipment (for example, the large solar arrays on the ISS) reduces the effective field of view from a hemisphere to a narrow cone (Fig. 1b).

L1 and L2 GPS signals are first downconverted to an intermediate frequency (IF) near 298.75 MHz; then, a dual-channel analog-to-digital converter synchronously samples the signals at 5.714286 MHz with two-bit (sign and magnitude)

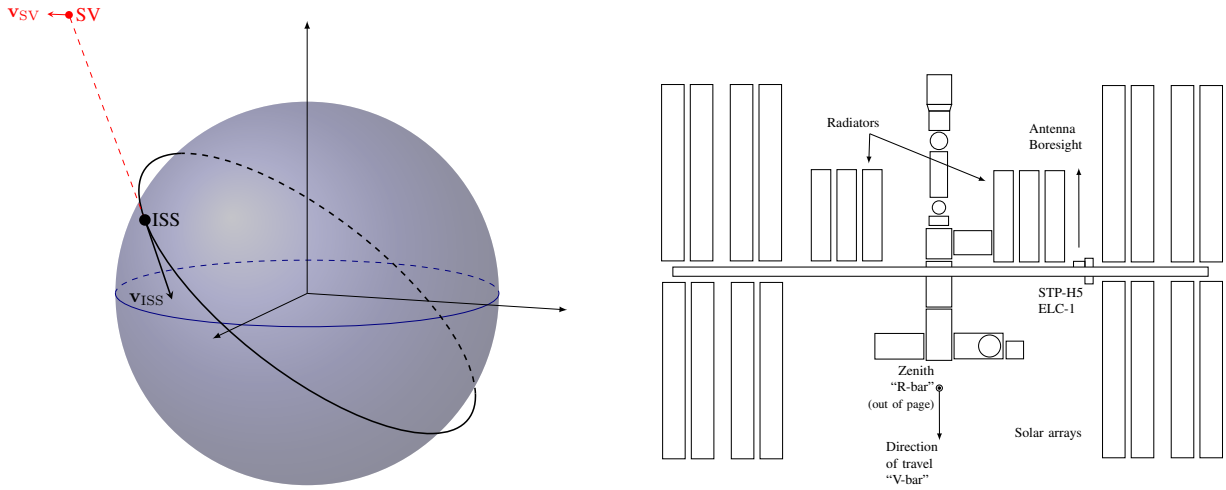

(a) The orbit of the International Space Station.

(b) The International Space Station.

Fig. 1: Overview (not to scale)

quantization [26]. An automatic gain control (AGC), the state of which is not reported by FOTON, scales the amplitude of the incoming signal to minimize quantization losses. After the observed IF signals are correlated with local replicas, the accumulator coherently sums the correlator output over a 10 ms accumulation interval to produce in-phase and quadrature (IQ) accumulations, which are the subject of the work in this paper. As is implied by the accumulation interval, the IQ data is reported at 100 Hz for each tracked signal. The FOTON receiver front-end is also capable of low level raw captures of the 5.714286 Msps IF samples for an interval of up to 70 seconds, limited chiefly by data storage and downlink capacity [10], [28].

Measurement Model

Effect of Interference on Received Power: The received ensemble of noise-free authentic GNSS signals can be modeled as

$$r_{\mathrm{A}}(t) = \sum_{i \in \mathcal{I}(t)} \sqrt{P_i} s_i(t)$$

where $s_i(t)$ is a unity-power representation of the ith signal, P_i is its received power, and $\mathcal{I}(t) = \{1, \ldots, N(t)\}$, with $N(t)$ being the number of authentic in-band GNSS signals present at time t [11], [29].

Receiver noise $r_{\mathrm{N}}(t)$ is modeled as a complex zero-mean Gaussian white noise process with (two-sided) spectral density N_0 [11].

In the absence of interference, and for $N(t) = 1$, the carrier-to-noise ratio C/N_0 of the sole authentic signal is $C/N_0 = P_1/N_0$. When $N(t) > 1$, because of multiple-access interference, the C/N_0 — now Carrier-to-Interference-and-Noise Ratio (CINR) — is for the ith signal

$$\mathrm{CINR}_i = \frac{P_i}{N_0 + M_{0i}} \tag{1}$$

where M_{0i} is the multi-access interference density caused by the $N(t) - 1$ authentic signals other than the ith one. If all these signals were BPSK(1) (like GPS L1 C/A) and identical in chipping interval T_C, then

$$M_{0i} = \frac{2}{3} T_C \sum_{j \in \mathcal{I}(t) \setminus i} P_{\mathrm{A}j} \tag{2}$$

where $\mathcal{I}(t) \setminus i$ denotes $\mathcal{I}(t)$ with i removed. In the usual case, not considered here, that the multi-access signals are a heterogeneous mix of signals with different waveforms from different GNSS constellations, M_{0i} will be rather more complicated.

When an interference signal is present, it is denoted $r_{\mathrm{I}}(t)$ and has power P_{I}. Its effect on the total receiver noise density depends on its spectral shape and on the shape of the local code replica. Denoting by $S_{r_I}(f)$ the power spectrum of $r_{\mathrm{I}}(t)$ and by $S_{C_l}(f)$ the power spectrum of the local code replica, then, when interference is present, CINR_i becomes

$$\mathrm{CINR}_i = \frac{P_i}{N_0 + M_{0i} + I_0} \tag{3}$$

where I_0 is the interference signal's equivalent noise density, is

$$I_0 \coloneqq \int_{-W_{\mathrm{FE}}/2}^{W_{\mathrm{FE}}/2} S_{r_I}(f) S_{C_l}(f)\ df \tag{4}$$

with W_{FE} being the RF front-end's noise equivalent bandwidth. For example, for in-band tone interference and a BPSK(1) desired signal, $I_0 = P_{\mathrm{I}} T_C$, where T_C is the chipping rate of the authentic signal's spreading code. For matched-spectrum interference and a BPSK(1) desired signal, $I_0 = \frac{2}{3} P_{\mathrm{I}} T_C$ [11].

The RF front end's AGC ensures that the total received power is held constant. Assuming that r_{A}, r_{N}, and r_{I} are uncorrelated with each other, and that the individual authentic signals in r_{A} are mutually uncorrelated, the total received power is

$$P_{\mathrm{T}} = P_{\mathrm{N}} + P_{\mathrm{I}} + \sum_{i \in \mathcal{I}(t)} P_i \tag{5}$$

Note that $P_{\mathrm{N}} = N_0 W_{\mathrm{FE}}$. Let $\bar{P}_{\mathrm{T}}$ be the desired setpoint for post-AGC power. Then the AGC gain factor β ensures $\beta^2 P_{\mathrm{T}} = \bar{P}_{\mathrm{T}}$ [29, Fig. 1]. Let the desired power setpoint be

$$\bar{P}_{\mathrm{T}} = P_{\mathrm{N}} + \sum_{i \in \mathcal{I}(t)} P_i \tag{6}$$

Under H_0, $P_{\mathrm{T}} = \bar{P}_{\mathrm{T}}$ and so $\beta = 1$. If an interference signal is present with power P_{I}, then P_{T} will increase and β will decrease to ensure $\beta^2 P_{\mathrm{T}} = \bar{P}_{\mathrm{T}}$. Thus,

$$\beta = \left(\frac{\bar{P}_{\mathrm{T}}}{P_{\mathrm{T}}} \right)^{1/2} = \left(\frac{P_{\mathrm{N}} + \sum_{i \in \mathcal{I}(t)} P_i}{P_{\mathrm{N}} + \sum_{i \in \mathcal{I}(t)} P_i + P_{\mathrm{I}}} \right)^{1/2} \tag{7}$$

When $P_{\mathrm{I}} \gg P_{\mathrm{N}} + \sum_{i \in \mathcal{I}(t)} P_i$, this expression can be simplified; however, this work is concerned with weak-interference cases, and so will apply this expression without simplification.

Complex Accumulation Products: The authentic accumulation product $\xi_{\mathrm{A}ik}$ for the ith signal, assuming it's being accurately tracked by the code tracking loop, can be written as

$$\xi_{\mathrm{A}ik} = \sqrt{P_{ik}} \exp(j \Delta\theta) \tag{8}$$

where P_{ik} is the power of the ith authentic signal averaged over the kth accumulation interval. The full complex accumulation for the ith signal can be written

$$S_{ik} = I_{ik} + jQ_{ik} = \beta_k \left[\xi_{\mathrm{A}ik} + \tilde{n}_{Iik} + j\tilde{n}_{Qik} \right] \tag{9}$$

where β_k is the average of $\beta(t)$ over the kth accumulation interval, and where $\tilde{n}_{Iik}$ and $\tilde{n}_{Qik}$ are the in-phase and quadrature noise terms, which are modeled as mutually-independent, white, zero-mean, Gaussian random variables with variance

$$\mathbb{E}\left[\tilde{n}_{Iik}^2\right] = \mathbb{E}\left[\tilde{n}_{Qik}^2\right] = \frac{N_0 + M_{0i} + I_0}{2T} \tag{10}$$

where T is the accumulation interval.

Suppose the carrier tracking loop is accurately tracking the ith signal. Then $\Delta\theta \approx 0$ and S_{ik} can be written

$$S_{ik} = I_{ik} + jQ_{ik} = \beta_k \left[\sqrt{P_{ik}} + \tilde{n}_{Iik} + j\tilde{n}_{Qik} \right]$$

In this case, the measured in-phase component of the signal can be modeled as

$$I_{ik} = \beta_k \left[\sqrt{P_{ik}} + \tilde{n}_{Iik} \right] \tag{11}$$

Note that the numerator in equation (10) is approximately proportional to P_{T}. As a consequence, the AGC-adjusted noise components $n_{Iik} = \beta\tilde{n}_{Iik}$ and $n_{Qik} = \beta\tilde{n}_{Qik}$ maintain nearly constant variance, denoted σ_{IQ}^2, despite changes in the received power. This fact allows the model in (11) to be simplified as follows:

$$z_{ik} \coloneqq I_{ik}/\sigma_{IQ} = \beta_k\sqrt{P_{ik}}/\sigma_{IQ} + n_{Iik}/\sigma_{IQ} \tag{12a}$$

$$= \beta_k\rho_{ik} + n_{ik} \tag{12b}$$

where

$$\mathbb{E}\left[n_{ik}n_{ij}\right] = \delta_{kj}, \quad \mathbb{E}\left[n_{ik}n_{jk}\right] = c_{ij,k} \tag{13}$$

with δ_{kj} being the Kroenecker delta, which is unity for $k = j$ and otherwise zero. The model in (12b) will serve as the basis of this paper's detection tests.

Note that the ith signal's accumulation interval extends from $t_{i,k-1}$ to $t_{i,k} = t_{i,k-1} + T_{i,k}$, where $T_{i,k}$ is approximately equal to the nominal accumulation interval T but varies slightly in response to carrier Doppler. The accumulation intervals for all signals are approximately aligned so that $\|t_{ik} - t_{jk}\| < T$ for $i \in \mathcal{I}(t_k)$ and $k \in \{1, \ldots, K\}$. Let N_k be the number of signals tracked continuously throughout the kth interval, and let $\mathcal{I}_k = \{1, \ldots, N_k\}$. The correlation coefficients $c_{ij,k},\ i, j \in \mathcal{I}_k$ are typically nonzero due to overlap among the tracked signals' kth accumulation intervals.

In summary, using the in-phase component I of the complex correlation data is desirable for three reasons: it (i) is easier to model, (ii) is better for accurately determining the onset of an interference event because it has not been low-pass filtered by the receiver like FOTON's reported C/N_0 observable, and (iii) has an advantage in detection sensitivity for low C/N_0 (less than 35 dB-Hz).

Problem Statement

The transient interference detection problem can now be stated as follows. Let $\boldsymbol{z}_k = [z_{1k}, \ldots, z_{N_k k}]^{\mathsf{T}}$ be the vector of normalized in-phase accumulation products over each signal's kth interval. Consider a set of such measurements $\mathcal{Z} = \{\boldsymbol{z}_k\}_1^K$. Denote by $\mathcal{Z}_{ab} = \{\boldsymbol{z}_k\}_{k=a}^b$ the subset of measurements from $k = a$ to $k = b$, and denote by $\mathcal{Z}_{ab}^c = \mathcal{Z}\backslash\mathcal{Z}_{ab}$ its complement. Assume that at most one interference event occurs within $\mathcal{Z}$, and let a and b denote the first and last indices of its span under H_1. Also let $\mathcal{K} = \{1, \ldots, K\}$, $\mathcal{K}_{ab} = \{k \mid \boldsymbol{z}_k \in \mathcal{Z}_{ab}\}$, and $\mathcal{K}_{ab}^c = \{k \mid \boldsymbol{z}_k \in \mathcal{Z}_{ab}^c\}$.

The null and alternate hypotheses for the transient interference detection problem may now be written as

$$H_0 : \boldsymbol{z}_k \sim \mathcal{N}(\boldsymbol{\rho}_k, P_k), \quad k \in \mathcal{K}_{ab}^c \tag{14a}$$

$$H_1 : \boldsymbol{z}_k \sim \mathcal{N}(\beta_k\boldsymbol{\rho}_k, P_k), k \in \mathcal{K}_{ab}. \tag{14b}$$

where $0 < \beta_k < 1$ is the AGC attenuation factor, $\boldsymbol{\rho}_k = [\rho_{1k}, \ldots, \rho_{N_k k}]^{\mathsf{T}}$ and where $P_k \in \mathbb{R}^{N_k \times N_k}$ is the covariance matrix for $\boldsymbol{z}_k$ under either hypothesis, which is assumed to be known.

Transient Change Detection

A windowed sum of log-likelihood ratios under the assumption of a multivariate Gaussian mean change was used to detect interference events. Following the work of Egea-Roca et al. [30, §III, §IV.B], the log-likelihood ratio at sample epoch k, $\boldsymbol{y}_k$ can be written as

$$y_k = (\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k})^T P_k^{-1} \left(\boldsymbol{z}_k - \frac{1}{2}(\tilde{\boldsymbol{\mu}}_{1,k} + \boldsymbol{\mu}_{0,k})\right) \tag{15}$$

where $\boldsymbol{z}_k$ is the measurement vector consisting of the in-phase component of the complex accumulation at sample epoch k, $\boldsymbol{\mu}_{0,k} = \mathbb{E}[\boldsymbol{z}_k|H_0]$, $\tilde{\boldsymbol{\mu}}_{1,k} = \mathbb{E}[\boldsymbol{z}_k|H_1]$, and P_k is the covariance matrix of the measurement vector, under the assumption that the distributions of both hypotheses have the same covariance. Note that y_k is a scalar (univariate) normal random variable.

Note that the set of satellites visible changes over time. The measurement $\boldsymbol{z}_k = [z_{1k}, \ldots, z_{N_k k}]^{\mathsf{T}}$ is of cardinality N_k. When the cardinality of the measurement is variable, and may change from sample epoch to sample epoch, it is

necessary to normalize the log likelihood by the cardinality (or 'dimension') N_k of the measurement at the kth sample epoch. Taking this step yields

$$y_k = \frac{1}{N_k} \left(\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k}\right)^T P_k^{-1} \left(\boldsymbol{z}_k - \frac{1}{2}(\tilde{\boldsymbol{\mu}}_{1,k} + \boldsymbol{\mu}_{0,k})\right) \tag{16}$$

The mean and variance of this detection statistic can be expressed as

$$\mu_{y_k|0} = -\frac{1}{2N_k} \left(\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k}\right)^T P_k^{-1} \left(\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k}\right)) \tag{17}$$

$$\sigma_{y_k}^2 = \frac{1}{N_k^2} \left(\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k}\right)^T P_k^{-1} \left(\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k}\right)) \tag{18}$$

$$\mu_{y_k|1} = \frac{1}{N_k} \left(\tilde{\boldsymbol{\mu}}_{1,k} - \boldsymbol{\mu}_{0,k}\right)^T P_k^{-1} \left(\boldsymbol{\mu}_{1,k} - \frac{1}{2}(\tilde{\boldsymbol{\mu}}_{1,k} + \boldsymbol{\mu}_{0,k})\right) \tag{19}$$

where $\mu_{y_k|0} = \mathbb{E}[y_k|H_0]$, $\mu_{y_k|1} = \mathbb{E}[y_k|H_1]$, and $\sigma_{y_k}^2 = \mathrm{Var}(y_k)$ under both hypotheses.

The finite moving average metric — actually, a windowed sum — used to detect interference events is written as

$$\mathrm{FMA}_k(m) = \sum_{i=k-m+1}^{k} y_i \tag{20}$$

for window length m, and is compared to the thresholds $h_k^{\mathrm{start}}(\tilde{\alpha}, m)$ and $h_k^{\mathrm{end}}(\tilde{\alpha}, m)$ to mark the beginning and end of an event, respectively.

$$h_k^{\mathrm{start}}(\tilde{\alpha}, m) = \sqrt{\sum_{i=k-m+1}^{k} \sigma_{y_i}^2} \cdot \Phi^{-1}\left[(1-\tilde{\alpha})^{1/m_\alpha}\right] + \sum_{i=k-m+1}^{k} \mu_{y_i|0} \tag{21}$$

$$h_k^{\mathrm{end}}(\tilde{\alpha}, m) = -\sqrt{\sum_{i=k-m+1}^{k} \sigma_{y_i}^2} \cdot \Phi^{-1}\left[(1-\tilde{\alpha})^{1/m_\alpha}\right] + \sum_{i=k-m+1}^{k} \mu_{y_i|1} \tag{22}$$

In these thresholds, $\tilde{\alpha}$ is the desired false alarm probability, Φ^{-1} is the inverse cumulative distribution function of a standard normal random variable, and m_α is the duration for which the false alarm probability is desired [30, Appendix A]. In practice, $m_\alpha \geq m$; in this work, $m_\alpha = m$ for clarity.

A candidate interference event is considered to have begun at sample epoch i if $y_i > h_i^{\mathrm{start}}(\tilde{\alpha}, m)$. Thenceforth, up to sample epoch $i + M$, where M is the maximum considered interference event duration, the event is considered to have ended if (i) $y_j > h_j^{\mathrm{end}}(\tilde{\alpha}, m)$ for $j \in i..i+M-1$ and (ii) $y_l < h_l^{\mathrm{end}}(\tilde{\alpha}, m)$ for $l > j, l \in i..i+M$. A candidate event is discarded if y_k never exceeds $h_k^{\mathrm{end}}(\tilde{\alpha}, m)$ at any point $k \in \{i..i+M\}$.

The distribution of the measurement $\boldsymbol{z}_k$ under the null hypothesis H_0 is computed from the statistics of collected data. The mean vector $\boldsymbol{\mu}_{0,k} = \boldsymbol{\mu}_0 \ \forall k \in \mathcal{K}_{(\mathrm{Year,DoY})}$, where $\mathcal{K}_{(\mathrm{Year,DoY})}$ is the set of sample epochs that occur during the specified year and day-of-year (in the UTC time standard). To reduce the likelihood that the "control" data that is used to compute the null hypothesis is itself subject to interference events, only data captured over the open ocean is incorporated.

The actual distribution of the measurement $\boldsymbol{z}_k$ under the alternate hypothesis H_1 is *a priori* unknown, and in fact is variable due to its dependence on the strength of the interference source. For this reason, $\tilde{\boldsymbol{\mu}}_{1,k}$ can be thought of as the mean of the alternate hypothesis for the minimum change one would like to detect. In contrast, $\boldsymbol{\mu}_{1,k}$ is the assumed actual expected value of $\boldsymbol{z}_k$ under the alternate hypothesis. Absent careful characterization of known GNSS interference events, it suffices to make the approximations $\beta_k = \beta \ \forall k \in \mathcal{K}_{(\mathrm{Year,DoY})}$ and $\tilde{\boldsymbol{\mu}}_{1,k} = \boldsymbol{\mu}_1 = \beta\boldsymbol{\mu}_0 \ \forall k \in \mathcal{K}_{(\mathrm{Year,DoY})}$.

There are three user-defined parameters of note (Table I). The first, the false alarm probability $\tilde{\alpha}$, is familiar in hypothesis testing problems [31]; in this context, however, it is perhaps better described as an upper bound to the probability of false alarm [30, Theorem 1, Appendix A]. The choice of window length m is of critical importance. If m is too small, then the detector is over-sensitive to short-duration anomalies like scintillation. On the other hand, if m is greater than the length of an interference event, the event may not be detected. A final detector parameter to set is the inferred AGC gain factor $\beta \in [0, 1]$. A value of β closer to zero (in practice, less than 0.3) will trigger the detection threshold for only the most extreme attenuation of the measured I data. A higher value of β (in practice, greater than 0.9), yields an over-sensitive detector whose flagged intervals require manual inspection.

TABLE I: User-defined parameters.

Parameter	Value
$\tilde{\alpha}$	0.01
m	5000
β	0.75

Figure 2 demonstrates the transient change detection approach when applied to simulated data. Note that the data is composed of several different signals, each of which is only present for a portion of the entire interval. Thus, at any given sample epoch, a measurement consists of all the signals available at that epoch and will have a cardinality equal to the total number of signals available at that epoch. The change event is clearly visible in the finite moving average detection statistic. The detection delay — note the distance between the true start of the event and the estimated start of the event — is largely a result of the choice of false alarm probability. Observe the gap in the detection statistic where there is no data available, or not enough data within the window to compute a valid windowed sum. This is a common occurrence throughout the real data collected by FOTON: some intervals have no admissible data because all visible satellites were likely to be obstructed (see Section).

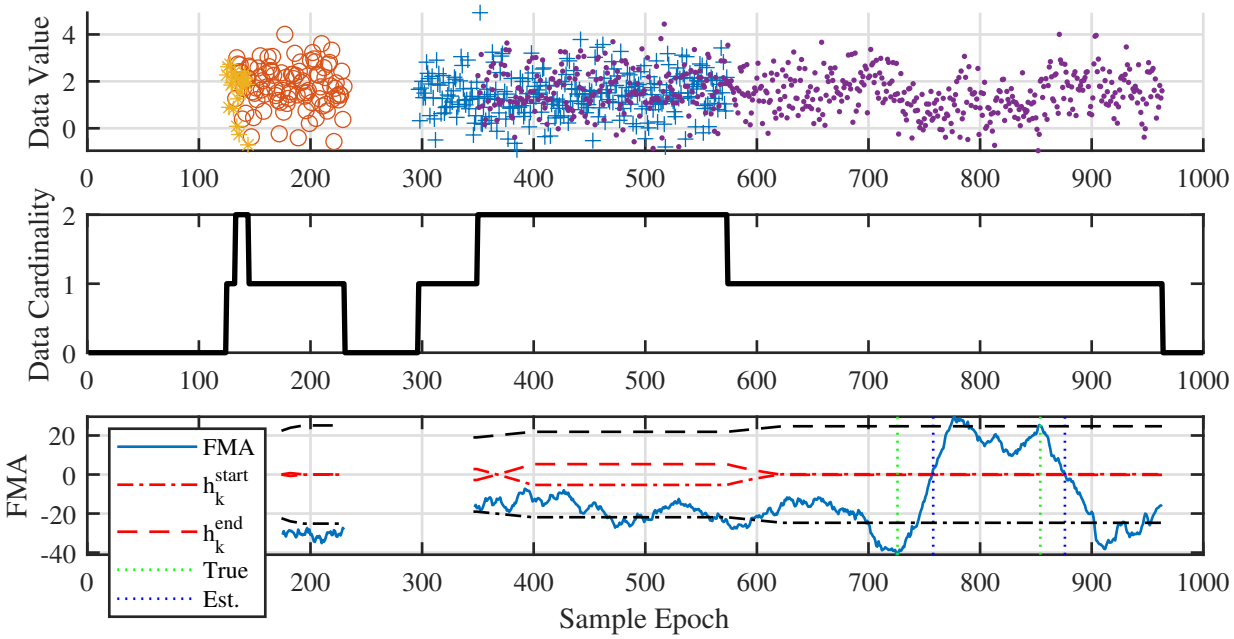

Fig. 2: Application of transient change detection approach to simulated data.

RESULTS

Identifying Possible GNSS Interference

Murrian et al. [10, §III] details a serendipitous recorded by the FOTON receiver of GNSS interference in the vicinity of the Eastern Mediterranean. This data capture was initiated for occultation purposes, but the raw frontend samples collected permitted Doppler-based geolocation of the interference source, which was determined to be Khmeimim Air Base in Syria. Further analysis of 1 Hz receiver observables recorded by FOTON over multiple years discovered that this interference source was remarkably powerful and consistent throughout the entire dataset [10, §IV]. One flagged interval identified (Fig. 3) is likely to result from this interference source. The interference identified in Figs. 3-5 is also corroborated by third-party reporting: the regions overflown by the ISS in the ground tracks pictured are identified as hot spots in a report that compiled data collected by Airbus aircraft [3].

A feature highly suggestive of terrestrial interference is the rapid recovery of the in-phase component's magnitude at the end of the interval. The point of this rapid increase could correspond to the time at which an interference source on the ground disappears over the horizon after overflight by the ISS. This feature appears in many of the flagged

intervals, and all of the examples included in this work (Figs. 3, 4, and 5). The more gradual apparent onset of the interference could be explained in part by the fact that the antenna faces rearward.

The primary challenge in identifying candidate interference events is the paucity of useful data. Due to the narrow field of view of FOTON's antenna, only a handful of GPS satellites may be visible at any given time. Figs. 3 and 4 demonstrate that there are intervals for which no useful data can be applied to the task of interference detection; nevertheless, it is possible to detect the event as long as there is sufficient data (at least one signal) over the interval during which the interference is thought to occur.

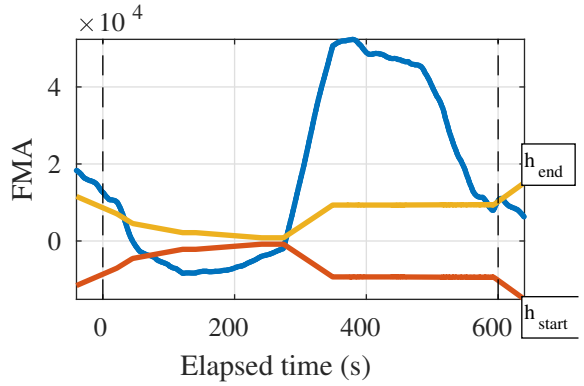

(a) Finite-moving-average detection metric and thresholds.

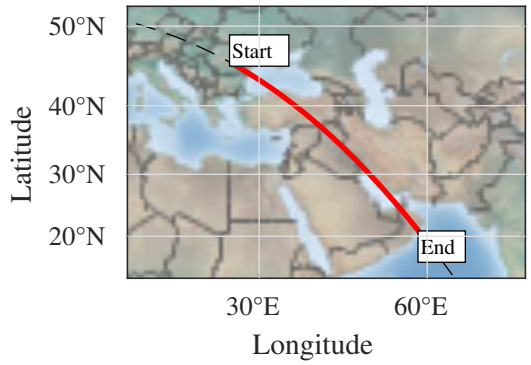

(b) Ground track of the ISS.

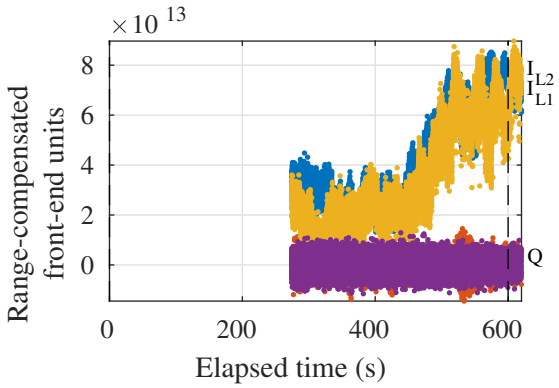

(c) Range-adjusted IQ data originating from TXID 6.

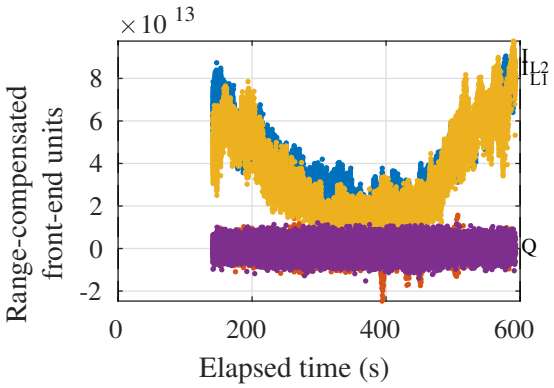

(d) Range-adjusted IQ data originating from TXID 19.

Fig. 3: GNSS interference thought to originate from Western Syria on September 16, 2019. Vertical dotted lines indicate the extent of the flagged event.

Ruling out alternatives

In order to ensure that the events detected by this paper's approach represent likely terrestrial interference events, as opposed to signal fading caused by blockage or multipath, it is necessary to restrict the search to a subset of the IQ data. At each sample epoch, only signals from GPS satellites that met the following criteria were admitted to the detection statistic (Fig. 6):

1) The vector to the GPS satellite from FOTON lies within a cone of at most a specified angle ϕ as measured from the boresight (assumed to be the negative velocity vector), in order to reduce the impact of blockage and multipath at the ISS.
2) The line connecting the GPS satellite with the ISS has a minimum straight-line tangent altitude (SLTA) above a specified threshold, in order to avoid including signals that are soon to set behind the limb of the Earth.

There is an inherent trade-off to filtering the data by these methods: more stringent selection will ensure that the signals are maximally useful for detecting interference. On the other hand, fewer usable signals increases the proportion of sample epochs for which there is no usable data at all, making it more likely for interference events to go undetected.

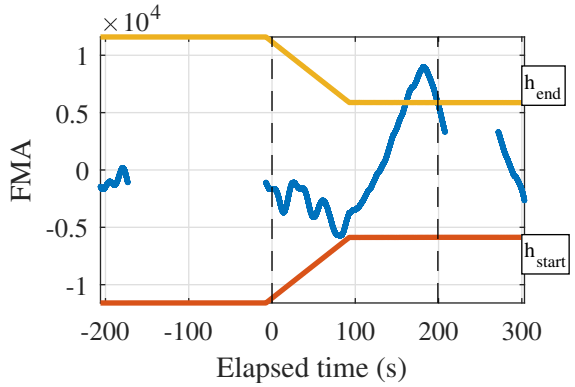

(a) Finite-moving-average detection metric and thresholds.

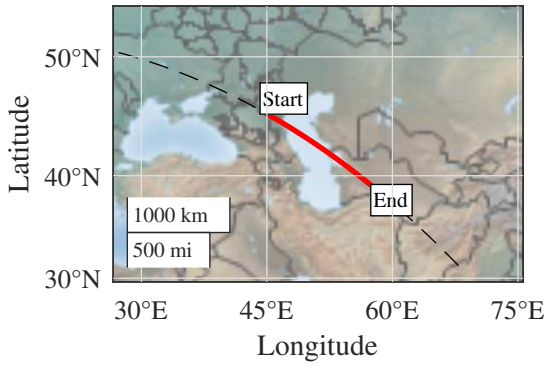

(b) Ground track of the ISS.

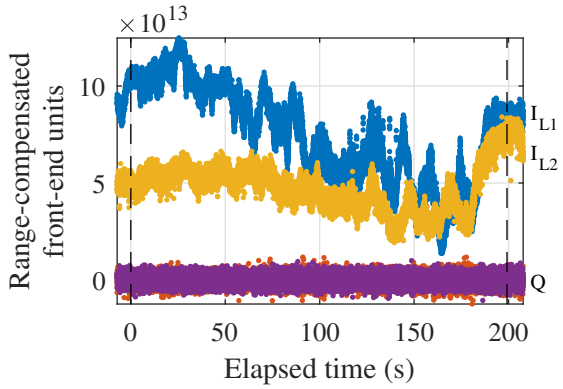

(c) Range-adjusted IQ data originating from TXID 17.

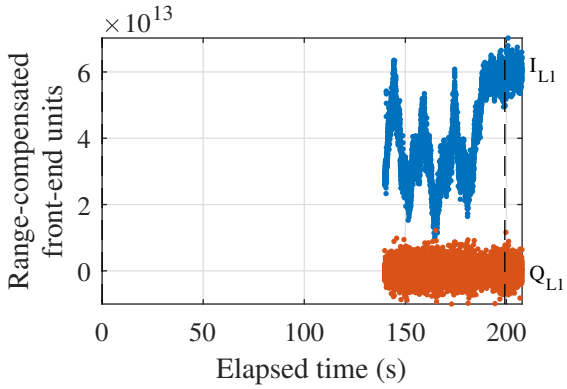

(d) Range-adjusted IQ data originating from TXID 19.

Fig. 4: GNSS interference thought to originate from the Black Sea region on March 14, 2018. Vertical dotted lines indicate the extent of the flagged event.

Future Work

One of the main challenges to detecting GNSS interference from the ISS was the signal obstruction and multipath environment close to the antenna caused by the changing geometry of the radiators and solar panels. This could be ameliorated by taking the receiver off the ISS: one could fly a LEO spacecraft for the explicit purpose of detecting, recording, and localizing GNSS interference. Such a spacecraft could use a version of the technique presented in this paper to autonomously initiate a raw data capture for downlink and postprocessing.

Since the publication of this work and the work that preceded it, jamming and spoofing strategies may have changed in response to scrutiny. Accordingly, interference detection techniques may need to be updated in order to detect new forms of interference in navigation bands.

CONCLUSIONS

Terrestrial GNSS interference can be powerful enough to affect receivers operating in low Earth orbit. The ability to detect the onset and end of an interference event elevates it from being a nuisance into being a valuable resource. Low Earth orbit is thenceforth transformed into a vantage point for identifying and geolocating the source of this interference. An opportunity such as this does not come without its challenges, however: depending on the detection statistic used, any decision rule must take into account occultation by the Earth and local multipath, among other factors. The approach demonstrated in this paper adapts transient change detection theory, computing the likelihood of obtaining a finite window of measurements conditioned on the presence or absence of possible interference. This method is applied to search for episodic decreases in the strength of authentic received signals. Instead of having to settle for the carrier-to-noise ratio values reported by the receiver, it is possible to take advantage of the lower-level in-phase accumulation as a higher-rate substitute for carrier-to-noise ratio when detecting apparent decreases in received

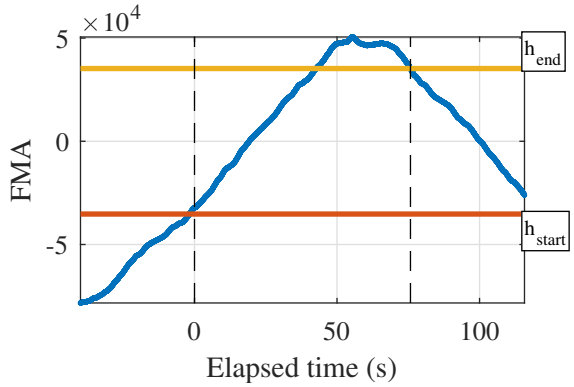

(a) Finite-moving-average detection metric and thresholds.

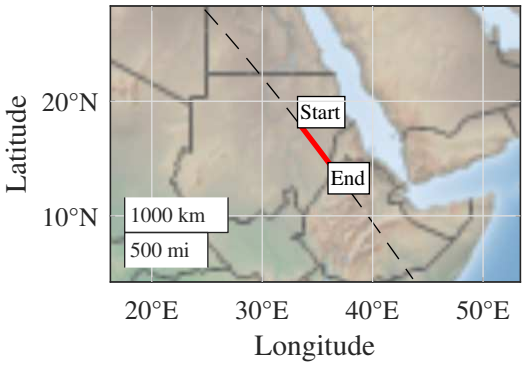

(b) Ground track of the ISS.

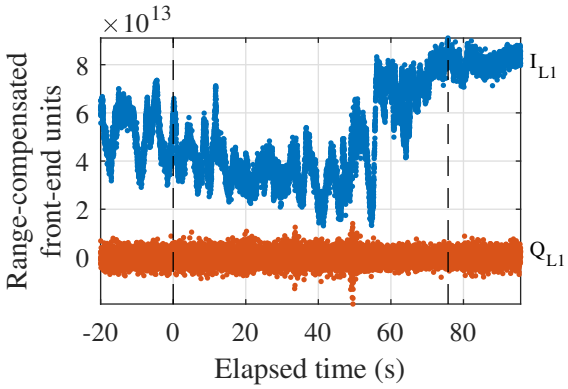

(c) Range-adjusted IQ data originating from TXID 2

Fig. 5: GNSS interference thought to originate from Northern Africa on March 14, 2018. Vertical dotted lines indicate the extent of the flagged event.

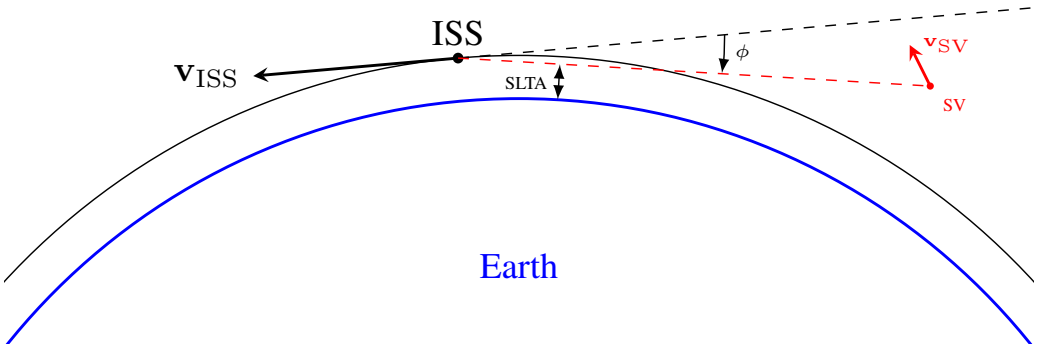

Fig. 6: Straight-line tangent altitude (SLTA) and receiver boresight zenith angle ϕ (not to scale)

power. A closer look at flagged intervals — alongside corroborative third-party reporting on the ground — suggests that terrestrial GNSS interference is the likely culprit.

ACKNOWLEDGMENTS

The STP-H5/GROUP-C experiment was integrated and flown under the direction of the Department of Defense Space Test Program.

Work at UT Austin was supported by the U.S. Department of Transportation (USDOT) under Grant 69A3552047138 for the CARMEN University Transportation Center (UTC). It was also supported by affiliates of The University of Texas Wireless Networking and Communications Group.

REFERENCES

[1] Liu, Z., Lo, S., and Walter, T., "GNSS Interference Characterization and Localization Using OpenSky ADS-B Data," *Multidisciplinary Digital Publishing Institute Proceedings*, Vol. 59, No. 1, 2020.

[2] Goward, D., "Russia ramps up GPS jamming along with troops at Ukraine border," *GPS World*, Apr 2021.

[3] *Does radio frequency interference to satellite navigation pose an increasing threat to network efficiency, cost-effectiveness and ultimately safety?*, Eurocontrol Aviation Intelligence Unit, Mar 2021.

[4] Garcia, M. A., Dolan, J., and Hoag, A., "Aireon's initial on-orbit performance analysis of space-based ADS-B," *2017 Integrated Communications, Navigation and Surveillance Conference (ICNS)*, 2017, pp. 4A1–1–4A1–8.

[5] Stader, J. and Gunawardena, S., "Leveraging worldwide, publicly-available data to create an automated satnav interference detection system," *Proceedings of the 2021 International Technical Meeting of The Institute of Navigation*, 2021, pp. 69–83.

[6] Sarda, K., CaJacob, D., Orr, N., and Zee, R., "Making the invisible visible: Precision RF-emitter geolocation from space by the HawkEye 360 Pathfinder mission," *Proceedings of the 32nd Annual AIAA/USU Conference on Small Satellites,* Next on the Pad, No. SSC18-II-06, 2018.

[7] Werner, D., "First Hawkeye 360 satellites pinpointing signals," *SpaceNews*, Feb 2019.

[8] Berghel, H., "Wireless infidelity I," *Commun. ACM*, Vol. 47, 09 2004, pp. 21–26.

[9] Psiaki, M. L. and Humphreys, T. E., "GNSS Spoofing and Detection," *Proceedings of the IEEE*, Vol. 104, No. 6, 2016, pp. 1258–1270.

[10] Murrian, M. J., Narula, L., Iannucci, P. A., Budzien, S., O'Hanlon, B. W., Powell, S. P., and Humphreys, T. E., "First Results from Three Years of GNSS Interference Monitoring from Low Earth Orbit," *Navigation, Journal of the Institute of Navigation*, 2021, To be published.

[11] Humphreys, T. E., "Interference," *Springer Handbook of Global Navigation Satellite Systems*, Springer International Publishing, 2017, pp. 469–503.

[12] Chandola, V., Banerjee, A., and Kumar, V., "Anomaly detection: a survey," *ACM Computing Surveys*, Vol. 41, No. 3, 2009, pp. 1–72.

[13] Rajabzadeh, Y., Rezaie, A. H., and Amindavar, H., "A dynamic modeling approach for anomaly detection using stochastic differential equations," *Digital Signal Processing*, Vol. 54, 2016, pp. 1–11.

[14] Xu, Z., Kersting, K., and Von Ritter, L., "Stochastic online anomaly analysis for streaming time series," *IJCAI*, 2017, pp. 3189–3195.

[15] Miralles, D., Bornot, A., Rouquette, P., Levigne, N., Akos, D. M., Chen, Y.-H., Lo, S., and Walter, T., "An assessment of GPS spoofing detection via radio power and signal quality monitoring for aviation safety operations," *IEEE Intelligent Transportation Systems Magazine*, Vol. 12, No. 3, 2020, pp. 136–146.

[16] Bastide, F., Chatre, E., and Macabiau, C., "GPS interference detection and identification using multicorrelator receivers," *Proceedings of the 14th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2001)*, 2001, pp. 872–881.

[17] O'Mahony, G., O'Mahony, S., Curran, J. T., and Murphy, C. C., "Developing a low-cost platform for GNSS interference detection," *Proceedings of the 2015 European Navigation Conference, Bordeaux, France*, 2015.

[18] Ndili, A. and Enge, P., "GPS receiver autonomous interference detection," *IEEE 1998 Position Location and Navigation Symposium (Cat. No. 98CH36153)*, IEEE, 1996, pp. 123–130.

[19] Manfredini, E. G., Akos, D. M., Chen, Y.-H., Lo, S., Walter, T., and Enge, P., "Effective GPS spoofing detection utilizing metrics from commercial receivers," *Proceedings of the 2018 International Technical Meeting of The Institute of Navigation*, 2018, pp. 672–689.

[20] Hewitson, S. and Wang, J., "GNSS receiver autonomous integrity monitoring (RAIM) performance analysis," *GPS Solutions*, Vol. 10, No. 3, 2006, pp. 155–170.

[21] Hegarty, C., Odeh, A., Shallberg, K., Wesson, K., Walter, T., and Alexander, K., "Spoofing detection for airborne GNSS equipment," *Proceedings of the 31st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2018)*, 2018, pp. 1350–1368.

[22] Poor, H. V. and Hadjiliadis, O., *Quickest detection*, Cambridge University Press, 2008.

[23] Truong, C., Oudre, L., and Vayatis, N., "Selective review of offline change point detection methods," *Signal Processing*, Vol. 167, 2020, pp. 107299.

[24] Page, E., "A test for a change in a parameter occurring at an unknown point," *Biometrika*, Vol. 42, No. 3/4, 1955, pp. 523–527.

[25] Sen, A. and Srivastava, M. S., "On tests for detecting change in mean," *The Annals of statistics*, 1975, pp. 98–108.

[26] Lightsey, E. G., Humphreys, T. E., Bhatti, J. A., Joplin, A. J., O'Hanlon, B. W., and Powell, S. P., "Demonstration of a Space Capable Miniature Dual Frequency GNSS Receiver," *Navigation*, Vol. 61, No. 1, Mar. 2014, pp. 53–64.

[27] Joplin, A. J., Lightsey, E. G., and Humphreys, T. E., "Development and Testing of a Minaturized, Dual-Frequency GPS Receiver for Space Applications," *Proceedings of the ION International Technical Meeting*, Long Beach, CA, Jan. 2012.

[28] Budzien, S. A., Humphreys, T. E., Powell, S. P., O'Hanlon, B. W., Bishop, R. L., and Stephan, A. W., "GPS Radio Occultation and Ultraviolet Photometry-Colocated (GROUP-C) early orbit testing results," official memorandum, Naval Research Laboratory, Washington, DC, USA, 2020 [Online].

[29] Wesson, K. D., Gross, J. N., Humphreys, T. E., and Evans, B. L., "GNSS Signal Authentication Via Power and Distortion Monitoring," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 54, No. 2, April 2018, pp. 739–754.

[30] Egea-Roca, D., López-Salcedo, J. A., Seco-Granados, G., and Poor, H. V., "Performance bounds for finite moving average tests in transient change detection," *IEEE Transactions on Signal Processing*, Vol. 66, No. 6, 2018, pp. 1594–1606.

[31] Van Trees, H. L., *Detection, Estimation, and Modulation Theory*, Wiley, 2001.