Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
Pending assignment.	N/A	N/A

4. Title and Subtitle	5. Report Date
GNSS Interference: Situational Awareness and LEO Backup	September 30, 2024
	6. Performing Organization Code N/A

7. Author(s)	8. Performing Organization Report No.
Todd E. Humphreys, P.h.D.; Zachary L. Clements; and Wenkai Qin	N/A

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
The University of Texas at Austin 1616 Guadalupe St. UTA Bldg. Suite 3.302 Austin, TX 78701	N/A
	11. Contract or Grant No. 69A3552348327

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
The Ohio State University Address: 281 W Lane Ave, Columbus, OH 43210	Final Report (June 2023 to August 2024)
	14. Sponsoring Agency Code N/A

15. Supplementary Notes

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

16. Abstract

Our research has addressed the major goals of the CARMEN+ UTC, mainly in (1) identifying and analyzing existing and emerging cybersecurity threats to highly HATS, and (2) developing and experimentally verifying cyber-resilient mitigation methods. GNSS is fragile: its service is easily denied by jammers or deceived by spoofers. The civilian aviation and maritime industries are seeing significant electronic warfare spillover from nearby conflict zones. A technique was developed and experimentally verified to geolocate GNSS spoofers with a single-satellite from Low Earth Orbit (LEO). Additionally, we have been studying the possibility of using signals from mega constellations of LEO satellites designed for broadband communications as a backup to traditional GNSS for PNT. Our findings are presented in this report. Commercial partnerships with existing LEO constellations can be fruitful, as they can provide near real-time GNSS interference monitoring and can serve as backup PNT services when GNSS is denied.

17. Key Words	18. Distribution Statement
GNSS interference, LEO PNT, interference geolocation	No restrictions. This document is available to the public through NTIS: National Technical Information Service Springfield, Virginia 22161

19. Security Classif.(of this report)	20. Security Classif.(of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	45 pages	

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

CARMEN+ UTC

Center for Automated Vehicle Research with Multimodal AssurEd Navigation

University Transportation Centers Program

Final Report: GNSS Interference: Situational Awareness and LEO Backup

P.I.	Project Info:
Todd E. Humphreys University of Texas at Austin Department of Aerospace Engineering and Engineering Mechanics	Grant No. 69A3552348327
	DUNS: 832127323
	EIN #: 31-6025986
	Project Effective: June 1, 2023 Project End: August 30, 2024 Submission: September 30, 2024

Consortium Members:

The Ohio State University College of Engineering

UCI University of California, Irvine

TEXAS The University of Texas at Austin

North Carolina Agricultural and Technical State University

DISCLAIMER

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

Abstract

This report offers the technical development resilient PNT under GNSS interference. Our research has addressed the major goals of the CARMEN+ UTC, mainly in (1) identifying and analyzing existing and emerging cybersecurity threats to highly HATS, and (2) developing and experimentally verifying cyber-resilient mitigation methods.

This report offers an experimental demonstration of single-satellite single-pass geolocation of a terrestrial Global Navigation Satellite System (GNSS) spoofer from Low Earth Orbit (LEO). The proliferation of LEO-based receivers can provide unprecedented spectrum awareness, enabling persistent GNSS interference detection and geolocation. Accurate LEO-based single-receiver emitter geolocation is possible when a range-rate time history can be extracted, traditionally accomplished through Doppler measurements. However, Doppler-based measurement techniques assume the emitter transmits at a quasi-constant center frequency. This assumption is not true for GNSS spoofers, as they transmit an ensemble of spoofing signals wherein each spoofed signal's carrier frequency contains an unique unknown time-varying frequency component that imitates the Doppler corresponding to the spoofed navigation satellite and spoofed location. This report presents a technique that removes the unknown time-varying frequency component across each signal so that the range-rate time history between receiver and transmitter can be extracted and exploited for geolocation. If a GNSS receiver allows itself to be spoofed, the range-rate between the receiver and the spoofer will manifest in the GNSS receiver's clock drift estimate. This technique is verified by a controlled experiment in partnership with Spire Global, in which a LEO-based receiver captures GNSS spoofing signals transmitted from a known ground station on a non-GNSS frequency band.

This report develops an antenna system to support use of signals from low Earth orbit (LEO) megaconstellations to conduct satellite beam and channel occupancy studies. The antenna system is based on an articulated horn antenna that can reposition itself from one LEO satellite vehicle (SV) to another in less than a second, which allows quasi-continuous LEO signal tracking despite the frequent satellite handoffs that occur in a LEO-based communications network. It offers an alternative to phased-array antennas for LEO megaconstellation signals, which are not yet commercially available at a reasonable cost. The system automatically directs its horn antenna towards transmitting LEO SVs such that acquisition and tracking through a matched filter can be performed using known portions of the signal structure. This system serves as a proof-of-concept and development base for a full LEO position, navigation, and timing (PNT) solution. Experimental results are presented showing successful sequential acquisition of time of arrival measurements from SpaceX's Starlink constellation as well as preliminary results for beam and channel occupancy, observing that up to three SVs may be simultaneously transmitting assigned beams toward a given user service cell.

A signal capture and analysis technique for extracting precise timing information from the Starlink communications megaconstellation's downlink transmissions was developed. Several characterizations of the Starlink frame clock adjustment pattern and stability are presented. The frame clock is adjusted at a regular 1 Hz cadence, and the adjustments are nearly discontinuous in nature. A composite clock Allan deviation analysis indicates that the Starlink frame clock has best-case stability characteristic of a temperature-controlled crystal oscillator. A further high-frequency clock instability analysis is conducted via polynomial trend removal, and indicates that the Starlink frame clock could hypothetically support a global position, navigation, and timing (PNT) mission when performing nominally, but manifests episodic

oscillatory and excursive behavior that would severely degrade opportunistic positioning and timing based on pseudoranges formed from the frame clock. Examples of such oscillatory and excursive patterns are shown and other aspects of the phenomena are discussed.

Acknowledgements

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

We would like to thank Spire Global for their collaboration in performing the spoofer geolocation experiment.

Executive Summary

Global Navigation Satellite Systems (GNSS) such as GPS provide meter-accurate positioning while offering global accessibility and all-weather, radio-silent operation. However, GNSS is fragile: its service is easily denied by jammers or deceived by spoofers. GNSS signals are especially vulnerable to jamming and spoofing because they are extremely weak: near the surface of Earth, they have no more flux density than light received from a 50 W bulb at a distance of 2000 km. GNSS jamming and spoofing is no longer strictly limited to the military battlegrounds or just as an academic topic. The civilian aviation and maritime industries are seeing significant electronic warfare spillover from nearby conflict zones. Without proper countermeasures, victim GNSS receivers can be rendered useless.

GNSS receivers in Low Earth Orbit (LEO) are a proven asset for detecting, classifying, and geolocating terrestrial GNSS interference that can be a danger to civil aviation, maritime, or ground vehicle traffic. Emitter geolocation from LEO offers worldwide coverage with a frequent refresh rate, making it possible to maintain a common operating picture of terrestrial sources of interference. We have developed and experimentally verified several techniques for LEO-based interference monitoring.

So far, our key findings for this research thrust include: (1) Since August 2023 there has been an alarming rise of GNSS spoofing across the Middle East and Eastern Europe. For the first time, GNSS spoofing has significantly affected commercial aviation, (2) it is possible to locate sources of GNSS spoofing to within 500 meters using pseudorange and Doppler measurements from LEO GNSS receivers provided that the spoofed navigation data are also extractable, and (3) it is possible to locate sources of GNSS spoofing to within 3km with only a single spoofed channel's Doppler time history.

Since the CARMEN+ UTC kickoff, we have been studying the possibility of using signals from mega constellations of LEO satellites designed for broadband communications as a backup to traditional GNSS for PNT. Before the kickoff, we had already identified structures in Starlink signals that will be useful for

measuring the time of arrival (TOA) of Starlink frames. But whether such frames could be used to form pseudorange measurements that are useful for PNT remained an open question.

The answer depends on the beam patterns and scheduling, proportion of predictable frame content, and frame timing properties of the Starlink signals. We have found numerous serious anomalies in Starlink frame timing. These complicate formation of easily-modeled pseudorange measurements. Our findings for this reporting period are as follows:

1. Up to 16 unique Starlink satellites could, in principle, simultaneously illuminate the same service cell – two on each of 8 channels – without violating the FCC bounds on effective power flux density (EPFD). This is an encouraging result as regards exploiting Starlink as a PNT source. In practice, May contain trade secrets or commercial or financial information that is privileged or confidential and exempt from public disclosure. we have verified illumination by up to three so-called assigned beams using an agile antenna system.

2. We have developed and verified a simple heuristic by which Starlink satellites likely to be directing an assigned beam to a given service cell may be predicted. Our heuristic is 94% effective at finding at least one assigned beam in each fixed assignment interval.

Our research has addressed the major goals of the CARMEN+ UTC, mainly in (1) identifying and analyzing existing and emerging cybersecurity threats to highly HATS, and (2) developing and experimentally verifying cyber-resilient mitigation methods. Commercial partnerships with existing LEO constellations can be fruitful, as they can provide near real-time GNSS interference monitoring and can serve as backup PNT services when GNSS is denied.

Table of Contents

Publications

1. Z. Clements, I. Goodridge, P. Ellis, M. J. Murrian, and T. E. Humphreys, “Demonstration of single-satellite GNSS spoofer geolocation,” in Proceedings of the ION International Technical Meeting, (Long Beach, CA), pp. 361–373, 2024

2. W. Qin, Z. M. Komodromos, and T. E. Humphreys, “An agile, portable antenna system for LEO megaconstellation-based PNT,” in Proceedings of the ION GNSS+ Meeting, 2023

3. W. Qin, Z. M. Komodromos, and T. E. Humphreys, “An analysis of the short-term time stability of the Starlink Ku-band downlink frame clock,” in Proceedings of the IEEE International Conference on Wireless for Space and Extreme Environments (WISEE 2024), 2024. Submitted for review.

4. A. M. Graff and T. E. Humphreys, “OFDM-based positioning with unknown data payloads: Bounds and applications to LEO PNT,” IEEE Transactions on Wireless Communications, 2024. Submitted for review.

5. Z. M. Komodromos, W. Qin, and T. E. Humphreys, “Signal simulator for Starlink Ku-Band downlink,” in Proceedings of the ION GNSS+ Meeting, pp. 2798–2812, 2023

Media Features

1. **Electronic Warfare Spooks Airlines, Pilots and Air-Safety Officials,** WALL STREET JOURNAL, September 2024, link
2. **Israel GPS 'spoofing' against missiles disrupts civilian life, aviation in Lebanon and Middle East**, ABC AUSTRALIA, link
3. **Why GPS Is Under Attack**, NEW YORK TIMES, July 2024, link
4. **An Israeli Air Base is a Source of GPS 'Spoofing' Attacks, Researchers Say,** NEW YORK TIMES, July 2024 link
5. **How GPS Warfare is Playing Havoc with Civilian Life**, FINANCIAL TIMES, May 2024, link
6. **The Dangerous Rise of GPS Attacks**, WIRED, April 2024, link
7. **Israel Fakes GPS Locations to Deter Attacks, but it also Throws off Planes and Ships**, NPR, April 2024, link
8. **FCC Probing US Phones That Use Signals From Foreign Satellites**, BLOOMBERG, March 2024, link
9. **The Sea Creatures That Opened a New Mystery About MH370**, NEW YORK MAGAZINE, March 2024, link
10. **War-Zone GPS Spoofing Is Threatening Civil Aviation**, FOREIGN POLICY, March 2024, link
11. **Unprecedented GPS jamming attack affects 1600 aircraft over Europe**, NEWSCIENTIST, March 2024, link
12. **GNSS Jamming and Spoofing Events Present a Growing Danger**, AVIATION INTERNATIONAL NEWS, March 2024, link
13. **Russia is using SpaceX's Starlink satellite devices in Ukraine, sources say**, DEFENSE ONE, February 2024, link
14. **As Baltics see spike in GPS jamming, NATO must respond,** BREAKING DEFENSE, January 2024, link
15. **GPS Spoofing Is Now Affecting Airplanes In Parts Of Europe**, FORBES, January 2024, link
16. **From Russia with love for Christmas: Jamming Baltic GPS**, GPS WORLD, January 2024, link
17. **Circle Spoofing Comes to Aviation – first the Baltic, now the Mediterranean**, RESILIENT NAVIGATION AND TIMING FOUNDATION, January 2024, link
18. **Air Travel Is Not Ready for Electronic Warfare**, NEW YORK MAGAZINE, January 2024, link
19. **GPS Spoofing in the Middle East Is Now Capturing Avionics**, FORBES, December 2023, link
20. **Disturbed GPS Signals Hinder Civil Aviation**, DER SPIEGEL," December 2023, link
21. **Commercial Flights Are Experiencing 'Unthinkable' GPS Attacks and Nobody Knows What to Do**, VICE, November 2023, link
22. **Israel's Using Widespread GPS Tampering to Deter Hezbollah's Missiles**, POLITICO, October 2023, link
23. **GPS Spoofing Thickens the Fog of War,** POLITICO, October 2023, link
24. **Israel Ramps Up GPS Jamming to Thwart Hezbollah, Hamas Drone Attacks**, HAARETZ, October 2023, link

Single-Satellite GNSS Spoofer Geolocation from Low Earth Orbit

Introduction

The combination of easily-accessible low-cost GNSS spoofers and the emergence of increasingly-automated GNSS-reliant systems prompts a need for multi-layered defenses against GNSS spoofing. GNSS spoofers broad- cast an ensemble of false GNSS signals intending that the victim receiver(s) will accept them as the authentic GNSS signals and subsequently infer a false position fix and/or a clock offset [1], [2]. A successful spoofing attack may lead to devastating consequences.

The academic community has long warned the public about the threat of GNSS spoofing [3]–[5]. Within the past decade, significant progress in has been made in onboard GNSS spoofing detection and mitigation [1]. Reliable spoofing detection techniques even exist for challenging environments such as dynamic platforms in urban areas where strong multipath and in-band noise are common [6]–[11]. Consistency checks between the estimated signal and onboard inertial sensors can provide quick and reliable spoofing detection [12]–[14]. Monitoring the clock state can also be used to detect spoofing [15]. Cryptographic authentication techniques are currently being developed and implemented [16], [17].

Although the recent advances in onboard GNSS spoofing detection have been inspiring, many of the older GNSS receivers in current operation are unable to incorporate these defenses, leaving them vulnerable to attacks. For example, the civilian maritime and airline industries are encountering GNSS jamming and spoofing at an alarming rate. Anomalous positioning information broadcast by ships in Automatic Identification System (AIS) messages, and airplanes in Automatic Dependent Surveillance-Broadcast (ADS-B) messages are indicative of wide-spread jamming and spoofing. Ships and airplanes near the Eastern Mediterranean, the Baltic region, and Shanghai have fallen victim to spoofing, as they seemingly teleport to new locations.

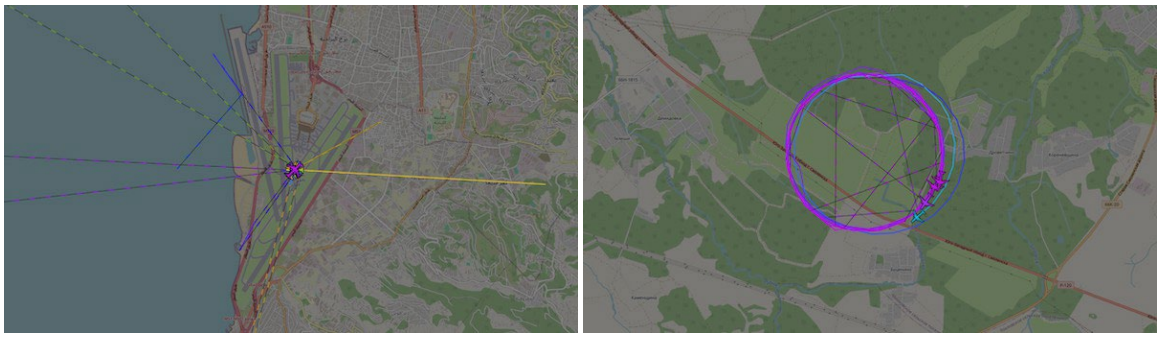

Fig. 1: Screenshots from the website ADS-B Exchange [18], the world's largest community of unfiltered ADS- B/Mode S/MLAT feeders. These screenshots show recent incidents of GNSS spoofing affecting aviation. The majority of spoofing induces false static locations, typically the location of airports (left). But within the past month, GNSS receivers have been spoofed to make them appear to be flying in circles (right).

Fig. 1 displays recent examples of aircraft ensnared by to GNSS spoofing, likely as unintended targets caught in the electronic warfare crossfire near ongoing conflict zones. From the ADS-B logs, one can infer that the most common spoofing attack is to transmit an ensemble of spoofing signals that is consistent with a single static location, typically a major airport. This type of spoofing is most likely used as a defense against commercial off- the-shelf drones, as these drones have built-in protocols to avoid protected airspace (e.g., surrounding airports). In a more alarming trend, spoofing attacks have caused aircraft to be spoofed in circular trajectories. These spoofers appear to not be targeting individual GNSS receivers, but rather broadcasting their signals for general GNSS denial. However, an attacker could in theory tailor a spoofing trajectory for a specific target, causing a gradual pull-off from its true trajectory, luring the victim into restricted airspace. Given that many currently-deployed GNSS receivers are unable to defend themselves even against easy-to-detect broad-area spoofing attacks, such targeted attacks are a clear and present threat.

The traditional approach for GNSS security has been to develop onboard receiver spoofing detection and mitigation techniques. The future of GNSS security takes a more active approach: global, accurate, and persistent localization of the emitters threatening GNSS receivers. The proliferation of LEO-based receivers provides unprecedented spectrum awareness, enabling GNSS interference detection, classification, and geolocation [19]–[24]. Dedicated LEO constellations provide worldwide coverage with frequent revisit rates, allowing for an always-updating oper- ating picture. Several commercial enterprises have seized the opportunity to deploy constellations of LEO satellites to provide spectrum monitoring and emitter geolocation as a service (e.g., Spire Global and Hawkeye360).

With multiple time-synchronized receivers, geolocation of emitters producing arbitrary wideband signals is possible and has been extensively studied [22], [23], [25]–[27]. Multiple time-synchronized receivers can exploit time- and frequency-difference-of-arrival (T/FDOA) measurements to estimate the emitter location. The authors of ththe current paper were able to geolocate over 30 GNSS interference sources across the Eastern Mediterranean and Ukraine from a dual-satellite time-synchronized capture [22], [23]. However, planning simultaneous multi- satellite captures to enable T/FDOA-based and direct geolocation can be difficult to coordinate and expensive, whereas single-satellite collects are straightforward and less costly. This paper focuses on single-satellite platforms.

Accurate single-satellite-based emitter geolocation is possible from Doppler measurements alone, provided that the emitter is transmitting at a quasi-constant frequency [20], [21], [28]–[30]. However, accurate single-satellite geolocation of emitters with arbitrary waveforms is impossible in general: if the signal's carrier cannot be tracked, only coarse received-signal-strength techniques can be applied for geolocation. In 2018, members of The University of Texas at Austin Radionavigation Lab (UT RNL) were able to geolocate a powerful 70-watt matched-code jammer operating in Syria to better than 300 meters using Doppler-based techniques [20]. One of the crucial assumptions of Doppler-based single-satellite geolocation is that the emitter transmits at a quasi- constant carrier frequency. Under this assumption, and assuming perfect stability of the receiver clock, the received Doppler is equivalently the range-rate, up to a constant bias and scaling. If a transmitter introduces any significant level of complexity to the carrier-phase behavior, such as frequency modulation or clock dithering, the accuracy of Doppler-based single-satellite techniques degrades.

GNSS spoofers must be treated differently, as they do not transmit at a constant center frequency: they add an extra unknown time-varying frequency component to each spoofed signal, imitating the range-rate between the corresponding spoofed GNSS satellite and the counterfeit spoofed location. This added unknown time-varying frequency component renders raw observed Doppler-based

geolocation for GNSS spoofers inaccurate. One of the key results in [21] is a technique that removes the unknown time-varying frequency component added by GNSS spoofers so that a range-rate time history can be extracted for geolocation. Furthermore, an analysis of how actual transmitter clock error and transmitter motion degrade the geolocation estimate is performed in [21]. A single-receiver spoofer geolocation technique based on counterfeit clock observables is also presented in [31], and makes a similar observation to [21], namely, that the spoofed clock bias of a mobile drone can be used for geolocation. However, [31] only considers the spoofed pseudorange measurements, whereas [21] and the current paper incorporate both pseudorange and Doppler measurements, and [31] depends on a static initialization period, which is not possible in LEO.

The key observation of [21] is that each spoofed navigation signal will share a common frequency shift due to the range-rate between the LEO receiver and terrestrial spoofer. If a GNSS receiver processes enough spoofing signals to form a navigation solution, the estimator will lump the common frequency shift of each signal from the shared range-rate into the receiver clock offset rate (clock drift) estimate. Therefore, the time history of the spoofed receiver clock offset rate can be exploited for geolocation because the range-rate between LEO receiver and terrestrial spoofer is embedded in this measurement.

This paper offers an experimental demonstration of the single-satellite single-pass geolocation technique intro- duced in [21]. This demonstration is the first of its kind in the public domain. In this experiment, conducted in partnership with Spire Global, an ensemble of self-consistent spoofing signals was transmitted from a ground station and captured by an overhead LEO receiver. The transmitted signals were centered at S-band to avoid interference in the GNSS bands and for FCC and ITU compliance. The GNSS spoofing signals were processed by the UT RNL GRID receiver [32], [33] to generate a clock offset rate time history, followed by geolocation of the GNSS spoofer.

GNSS Spoofing Signals

The goal of a broadcast GNSS spoofer is to deceive the victim receiver(s) into inferring a false position, velocity, and timing (PVT) solution, denoted $\tilde{\mathbf{x}} = \left[\, \mathbf{r}_{\tilde{\mathrm{R}}}^{\mathsf{T}},\ \delta t_{\tilde{\mathrm{R}}},\ \mathbf{v}_{\tilde{\mathrm{R}}}^{\mathsf{T}},\ \delta \dot{t}_{\tilde{\mathrm{R}}} \,\right]^{\mathsf{T}}$, where $\mathbf{r}_{\tilde{\mathrm{R}}}$ is the spoofed position in Earth-centered-Earth-fixed (ECEF) coordinates, $\delta t_{\tilde{\mathrm{R}}}$ is the spoofed clock bias, $\mathbf{v}_{\tilde{\mathrm{R}}}$ is the spoofed velocity, and $\delta \dot{t}_{\tilde{\mathrm{R}}}$ is the spoofed receiver clock drift. To achieve a successful attack, the spoofer must generate an ensemble of self-consistent signals. To this end, the attacker must (1) select a counterfeit PVT solution for the victim to infer, (2) select an ensemble of GNSS satellites to spoof, and (3) for each spoofed navigation satellite, generate a signal with a corresponding navigation message, code phase time history, and carrier phase time history consistent with (1) and (2).

A general baseband signal model for broadcast spoofing signals is presented below. The ensemble of spoofing signals transmitted by the spoofer, denoted

$$x(t) = \sum_{i=1}^{N} s_i\,(t)$$

contains N spoofing signals, with the ith spoofing signal denoted $s_i(t)$. The ith spoofing baseband signal takes the form

$$s_i(t) = A_i D_i[t - \tau_i(t)] C_i[t - \tau_i(t)] \exp[j2\pi\theta_i(t)]$$

where, for the ith signal, A_i is the carrier amplitude, $D_i(t)$ is the data bit stream, $C_i(t)$ is the spreading code, $\tau_i(t)$ is the code phase, and $\theta_i(t)$ is the beat carrier phase. The Doppler of the ith spoofing signal is related to $\theta_i(t)$ by

$$\tilde{f}_i(t) = \frac{d}{dt}\theta_i(t)$$

where $\tilde{f}_i(t)$ is the Doppler of the ith spoofing signal. The spoofer adds an unique Doppler component to each spoofing signal that mimics the combined Doppler of the following components: (1) the range-rate between the spoofed satellite and spoofed position, (2) the spoofed receiver clock drift, and (3) the spoofed satellite clock drift. Additionally, the spoofed code phase and carrier phase time histories must be mutually consistent to avoid code-carrier divergence. Accordingly, the Doppler of the ith transmitted spoofing signal may be modeled as

$$\tilde{f}_i(t) = \quad -\frac{1}{\lambda}\hat{\mathbf{r}}_{\tilde{\mathrm{R}},i}^{\mathsf{T}}(t)\left(\mathbf{v}_{\tilde{\mathrm{R}}}(t) - \mathbf{v}_{\tilde{\mathrm{S}},i}(t)\right) - \frac{c}{\lambda}\left(\delta\dot{t}_{\tilde{\mathrm{R}}}(t) - \delta\dot{t}_{\tilde{\mathrm{S}},i}(t)\right)$$

where λ is the carrier wavelength, c is the speed of light, $\hat{\mathbf{r}}_{\tilde{\mathrm{R}},i}$ is the unit vector pointing from the ith spoofed navigation satellite to the spoofed position, both in ECEF coordinates, $\mathbf{v}_{\tilde{\mathrm{R}}}$ is the spoofed receiver velocity, $\mathbf{v}_{\tilde{\mathrm{S}},i}$ is the ith spoofed navigation satellite velocity, and $\delta\dot{t}_{\tilde{\mathrm{S}},i}$ is the spoofed clock drift of the ith navigation satellite. One can immediately appreciate that the Doppler frequency is different for each spoofing signal. Had this been a targeted spoofer, there would be an additional Doppler term that compensates for the relative motion between the victim and spoofer, but in the case of broadcast spoofing, this term is zero.

Received Doppler Model

First consider a scenario in which a moving receiver captures a transmitted signal having a quasi-constant center frequency. The received Doppler $f_{\mathrm{D}}(t)$ at the moving receiver can be modeled as

$$f_{\mathrm{D}}(t) = \quad -\frac{1}{\lambda}\hat{\mathbf{r}}^{\mathsf{T}}(t)\left(\mathbf{v}_{\mathrm{R}}(t) - \mathbf{v}_{\mathrm{T}}(t)\right) - \frac{c}{\lambda}\left(\delta\dot{t}_{\mathrm{R}}(t) - \delta\dot{t}_{\mathrm{T}}(t)\right)$$

where $\hat{\mathbf{r}}$ is the unit vector pointing from the transmitter to the receiver, $\mathbf{v}_{\mathrm{R}}$ is the velocity of the receiver, $\mathbf{v}_{\mathrm{T}}$ is the velocity of the transmitter, $\delta\dot{t}_{\mathrm{R}}$ is the clock drift of the receiver, and $\delta\dot{t}_{\mathrm{T}}$ is the clock drift of the transmitter.

Now consider a scenario in which a moving receiver captures an ensemble of transmitted spoofing signals from a stationary terrestrial spoofer. An analysis of how spoofer motion affects the geolocation solution is given in a prior version of this paper. But would-be spoofers are typically stationary; otherwise, they face the additional difficulty of compensating for their motion to avoid producing easily-detectable false signals. Therefore, a stationary spoofer will be assumed for the rest of this paper.

Each observed signal at the receiver will contain a common Doppler shift f_{D} due to the relative motion between the transmitter (spoofer) and the receiver. Each observed signal will also manifest a common frequency shift due to the clock drift of the transmitter and the clock drift of the receiver. Dropping the time indices for clarity, the observed Doppler of the ith spoofing signal at the moving receiver, f_i, may be written as

$$
\begin{aligned}
f_i = \quad & f_{\mathrm{D}} + \tilde{f}_i \\
= \quad & -\frac{1}{\lambda}\hat{\mathbf{r}}^{\mathsf{T}}\mathbf{v}_{\mathrm{R}} - \frac{c}{\lambda}(\delta\dot{t}_{\mathrm{R}} - \delta\dot{t}_{\mathrm{T}}) \\
= \quad & -\frac{1}{\lambda}\hat{\mathbf{r}}_{\tilde{\mathrm{R}},i}^{\mathsf{T}}\left(\mathbf{v}_{\tilde{\mathrm{R}}} - \mathbf{v}_{\tilde{\mathrm{S}},i}\right) - \frac{c}{\lambda}\left(\delta\dot{t}_{\tilde{\mathrm{R}}} - \delta\dot{t}_{\tilde{\mathrm{S}},i}\right)
\end{aligned}
$$

What makes single-satellite GNSS spoofer geolocation difficult is the $\tilde{f}_i$ term: it is typically unknown, time-varying, and different for each spoofing signal. In the case of the matched-code jammer discovered in, $\tilde{f}_i = 0$. One may suppose that the operator's intent in this case was not to deceive victim receivers into inferring false locations like a spoofer. When $\tilde{f}_i = 0$, the observed Doppler can be modeled as the range-rate between transmitter and receiver, with a constant measurement bias over the capture to account for the clock drift of the transmitter. Contrariwise, naive geolocation with the observed Doppler yields final position estimates that are biased because the spoofing signals contain the unmodeled $\tilde{f}_i(t)$ term. In the following section, a technique is presented that removes $\tilde{f}_i(t)$ and extracts $\hat{\mathbf{r}}^{\mathsf{T}}(t)\mathbf{v}_{\mathrm{R}}(t)$, the range-rate time history between transmitter and receiver, which can be exploited for geolocation.

Conceptual Overview of Broadcast GNSS Spoofer Geolocation

As discussed, the observed Doppler of the ith spoofing signal is a combination of the physical range-rate between the the transmitter and receiver, and a Doppler component that mimics the motion between the ith spoofed satellite and the spoofed position and velocity. This section presents and overview of the technique originally presented in for spoofer geolocation The common Doppler components across all spoofing signals are indicated below:

$$
\begin{aligned}
f_i = \quad & \underbrace{-\frac{1}{\lambda}\hat{\mathbf{r}}^{\mathsf{T}}\mathbf{v}_{\mathrm{R}} - \frac{c}{\lambda}(\delta\dot{t}_{\mathrm{R}} - \delta\dot{t}_{\mathrm{T}})}_{\text{common}} \\
= \quad & -\frac{1}{\lambda}\hat{\mathbf{r}}_{\tilde{\mathrm{R}},i}^{\mathsf{T}}\left(\mathbf{v}_{\tilde{\mathrm{R}}} - \mathbf{v}_{\tilde{\mathrm{S}},i}\right) - \frac{c}{\lambda}\left(\underbrace{\delta\dot{t}_{\tilde{\mathrm{R}}}}_{\text{common}} - \delta\dot{t}_{\tilde{\mathrm{S}},i}\right)
\end{aligned}
$$

All common Doppler terms can be lumped into a single term

$$
\gamma(t) = \frac{1}{c}\hat{\mathbf{r}}^{\mathsf{T}}(t)\mathbf{v}_{\mathrm{R}}(t) + \delta\dot{t}_{\mathrm{R}}(t) - \delta\dot{t}_{\mathrm{T}}(t) + \delta\dot{t}_{\tilde{\mathrm{R}}}(t)
$$

so that it may be written

$$
f_i = -\frac{1}{\lambda}\hat{\mathbf{r}}_{\tilde{\mathrm{R}},i}^{\mathsf{T}}\left(\mathbf{v}_{\tilde{\mathrm{R}}} - \mathbf{v}_{\tilde{\mathrm{S}},i}\right) - \frac{c}{\lambda}\left(\gamma - \delta\dot{t}_{\tilde{\mathrm{S}},i}\right)
$$

If a GNSS receiver were to process the ensemble of spoofing signals, its PVT estimator would infer the state $\mathbf{x}(t) = \left[\ \mathbf{r}_{\tilde{\mathrm{R}}}^{\mathsf{T}}(t),\ \delta t_{\tilde{\mathrm{R}}}(t),\ \mathbf{v}_{\tilde{\mathrm{R}}}^{\mathsf{T}}(t),\ \gamma(t)\ \right]^{\mathsf{T}}$, which is composed of the spoofed position, spoofed clock bias, spoofed velocity, and the new receiver clock drift $\gamma(t)$. The apparent clock drift $\gamma(t)$ has units of s/s and contains all common Doppler terms. The PVT estimator attributes common-mode frequency deviations across received signals to the receiver's clock drift. A brief review of PVT estimation from pseudorange and Doppler measurements is provided in . At each navigation epoch, the PVT estimator produces an optimal estimate of $\gamma(t)$. Importantly, $\gamma(t)$ is unaffected by the unknown non-common Doppler components from $\tilde{f}_i(t)$ for all $i \in \{1, \ldots, N\}$.

The time history $\gamma(t)$ (or its estimate) can ultimately be used for geolocation because it depends strongly on the range-rate between the LEO-based receiver and the terrestrial spoofer. In particular, information about the transmitter's location is embedded in the time history $\hat{\mathbf{r}}^{\mathsf{T}}(t)\mathbf{v}_{\mathsf{R}}(t)$. Based on the range-rate measurement model, a nonlinear least-squares estimator for the time history of $\gamma(t)$ is developed in the next section to estimate the spoofer's position.

For a targeted spoofing attack, there would be an extra Doppler term in $\gamma(t)$ that corresponds to the motion between the targeted victim and the spoofer. This term could potentially cause trouble for this paper's technique. If the relative velocity between the targeted victim and the spoofer over the capture is small compared to the relative velocity between the LEO receiver and the spoofer—a likely situation given that LEO orbital speeds are extreme compared with almost any terrestrial or near-Earth vehicle—then this paper's technique would suffer only a minor degradation in accuracy. Similarly, if the targeted victim's position and velocity were somehow accurately known to the LEO-based receiver, this technique would still work, but alterations to the estimator presented in the next section would have to be made. Finally, if the targeted victim receiver is stationary, this paper's technique can be applied without modification.

The other three terms in $\gamma(t)$, namely $\delta\dot{t}_{\mathsf{R}}(t)$, $\delta\dot{t}_{\mathsf{T}}(t)$, and $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ are nuisance terms that potentially degrade geolocation accuracy. Fortunately, their contributions are typically minor or can be estimated. Consider $\delta\dot{t}_{\mathsf{R}}(t)$. If the satellite's GNSS receiver and the radio frequency (RF) front-end capturing spoofing signals are driven by the same oscillator, $\delta\dot{t}_{\mathsf{R}}(t)$ is estimated by the onboard GNSS receiver and can be compensated.

It is worth mentioning that one of the core assumptions in any geolocation system is that the capture platform has knowledge of its PVT, otherwise, geolocation is impossible. In this scenario, the LEO-based receiver has access to its PVT from an onboard GNSS receiver that is robust to terrestrial interference. Despite the presence of spoofing signals, code- and carrier-tracking of the authentic GNSS signals is maintained due to sufficient separation of the false and authentic signals in the code-Doppler space. Furthermore, robustness is achieved if a zenith-facing antenna feeds the onboard GNSS receiver's RF front-end, as the gain pattern towards Earth will be limited. Finally, PVT can be trivially maintained by a multi-GNSS receiver when only single-constellation spoofing signals are present.

The terms $\delta\dot{t}_{\mathsf{T}}(t)$ and $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ originate from the spoofer. Specifically, $\delta\dot{t}_{\mathsf{T}}(t)$ originates from the spoofer's hardware, while $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ originates from the spoofer's software. $\delta\dot{t}_{\mathsf{T}}(t)$ arises due to the clock drift in the spoofer. It can often be accurately modeled as constant over short (e.g., 60-second) capture intervals and estimated as part of the geolocation process. The spoofed clock drift $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ arises from the spoofer's attack configuration. It $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ can be troubling for geolocation, but a potential attacker would typically opt to keep $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ near constant, because if $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ grows too rapidly to be explained by the expected variation in clock drift for the receiver's oscillator type, the victim receiver could flag the anomaly and thereby detect the spoofing attack.

This constraint can be generalized to the sum $\delta\dot{t}_{\mathsf{T}}(t) + \delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ and summarized as follows: if the spoofer allows extraordinary frequency instability in its own oscillator so that $\delta\dot{t}_{\mathsf{T}}(t)$ changes too rapidly, or if it attempts to induce a quickly-varying spoofed clock drift so that $\delta\dot{t}_{\tilde{\mathsf{R}}}(t)$ changes too rapidly, geolocation accuracy is degraded but, on the other-hand, the spoofing attack becomes trivially detectable.

Experimental Spoofer Geolocation with γ(t)

The transmitted spoofing signals were captured by the LEO-based receiver and processed with the UT RNL's GRID software-defined GNSS receiver. GRID was also able to compute the spoofed PVT solution corresponding to the top of the Aerospace Engineering building. The position solution is slightly biased, but that is explained by the code-carrier divergence from the S-band carrier . On GRID's display, the large Doppler shifts across each signal and the 4,810 meter per second clock drift are immediately noticeable. Of course, no such oscillator on a GNSS receiver would experience a clock drift that extreme.

To coax GRID into properly processing the S-band spoofing signals, special modifications to the receiver's configuration and estimator had to be made. Reconfiguring such parameters is trivial withing GRID's software-defined architecture. The bandwidth of the receiver's delay lock loop (DLL) and phase lock loop (PLL) were increased to help maintain lock despite the code-carrier divergence introduced by the S-band carrier. The bandwidth of the DLL was set to 1.5 Hz and the bandwidth of the PLL was set to 40 Hz, introducing more noise. Furthermore, changes to the receiver's clock model had be made to account for the quickly drifting receiver clock, induced by the range-rate between LEO-based receiver and terrestrial transmitter. The clock model used set h_0=5×10^(-18) and h_(-2)=3×10^(-18), allowing the estimator to accept the quickly-varying clock drift.

Fig. 2: Shown here is the Doppler time history of each received spoofing signal. Also shown is the equivalent Doppler of γ(t) in black, which is used for geolocation.

Given all of this, γ(t) could be calculated over 17.75 seconds, and is shown in Fig. 2, along with the raw observed Doppler of each spoofing signal. The GNSS receiver allowed itself to be spoofed and the true range-rate between LEO-based receiver and terrestrial transmitter was lumped in the receiver's clock drift estimate.

Fig. 3: Top: Experimental setup. The groundtrack of the LEO-based receiver is shown as well as the transmission site. Bottom: Final spoofer position estimate (white) using γ(t). Shown in red is the true spoofer location. The error of the final estimate is 68 m. The emitter is contained within the 95% error ellipse, which has a semi-major of 6.7 km.

The time history of γ(t) was fed to the nonlinear least-squares estimator and the final position fix is shown in Fig. 3. The final position error was 68 meters, and most importantly, the true emitter position lay within the horizontal 95% error ellipse. The error ellipse is highly eccentric, but the error ellipse's eccentricity is dictated by the receiver-transmitter geometry. The Doppler post-fit residuals are zero-mean with a standard deviation of .12 m/s, which is exactly what is expected from a properly modeled system. This unprecedented experiment verifies this paper's developed technique.

Conclusion

This paper presented and verified single-satellite single-pass geolocation technique specifically for GNSS spoofers from LEO. The developed technique removed the unknown time-varying frequency component across each spoofing signal so that the range-rate time history between receiver and spoofer could be extracted and exploited for geolocation. This was accomplished by processing the spoofing signals and extracting a time history of the receiver clock drift. This paper also detailed a controlled experiment in partnership with Spire Global, in which a LEO-based receiver captured GNSS spoofing signals transmitted from a known ground station on a non-GNSS frequency band.

An Agile, Portable Antenna System for LEO Megaconstellation-Based PNT

Introduction

Recent research has shown increased interest in using low Earth orbit (LEO) satellite constellations as a means to obtain position, navigation, and timing (PNT) solutions. As noted in [37]–[39], LEO PNT offers improved robustness and redundancy against jamming and anti-satellite warfare, as well as more precise positioning when compared to traditional MEO approaches. As no such dedicated LEO PNT constellation exists yet, researchers have instead begun to opportunistically exploit commercial LEO communications. For instance, recent work has already demonstrated the feasability of Doppler-based positioning based on observations of signals sent by Starlink's constellation of over 3,000 LEO satellites [40]. However, despite its relative merits over Doppler-based positioning in terms of accuracy and time-to-fix, opportunistic pseudorange-based PNT remains an unproven solution. For instance, consider that the 240 MHz per-channel Starlink bandwidth possibly offers nanosecond- accurate timing, an improvement upon the 15 ms accuracy offered by Doppler-based positioning [40]–[42].

It is well-known that the accuracy of pseudorange-based PNT improves with the reduction in the dilution of precision [43]. Generally speaking, as the number of available pseudorange measurements to unique SVs increases, so too does the estimation problem's geometric strength, in turn enabling a more accurate PNT solution. Minimally,if d is the number of physical dimensions under consideration, including time, $n \geq d$ pseudorange measurements are required to obtain an observable set of coordinates. Yet multiple simultaneous pseudorange measurements using signals of opportunity has yet to be shown, partly due to the unique challenges encountered when non- cooperatively acquiring LEO communications signals.

These challenges are as follows: First, unlike traditional GNSS satellites that use publicly-known sequences to aid signal acquisition and tracking, LEO communications megaconstellations typically employ proprietary signal structures known only to the companies designing them. In fact, Starlink's basic signal structure was only recently unveiled, with significant work still remaining regarding additional exploitable signal features and signal timing [42]. As such, despite LEO satellites enjoying signal-to-noise ratios (SNR) that can be 30 dB higher than those of traditional GNSS satellites [44], signal acquisition remains difficult without pointing directly at an SV that is also directly transmitting to the receiver's location. Second, efficient operation of a satellite communications constellation dictates that each SV beamforms to cover a particular region on the ground, known as a service cell. The particular service cell that is covered hypothetically changes with user demand and SV movement as the SV passes over user terminals, among other factors. Consequently, from a non-cooperating receiver's perspective, it becomes difficult to uniquely determine which satellite(s) are beamforming towards the receiver, from the set of SVs passing overhead at any given time. Third, the fact that LEO satellites occupy orbits with altitudes much closer than those of traditional GNSS satellite constellations (∼550 vs. ∼20,200 km) means faster dynamics and a higher Doppler effect. In turn, the signal acquisition and lock problems become more challenging.

A simple alternative to receiver beamforming would be to employ signal processing techniques to acquire multiple low-SNR signals with a rigidly-affixed antenna, in the style of traditional GNSS signal tracking. However, the predictable portion of the Starlink frame signal is currently limited to a small

fraction of each frame, requiring a received SNR of at least -15 dB [45]. As such, the system described in this paper points the antenna to relax the requirement on known signal proportion, and leaves a mechanically simpler system for later development.

As SpaceX's Starlink megaconstellation currently boasts the most mature deployment amongst LEO broadband communications networks, a single receiver often has simultaneous line-of-sight access to multiple Starlink SVs. Thus, the Starlink constellation offers the best target for testing a rapid-switching antenna pointing system.

This paper's approach uses a dual-axis mount to focus an attached antenna on a Starlink SV passing overhead. With a matched filter, downstream signal processing performs acquisition of the received transmission using the predictable portion of the signal structure, as known from [42], and extracts a time of arrival (TOA) measurement.

Terminology

It will be useful to define several new terms related to opportunistic use of broadband LEO signals for PNT. As described in [39] and [42], Starlink Ku-band downlink signals are sent via directional beams from overhead satellites on one of eight 240-MHz channels. Similarly, OneWeb signals are sent via beam-channel combinations [46].

Assigned Beam: Beam directed by an SV toward the service cell in which the user's receiver is located.

Side Beam: Beam directed by an SV toward a service cell other than the one in which the user's receiver is located. Side-beam signals received by the user's receiver will be weaker than assigned-beam signals, but they may still be powerful enough to obtain accurate TOA measurements.

This paper makes three contributions. First, it provides a selection of known Starlink constellation quantities and qualities to inform platform design and experimental methodology. Second, it presents the new capture platform in terms of its hardware and software, before detailing data capture methods employed for the purpose of this study. Finally, it presents new results obtained with the proposed platform.

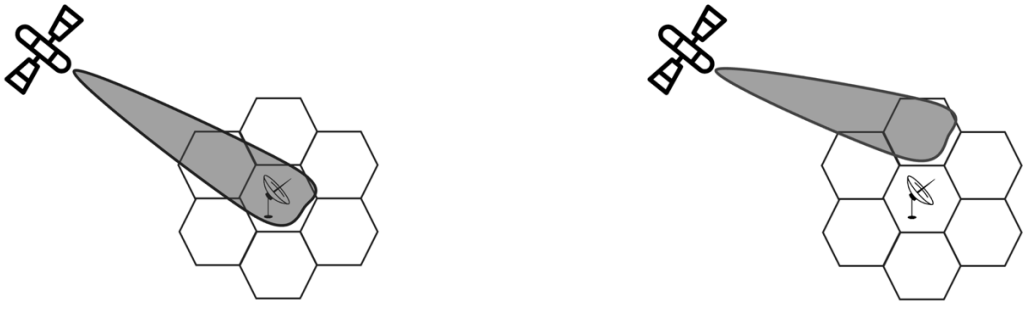

Fig. 4: Illustrated example of assigned and side beams. In each subfigure, the hexagonal grid represents a distribution of service cells on the ground. The antenna in the center represents an opportunistic receiver. If the SV transmits to the service cell occupied by the receiver, as is shown on the left, then we refer to the beam as an assigned beam. If the SV transmits to any other service cell, such as the neighboring service cell as shown on the right, then we refer to the beam as a side beam.

Starlink Characteristics

In support of the platform design and open questions discussed in the remainder of this paper, this section details various pertinent quantities and qualities of the Starlink constellation. It first focuses on important parameters and features of the Starlink signal structure, then introduces some new terminology to describe useful transmission patterns of the Starlink constellation.

Signal Structure

This paper adopts the terminology and Starlink-specific parameters as presented in [42]. For convenience, we provide a brief summary of the important signal properties.

Fig. 5: Starlink frame structure along time-frequency dimensions for a single channel (from [42]).

As shown in Fig. 7, each Starlink frame consists of 302 non-zero intervals of equal duration and one empty frame guard interval, all sent sequentially in time. A single full Starlink frame length takes T_"f" =1/750 s to send. While the first interval contains a time-domain sequence, all the following intervals contain OFDM symbols with N=1024 orthogonal subcarriers per symbol. The first two intervals are common to all frames sent by all SVs, referred to as the primary and secondary synchronization sequences (PSS and SSS, respectively). These known sequences can be generated according to [42]. The OFDM subcarriers collectively occupy a channel bandwidth of 240 MHz; Starlink has eight such channels laid out across the 10.7-12.7 GHz frequency band. It should be noted that the first two channels are yet unused.

These synchronization sequences are remarkable: similar to how a traditional GNSS receiver uses civil spreading codes, an opportunistic receiver can use the PSS and SSS to construct a local replica. Correlation against such a local replica yields TOA and Doppler measurements, the building blocks of PNT.

Beamforming Overview

Opportunistic use of Starlink signals carries a particular set of challenges when attempting to uniquely determine which satellites are transmitting to a given user service cell. Although dozens of satellites might be passing overhead at any time and each satellite can form up to 48 Ku-band downlink beams

[46], [47], any given service cell may only be targeted by one or two satellites [39]. Limitation to these few satellites causes the PNT problem geometry to become poor or even underdetermined. If possible, the ability to also acquire pseudoranges from SVs from outside this set could greatly increase the degree of multilateration, thereby recovering PNT accuracy.

Possibly, the use of known sequences to exploit signals destined for neighboring service cells could aid in this regard. Public filings such as [48] indicate the Starlink network transmits to individual hexagonal service cells about 20 km wide; however, the transmitted power does not necessarily fall sharply to 0 at service cell edges, especially when the satellite is far from the nadir position. This possibility is supported by a an example cross-correlation with the combined PSS and SSS.

We introduce the following terminology for the remainder of this paper. First, we use the term ”assigned beam” to describe any beam primarily directed at the service cell in which a given receiver resides. These beams tend to possess stronger SNR and often have observable power signatures with little to no signal processing required. Any of the eight Starlink channels that the beam occupies will be referred to as ”assigned channels.” In contrast, ”side beams” are designated as any beams that are not pointed to the occupied service cell. These beams tend to be weaker than assigned beams, but their transmissions can sometimes be used for pseudorange measurements after correlation with a local replica. Accordingly, ”side channels” represent any channels via which a given beam occupies.

To facilitate efficient data transmission to and from Starlink users, one can expect that beam and channel assignments change as the transmitting satellites pass overhead. To an opportunistic receiver, the expected assignment state after a change is not predictable; yet, the timing of the change is.

Observations made over the past year indicate that any beam and channel switching occurs at regular, predictable intervals. Beginning with the GPS second reset per week, all beam and channel switching occurs every 15 GPS seconds. To be clear, an assignment change is not guaranteed at every 15-second increment, but if one does occur, it will be on said increment. It then follows that no changes will occur within the 15-second interval between increments. As such, the 15-second intervals will be referred to in this paper as ”fixed assignment intervals.”

Signal Capture Platform

This section introduces the agile antenna system developed over the past several months for the purpose of this and future studies. It begins with a description of the hardware and provides major physical parameters considered in the mechanical design. It then details the software front-end and back-end.

Fig. 6: Physical hardware components used to construct the receiver assembly.

Hardware

We used an Interbotix WidowX Dual XM430 pan-tilt mount as the assembly actuator. A picture of the device is shown in Fig. 8. Three Dynamixel XM430 DC motors collectively move the turret - while only one motor is responsible for azimuth movement, two are responsible for elevation movement. Each motor has a stall torque of 4.10 Nm and a no-load speed of 276 deg/s, and both of the elevation motors were required for stall-free actuation of the mounted feedhorn [49]. The system receives using an NSI-MI manufactured standard gain pyramidal feedhorn fixed atop the pan-tilt turret with 80/20 T-slot alumnimum prototyping material. The antenna weighed 1.2 kg, had a 24.4 dBi nominal gain, and was tuned to the frequency range 10.0-15.0 GHz, which covered the Starlink allotted frequency bands of 10.7-12.7 GHz [42], [50], [51]. Again using 80/20 aluminum prototyping stock, we also constructed a wide base under the turret for stabilization under rapid switching from satellite to sallite. To maintain kinematic range of motion, the base was constructed with sufficient height to ensure the extremities of the low-noise block attached to the feedhorn would not collide with the ground in any orientation. To increase the maximum pointing elevation, the pan-tilt turret was cut and remounted onto the assembly base such that it no longer lay in the path of the low-noise block as the antenna precessed relative to the base.

The main advantage of the current capture equipment when compared to those used for [42] was its switching time. Whereas previous equipment often took up to 30 seconds to change observed satellites, the current antenna consistently switches in less than a second. This allowed for more satellites to be observed in a short period of time. Further, the antenna continuously tracks observed satellites with errors consistently less than 1 degree, even when satellites passing close to zenith (i.e. elevations close to 90 degrees) demanded relatively high azimuth rates.

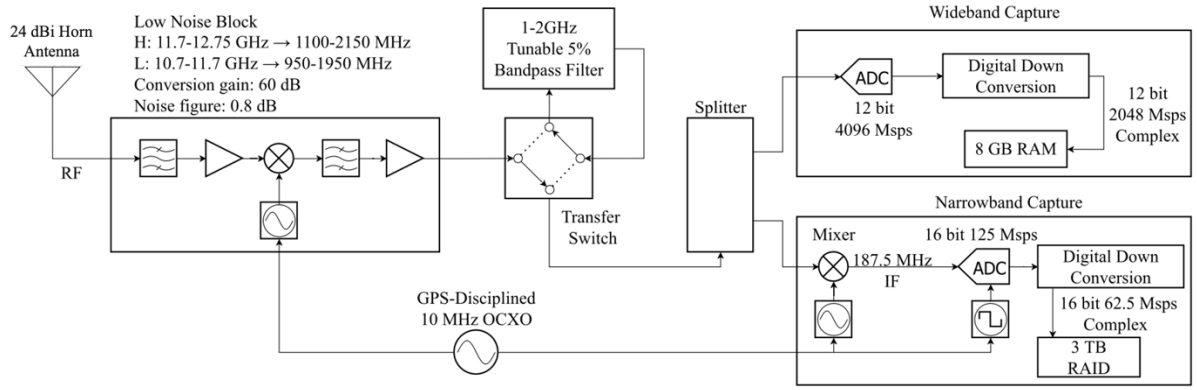

Fig. 7: Block diagram of the capture equipment used in this study.

Fig. 7 illustrates the capture equipment used, which is largely unaltered from that used in [42]. The pyramidal feedhorn first feeds signals to a low-noise block with a conversion gain of 60 dB and a noise figure of 0.8 dB [52].

There are two main capture modes available: a narrowband mode and a wideband mode. The narrowband mode offers data at a 16-bit complex sampling, but can only capture at 62.5 Msps. Conversely, the wideband mode does not offer extended data capture, but does allow a live spectrogram view that covers a total bandwidth of just greater than 1 GHz, a useful feature for making power-based observations of channel occupancy.

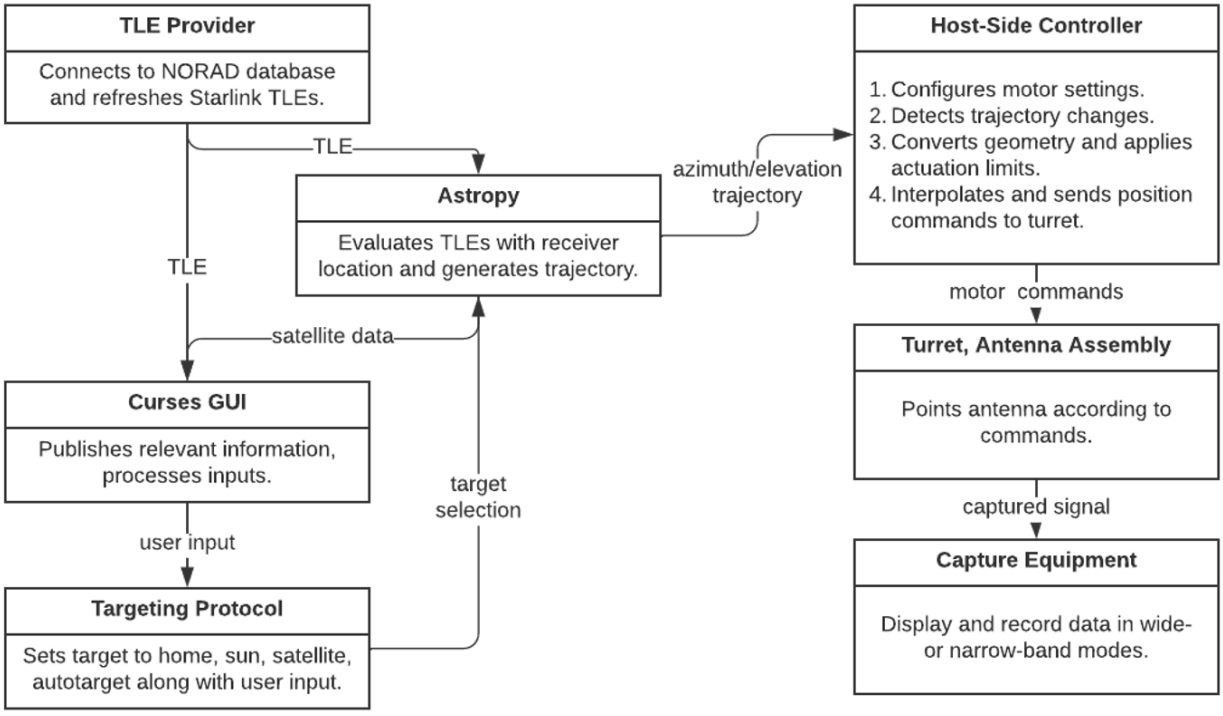

Fig. 8: Block diagram of software used to plan trajectories, track satellites, and actuate the antenna assembly.

Software

The software consists of two major modules: a trajectory manager and a motor controller as shown in Fig. 8. The modules run mutually asynchronously and communicate via shared storage. This section will present these modules separately.

The trajectory manager has several jobs. It (1) propagates orbits, (2) processes user inputs, (3) presents a graphical interface, (4) runs a targeting protocol, and (5) submits trajectories produced to the motor controller.

The trajectory manager first downloads all Starlink two-line elements (TLEs) from CelesTrak NORAD auto- matically on a daily basis, and then propagates the TLEs using an SGP4 simplified perturbation model. This industry-standard propagation technique has an error of approximately 1 km at epoch and grows between 1 and 3 km a day [53]. Though this simplified model does not provide trajectories with errors small enough for PNT, this does not prevent tracking of a satellite using a directed antenna receiver. Under worst-case scenarios, the

approximate upper bound on the tracking error due to orbit propagation is at most ∼0.42 degrees. This orbital propagation error is tolerable as the feedhorn has a beamwidth on the order of several degrees and the Dynamixel motor has a 0.25-degree backlash [54].

The trajectory manager employs a Python module called curses to drive the text-based user interface (TUI) and input handling. When compared to full graphical user interfaces (GUIs) such as PyQT, TUIs prove to be lightweight, easily adaptable, and quickly reconfigurable, all favorable qualities for rapid prototyping and automation. In our implementation, this allowed our combined propagation and interface loop to run consistently at 20 Hz, enabling live status updates and rapid actuator response to user commands. The quick access list of keyboard commands at the TUI's bottom is not pictured. Several important hotkeys initiate a tracking quickchange, toggle the autotargeting mode, and save the current configuration to storage. Ultimately, whether in automatic or manual targeting mode, the interface selects a satellite to track and provides the corresponding trajectory to the motor controller.

Upon software start, the motor controller conducts setup routine procedures, such as declaring the required Robot Operating System (ROS) publishers and subscribers for communicating with the pan-tilt mount and setting the PID gains. As the program runs and with each trajectory received, the motor controller processes the new data, translating azimuth-elevation pointing commands to pan-tilt motor commands based upon the known turret geometry. At every timestep, the controller linearly interpolates the motor command time history to the current time before instructing hardware to execute the command.

Satellite Tracking Algorithm

We developed the autotargeting algorithm to: (1) construct an initial upper limit for the number of satellites transmitting to a single service cell with assigned beams and (2) establish a simple baseline heuristic method that identifies satellites with assigned beams. It should be noted that as the capture hardware currently limits our data processing method to a live wideband spectrogram rather than cross-correlation-based detection, investigations of side beam occupancy are yet unviable. Nonetheless, power-based channel occupancy investigations of assigned beams are still viable.

To reiterate, the autotargeting algorithm explores the number of simultaneously assigned beams available to a receiver. Note, however, that simultaneous observation of multiple satellites using a narrow-beam antenna is generally unfeasible. The algorithm mitigates this shortcoming by leveraging the FAI. From the observation that no alterations in beam or channel assignments occur during the same FAI, we conclude that a single detection of a satellite assigned beam assures the assignment of the satellite for the duration of the entire FAI. Consequently, the total count of individual satellites with assigned beams observed is a sufficient substitute for the number of concurrently assigned beams. This, in turn, motivates us to maximize the number of satellites visited within a single 15-second interval.

The maximum number of satellites we can visit within the same FAI is limited by the dwell and switch times. If real-time matched filtering were possible, then we could assume only a short dwell time (<1 s) is required to determine the presence of an assigned beam. Yet, given that the observational methods involved human interpretation of a live spectrogram, we require a much more substantial dwell time of several seconds to either validate or dismiss channel occupancy. This imposed a limitation on the number of unique satellite observations per FAI. When considering three satellites for observation per FAI, the allotted combined dwell and switching time becomes 5 seconds, below which we would expect the rate of detection errors to rise substantially. Consequently, we set the number of satellites visited for a given FAI m=3 for this study.

Observations informed us that satellites with high signal strength and frequent transmissions tended to occupy elevated positions and exhibited prolonged presences in the sky, as opposed to satellites that barely skim our field of view at low, near-horizon elevations. Accordingly, we adopted a basic heuristic based on the integration of elevation versus time during the current FAI.

Let I be the set of all SVs above the horizon, and J be all the visited SVs in the kth FAI. Then, let θ_((t)) represent the elevation time history of the ith satellite and t_k the start time of the kth FAI. We then used the simple heuristic to determine which three satellites to observe during each FAI.

$$i_{\text{next}} = \text{argmax}_{i \in \mathcal{I} \setminus \mathcal{J}} \int_{t_k}^{t_{k+1}} \theta_i\,(t) dt$$

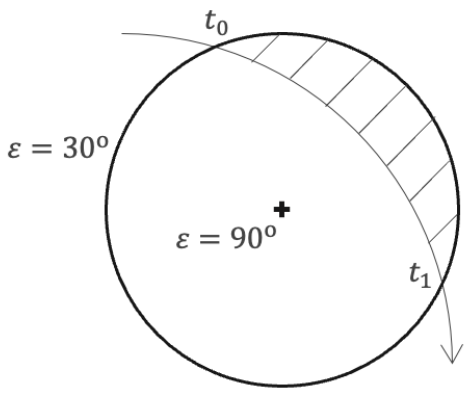

Fig. 9: Diagram of satellite track crossing the sky, where the ribbed area represents the heuristic used to determine satellite tracking patterns.

Results

Two sets of results were produced using the narrowband and wideband capture modes. First, we validated the tracking of a Starlink satellite despite comprehensive hardware and software updates. This confirmation was achieved via a cross ambiguity function (CAF) of known Starlink signal portions, the PSS and SSS, against a captured signal. Subsequently, we investigated the number of concurrently assigned beams within a given FAI.

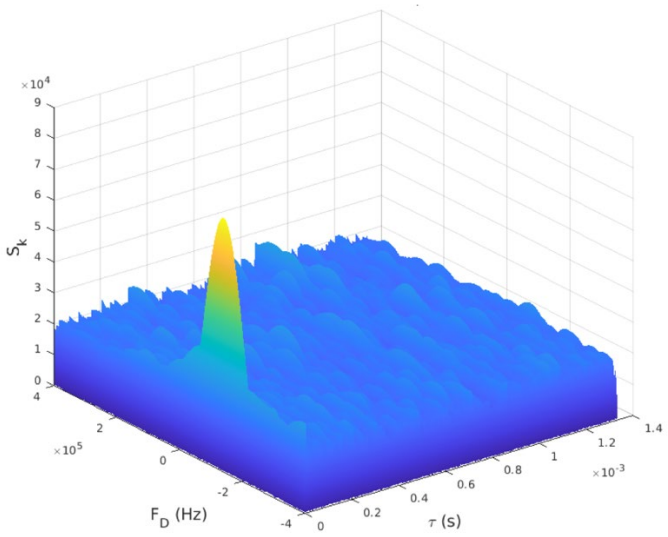

Fig. 10: Cross-ambiguity function results. The peak in cross-correlation at FD ≈ −10 kHz and τ ≈ 0.05 s proves successful capture of a Starlink signal.

$$S(f,k) = \sum_{n} r(n+k)c^{*}(n)e^{-2i\pi f \Delta t}$$

The equation above shows the cross-correlation of the received signal r(n) delayed by k samples against a local replica c(n). f is the Doppler frequency in Hz, and Δt represents the sample interval. As f approaches the true Doppler experienced and k results in a TOA measurement close to the true TOA, the cross-correlation between r and c constructively interfere, giving rise to a higher result S. After repetition across a wide search grid of varying Dopplers and TOAs, we can construct a CAF as shown in Fig. 10. Given that the local replica used was the time-concatenated PSS and SSS, the existence of a significant relative peak in the CAF proves successful acquisition of a Starlink satellite.

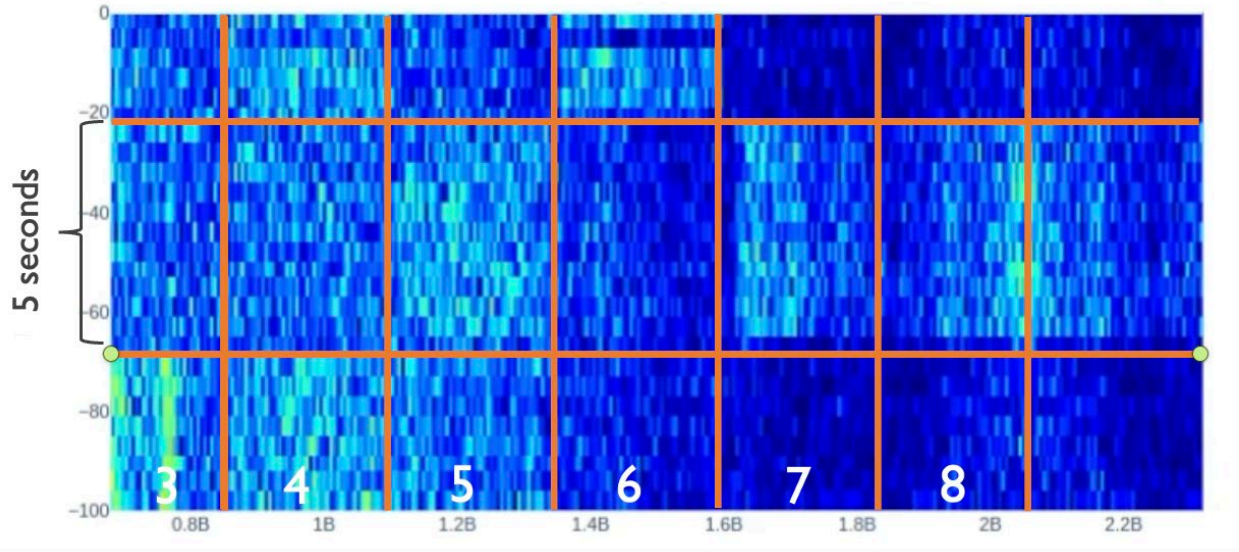

Fig. 11: Example of an active portion of a spectrogram. Six Starlink channels can be seen, from channel 3 through channel 8. The horizontal lines indicate the points at which the autotargeting algorithm changed satellites. This freeze frame shows a recorded observation of three unique satellites concurrently transmitting assigned beams across six channels.

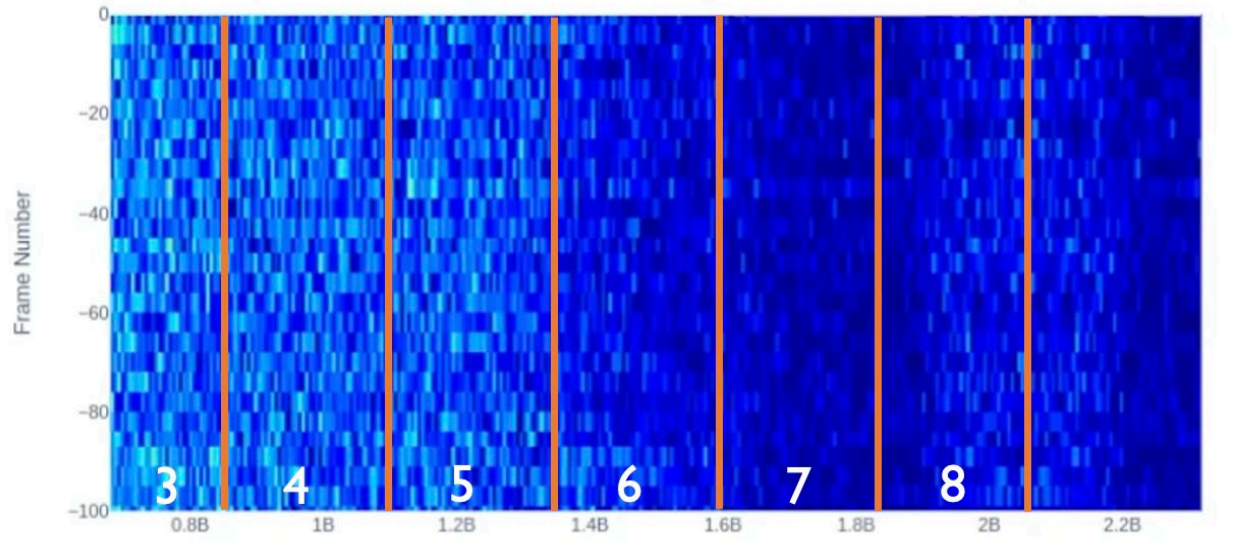

Fig. 12: Example of a spectrogram with no activity. It should be noted that the power is not uniform across the entire 1 GHz bandwidth, which can be partially attributed the antenna gain pattern.

To investigate the number of concurrently assigned beams and to measure the viability of our heuristic H, the turret assembly was run autonomously. For each of these 120 FAIs, the number of satellites with assigned beams was counted, up to the three visited per FAI. Figs. 11 and 12 show example

spectrograms at moments with full and absent assigned beam occupancy, respectively. Fig 13 is a bar chart showing the distribution of concurrent assigned beams per FAI.

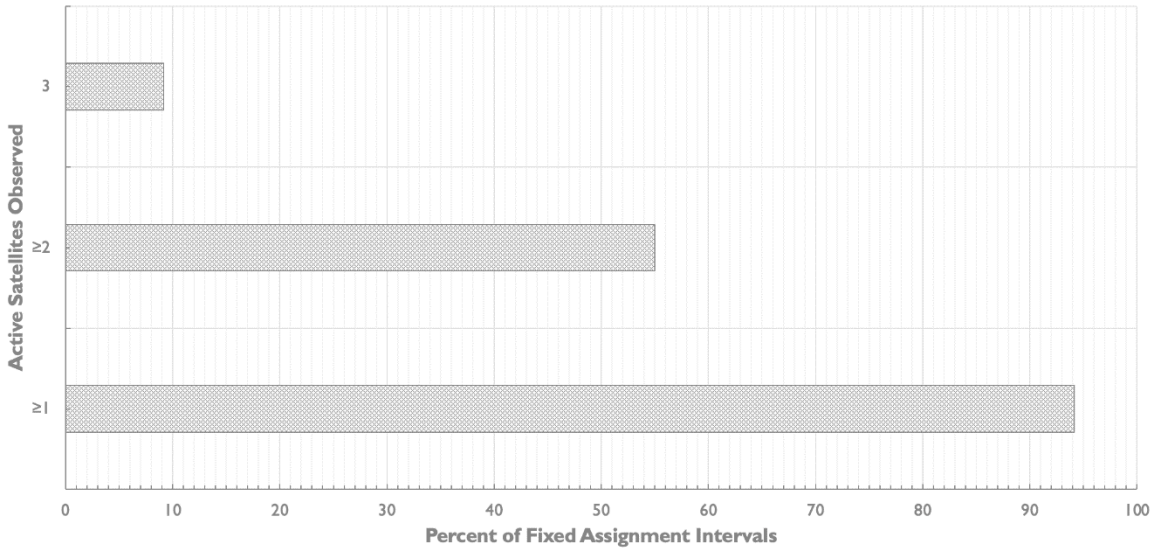

Fig. 13: Bar chart indicating the number of concurrent assigned beams observed per FAI, up to three.

Aggregated, the heuristic was successful in detecting at least one satellite per FAI in 94% of the cases, at least two satellites per FAI in 55% of cases, and all three guesses were successful in 9% of cases. This not only confirms the basic viability of the integrated elevation heuristic upon which more work can be done, but also shows that the potential number of main beams per FAI is at least three.

Discussion and Conclusion

In this study, a new set of capture equipment was designed, assembled, and tested to pursue a set of open questions whose answers will enable efficient operation of an opportunistic LEO PNT system focused on Starlink broadband transmissions. Briefly restating these open questions, we wish to find: (1) the number of unique SVs a user could possibly use for TOA measurement, (2) a simple pattern to predict assignment triads, and (3) a range of absolute received power into a hemispherical antenna from an assigned or side beam. This section will discuss the degree to which the current results have already answered these questions, and the next steps we are pursuing to fully address them.

With regards to the number of unique SVs from whose transmissions TOA measurements can be made, we confirmed the current upper bound to be three unique SVs transmitting assigned beams. With more exploration, we expect the number of usable assigned to rise, as well as the additional inclusion of the number of usable side beams.

We introduced a simple heuristic to help predict assignment triads per FAI. It acquires at least one assignment triad per FAI for 94% of FAIs. It should be noted, however, that observations were made in a populated urban area where demand is possibly higher than average. Users in areas with lower demand should expect differences in this measurement. Nonetheless, this heuristic provides the basis for further development.

While results for assigned and side beam power estimation have not yet been completed, work is underway. Using pyramidal feedhorn gain estimation techniques such as those presented in [55], [56], we can use known parameters of the capture system, Starlink communications system, and signals aggregated over a sufficient number of captures to estimate the expected power per unit area. This figure will allow us to determine if the required SNR is available to support wide beam capture for simultaneous acquisition of several SVs.

In terms of hardware and software development, the next priority is to use a live matched filter against the known portion of the Starlink signal in the control loop, from which we expect a reduction in required dwell time to determine transmission presence. This, in turn, should raise the number of SV visits per FAI several times, with benefits for the exploration of the former two open questions introduced.

An Analysis of the Short-Term Time Stability of the Starlink Ku-Band Downlink Frame Clock

Introduction

Global navigation satellite systems (GNSS) are currently the most prevalent form of technology used for position, navigation, and timing (PNT). However, traditional GNSS techniques remain vulnerable to jamming and spoofing attacks that can leave users stranded without the ability to navigate, as emphasized by recent events. According to an International Air Transport Association report published in September 2023, the number of radiofrequency interference events soared from 10,843 such events in 2021 to 49,605 in 2022. In the effort to protect radionavigation, researchers have shown much interest in the augmentation of traditional GNSS with low Earth orbit (LEO) communications mega-constellations. As these constellations offer higher power and wider bandwidth, they also provide intrinsic resistance against adversarial interference. Further, the two-way, high-rate connectivity afforded by communications constellations enables desirable features such as user authentication and zero age-of-ephemeris.

Researchers interested in a free-to-use radionavigation receiver have investigated opportunistic approaches to PNT, i.e., PNT information extraction with no user-provider cooperation and limited a priori knowledge regarding the satellite's ephemeris and signal. SpaceX's Starlink constellation is of particular interest: it offers the widest signal availability, serving millions of subscribers worldwide with its 6,000 satellites. Opportunistic approaches using Starlink signals have already proven fruitful—researchers in several groups have independently demonstrated Doppler-based positioning with accuracies better than 10 m. While further accuracy improvements can be achieved with a pseudorange-based approach, the suitability of Starlink signals for precise pseudorange-based PNT remains an open question. The answer depends on detailed characteristics of the constellation's structure and timing aspects, such as its signal modulation, clock drift, and the nature of its timing adjustments.

Much information regarding the signal structure has been recently publicized in [42]. Some studies have shown simulated impacts of various clock types in LEO [57]. Others have developed methods for predicting LEO clock corrections, such as those developed for the GRACE mission [58], [59]. Nonetheless, little literature is yet available properly characterizing the accuracy and behavior of Starlink constellation clocks, essential details necessary for design of an opportunistic LEO PNT receiver.

In this paper, we leverage the insights offered by [42] as well as a dual Starlink and GPS L1 C/A capture system to conduct a detailed study of the Starlink timing properties and makes several contributions elucidating the nature of Starlink's timing characteristics. Note that [42] already details the signal model and capture system used, and thus the related descriptions are forgone. Interested readers are encouraged to refer to [42]. We separate timing characteristics into two main categories: relative and absolute frame timing. Relative frame timing refers to aspects of the Starlink frame clock that may be deduced without reliance on true GPS time (TGT) timestamps, including correction methods and intrinsic clock stability. In contrast, absolute frame timing explores aspects of the Starlink frame clock that depend on knowledge of absolute time, particularly any alignment of signal features with TGT. This paper focuses on the former, although follow-up analysis using similar methods as those described in the following sections are possible.

This paper begins by introducing the terminology and clock models used throughout the remainder of this paper. Then, it describes the methods and analyses conducted to make several observations regarding the Starlink frame clock's relative timing characteristics. It begins with a detailed presentation of the regular corrections applied to the frame clock, the frequency and intensity of which inform the clock predictability and a user's ability to make receiver-side clock corrections. Then, it presents a study of the Starlink frame clock's short-term stability characteristics, which can inform a potential Starlink PNT receiver's feasible degree of precision.

Fig. 14: Normalized cross-correlation of received Starlink data against a local PSS + SSS replica yields sharp peaks at the beginning of each frame whose primary lobe is approximately 20 ns wide for a 55-MHz-bandwidth capture.

SYSTEM MODELS AND CONCEPTS

This section presents a model for the structure and timing concept framework that will be used to discuss the methods and results presented in the following sections. Terminology useful for clear representation of the captured data is first introduced. Then, a theoretical Starlink clock model based on well-established clock models is presented to describe and understand the timing aspects of frame transmission and receipt.

Terminology

To better describe the beams and signals present when opportunistically extracting PNT information from a LEO mega-constellation, we illustrate common data capture complications, reiterate terms presented in previous publications, and introduce several new terms.

We often captured signals that simultaneously included transmissions from both assigned and side beams, as shown in Fig. 14. Such captured signals included as many as four coexisting signals and required careful disambiguation. As such, we introduce new terms describing the composition of a captured signal.

Composite signal: A captured signal composed of transmissions from two or more beams at a significant power level.

Simplex signal: A captured signal that contains the transmission from only a single beam at a significant power level.

Dominant signal: Within a composite signal, the strongest transmission originating from a single beam.

Secondary signal: Within a composite signal, any transmission originating from a single beam present, other than the dominant signal.

While we do not exclude the possibility that any given simplex signal may actually be a composite signal with quiet secondary signals, we posit that these supposed secondary signals typically retain precorrelation SNR values below -15 dB, too weak to be useful for PNT. As such, such signals are not considered. We typically attempt to measure the SNR of all dominant and secondary signals present via comparison against the captured data's noise floor, which is ideally taken as the complex signal variance over an interval in which neither dominant nor secondary signals are present. In practice, it is impossible to say with absolute certainty that a given interval used as the noise floor does not contain a weak frame from a side beam.

Clock Models

As with any analysis of PNT systems based on radio wave propagation, unambiguous and precise models of the various clocks involved are key to understanding and characterizing the system. Let t represent true time, or time according to an ideal clock, such as is closely realized by GPS system time [60]. Let t_f represent time according to a satellite's frame clock, or the clock that governs the timing of frames transmitted by the satellite. Finally, let t_r represent time according to the receiver clock. These time representations may be related by

$$\begin{aligned} t &= t_\mathrm{f}(t) - \delta t_\mathrm{f}(t) \\ t &= t_\mathrm{r}(t) - \delta t_\mathrm{r}(t_\mathrm{r}) \end{aligned}$$

The frame clock offset $\delta t_\mathrm{f}(t)$ and the receiver clock offset $\delta t_\mathrm{r}(t_\mathrm{r})$ are respectively the amount by which frame time and receiver time lag true time t. The receiver clock offset $\delta t_\mathrm{r}(t_\mathrm{r})$ is represented as a function of t_r because it is natively measured in receiver time in the course of solving for a position and time solution. The frame clock offset $\delta t_\mathrm{f}(t)$ is taken to be function of t because a model for it must be shared among users, which requires a common time base.

The time derivative of $\delta t_\mathrm{f}(t)$ and of $\delta t_\mathrm{r}(t_\mathrm{r})$, both with respect to t and denoted $\dot{\delta t}_\mathrm{f}(t)$ and $\dot{\delta t}_\mathrm{r}(t_\mathrm{r})$, are called fractional frequency deviation, written generically as $y(t)$, on which clock stability analysis is based [61, Chapter 9].

Fig. 15: Frame sequence timing diagram.

The time-domain signal $x(t)$ introduced in [42, Section III-A] models the sequence of frames transmitted by a given satellite under the assumption of an ideal clock. With the introduction of t_f, it may be expressed more realistically as $x(t_\mathrm{f})$. Fig. 20 offers further details about the frame clock. Each frame as transmitted has duration T_f according to the frame clock. Within each FAI, the frame slot index increments from 0 to $N_\mathrm{a} - 1$, with $N_\mathrm{a} = 11250$ being the number of frame slots in a FAI. Each FAI starts at the beginning of frame slot 0 and lasts $N_\mathrm{a}T_\mathrm{f} = 15$ seconds. The interval of unoccupied frame slots at the beginning of each FAI, called the FAI guard interval $T_\mathrm{ag}(l) = N_\mathrm{ag}(l)T_\mathrm{f}$, spans a variable number of frame slots $N_\mathrm{ag}(l) \in [16{,}26]$. Note that, for any FAI index l, frame slot $N_\mathrm{ag}(l)$ is occupied by definition, but other frame slots may not be occupied.

Let $t_\mathrm{f}(l, m)$ be the frame clock time at the instant when the frame in the mth frame slot of the lth FAI begins to pass through the phase center of the satellite's downlink antenna, where l and m are zero-based indices. This will be defined as

$$t_\mathrm{f}(l, m) \triangleq 15l + mT_\mathrm{f}$$

The quantity $\delta t_\mathrm{f}(l, m)$ is the corresponding clock offset and $t^*(l, m)$ is the corresponding true time, such that $t^*(l, m) = t_\mathrm{f}(l, m) - \delta t_\mathrm{f}(l, m)$. Another of this paper's key findings is that $\delta t_\mathrm{f}(l, 0) \approx 0$. Stated differently, a Starlink satellite's frame clock departure from true time at the beginning of each FAI is small—typically less than a few ms.

The quantity $t_\mathrm{r}(l, m)$ will be taken to indicate the time of reception, according to the receiver clock, of the frame that was transmitted at true time $t^*(l, m)$, with $\delta t_\mathrm{r}(l, m)$ being the corresponding receiver clock offset and $t_*(l, m)$ being the corresponding true time. More precisely, $t_\mathrm{r}(l, m)$ is the receiver clock time at which the frame transmitted at $t^*(l, m)$ from the satellite's downlink antenna's phase center first reached the receiver antenna's phase center. The receipt time $t_\mathrm{r}(l, m)$ can be related to $t_*(l, m)$, $t^*(l, m)$, and $t_\mathrm{f}(l, m)$ by

$$\begin{aligned} t_*(l, m) &= t_\mathrm{r}(l, m) - \delta t_\mathrm{r}(l, m) \\ t^*(l, m) &= t_\mathrm{r}(l, m) - \delta t_\mathrm{r}(l, m) - \delta t_\mathrm{tof} \\ t_\mathrm{f}(l, m) &= t_\mathrm{r}(l, m) - \delta t_\mathrm{r}(l, m) - \delta t_\mathrm{tof} + \delta t_\mathrm{f}(l, m) \end{aligned}$$

where δt_tof is the frame's true time of flight from transmission to reception.

RELATIVE FRAME TIMING ANALYSIS

This section focuses on relative frame timing analysis of the Starlink frame clock, investigating aspects of the Starlink frame clock discernible without consideration of the TGT timestamp. Restated, this section

concerns the timing of frames compared to one another, rather than to a commonly-agreed upon reference time such as the GPS epoch. This section first describes the methods used to determine relative timing of Starlink TOA signals with high precision. Then, it characterizes the nature of Starlink frame clock correction and presents statistical bounds on the correction size. Finally, it investigates the short-term clock stability, presenting and discussing two commonly observed patterns of clock stability degradations dubbed oscillations and excursions.

Exploration of Starlink frame clock adjustment and stability requires processing of the Starlink signal captured using the system described in [42]. We augmented the capture system with a common GPS-disciplined 10-MHz oven-controlled crystal oscillator such that parallel, simultaneous Starlink and GPS L1 C/A data captures could be made with guaranteed perfect sample parity. We begin by determining the per-frame frame clock offset $\delta t_{\mathrm{f}}(l,m)$ from the captured signal.

The captured Starlink signal is cross-correlated against a local replica consisting of the primary and secondary synchronization signals (PSS and SSS, respectively), resulting in correlation magnitude data as presented in Fig. 14. As these signals are present in all Starlink downlink frames and remain the same for all SVs at all times [42], this process can be applied to all captures using the same local replica.

Careful analysis of the correlation magnitude data, as presented in Fig. 14, yields high-precision TOA for each frame received. We take the dominant signals' correlation peaks in RRT as the TOAs of interest $t_{\mathrm{r}}(l,m)$. The parallel GPS L1 C/A capture is processed to estimate the RFSAs' receiver clock deviation $\delta t_{\mathrm{r}}(t_{\mathrm{r}})$ for the same duration of time. Note that Doppler estimates are also produced per received frame at this point via a one- dimensional correlation-maximizing search. The Starlink and GPS captures' simultaneity allows direct application of $\delta t_{\mathrm{r}}(t_{\mathrm{r}})$ to the Starlink capture, granting access first to $\delta t_{\mathrm{r}}(l,m)$, then to the frame TOA in TGT $t_{*}(l,m)$ with an accuracy of <1 ns via (15). This process also produces absolute TGT timestamps; however, these are not necessary for analysis of relative frame timing.

To access the frame clock deviation -$\delta t_{\mathrm{f}}(l,m)$, we manipulate the clock models presented earlier. Rearrangement of (17) with substitution of (15) and (14) gives

$$-\delta t_{\mathrm{f}}(l,m) = t_{*}(l,m) - (15l + mT_{\mathrm{f}}) - \delta t_{\mathrm{tof}}$$

To complete the right-hand side, we first reconstruct the ideal frame clock progression mT_{f} by counting the index frame occupancy index m, assuming m=0 at the beginning of the capture. This step introduces an offset error under the likely condition that the capture's first sample does not coincide perfectly with the beginning of an FAI. However, as this analysis primarily directs attention to relative frame timing, this offset error may be omitted as it is are shared by all frames in the same capture. The same reasoning also justifies omission of the per-FAI offset 15l.

As such, the only missing element from the right-hand side is the signal time of flight δt_{tof}. While it is possible to precisely determine δt_{tof} for every frame received, the methods involved are cumbersome and unnecessary for analysis of short-term clock stability and adjustments. Instead, we fit cubic polynomial models on actual TOA vs. ideal TOA data for each capture, approximating orbital effects on δt_{tof} over the course of an FAI.

The ideal TOA for each frame was taken in a two step process: First, a sequential count of the frames received was taken, such that each frame was assigned a number, dubbed the frame slot number. Note that not all available frame transmission times, or slots, were occupied by frame transmissions. Thus,

when no frame was found for a duration of $1/F_{\mathrm{f}}$ s, we assume that the SV chose to not send a frame for the slot and incremented the slot number by 1 such that sequential continuity was preserved. Then, for each frame TOA, its corresponding slot number was multiplied by $1/F_{\mathrm{f}}$ s to obtain its ideal TOA.

The order of fit was chosen to ensure trends up to the length of a single FAI length were eliminated. However, it should be noted that polynomial trend removal is cause-blind, and so polynomial trends in the frame clock deviation endemic to clock instabilities and adjustments are removed alongside those caused by orbital effects. Fortunately, the short-term instabilities of interest generally do not manifest in low-order polynomial trends over the course of a typical capture length (5-15 s). Nonetheless, it remains prudent to retain a low order of fit so as to avoid overfitting and removing clock variations of interest. Cubic polynomials consistently achieved both goals across all captures, and were thus chosen to fit δt_{tof}.

Application of this process per capture yielded important conclusions regarding the clock adjustments, short-term clock stability, and common clock degradation patterns that are discussed in the following sections.

Fig. 16: Top: Time history of the frame arrival time after removal of a 3rd-order polynomial fit to eliminate constant frequency and frequency rate errors and variations due to satellite orbital motion. Bottom: A punctured but otherwise continuous frame arrival time history can be constructed by fitting a low-order polynomial to each piecewise segment (excluding settling time; green traces), then vertically shifting the original-data segments to achieve a function (black trace) having underlying first-order continuity.

Fig. 17: As the top plot from Fig. 16 but for a Block v2.0-Mini SV.

Characterization of 1 Hz Frame Clock Adjustment

Application of the above process unveiled a clock deviation trace as shown in the top plot of Fig. 16. Analysis across all signal captures yielded several noteworthy characteristics of the frame clock adjustment.

First, the cadence of adjustment remained consistent at 1 Hz, with no clear dependence on the adjustment magnitude or direction. This is remarkably consistent—in only one case did we find that adjustment happened much earlier or later than expected, when the adjustment occurred at a 130 ms delay from the expected cadence. Further, the adjustments were extreme in magnitude to the point of near discontinuity. While some PNT clock correction methods gently coax the clock deviations back towards TGT such that receivers do not experience sudden phase trauma or loss of lock, the Starlink frame clock follows a rather forceful method. All clock corrections commit their changes and settle within a fraction of a nanosecond, even as adjustments frequently

shifted frame clock timing by >100 ns. As shown in Fig. 16, the correction committed at the 5.5 s mark alters the frame clock deviation by more than 200 ns with only a ∼0.15 s settling time. Further, it is worth noting that not every opportunity for clock adjustment was taken. For instance, in the top plot of Fig. 16, a clock adjustment at

the 6.5 s mark would both seem appropriate and match the apparent adjustment cadence. Nonetheless, the clock deviation trace appears smooth, indicating that no adjustment was committed. Altogether, the above adjustment characteristics give the short-term clock deviation plot a fractured appearance. We suspect that these adjustments would cause sizable phase trauma if ingested by a PNT receiver.

To create a smooth clock deviation time history, we note the size of each clock adjustment and shifted the 1 Hz segments vertically to eliminate the discontinuities associated with frame clock adjustment. This results in the black trace as presented in the bottom plot of Fig. 16. This data proves to be useful for the further evaluation of high-frequency clock stability.

Here, we note that frame clock adjustment characteristics of Block v2.0-Mini SVs behave somewhat differently from those of Block v1.0 and v1.5. Block v2.0-Mini SV clock bias adjustments are different from those of v1.0 and v1.5 SVs in two main ways: (1) their adjustment magnitudes are much smaller

overall, and (2) the adjustment intervals do not occur as consistently at the 1-Hz rate. This suggests that while the underlying adjustment algorithms may remain the same, software upgrades resulted in more consistent clock behavior. An example Block v2.0-Mini frame clock deviation trace is presented in Fig. 17.

Next, we evaluated the clock adjustment-induced discontinuities' statistical normality. Combining 281 clock adjustment data point from over 50 independent captures, we conducted a Shapiro-Wilk goodness-of-fit test against the Gaussian model and arrived at $p = 0.3$. This indicated a reasonably approximate match [62]. Further analysis showed the mean clock adjustment was 20 ns with a standard deviation of 117 ns, and the adjustments did not appear to be quantized.

Further investigation revealed mixed results regarding bias rate and clock rate adjustments. While the frame clock deviation data does show relatively small bias rate changes at the same 1 Hz intervals as bias adjustment, the adjustment generally correlates to neither the clock bias adjustment nor the required bias rate change for closer alignment with TGT, suggesting that any bias rate changes are unintentional. It should be noted that such bias rates changes could also be explained by low-quality clocks, but the true cause is unclear. We discovered small clock rate changes, as well. Doppler tracking revealed regular carrier phase rate changes on the same 1-Hz cadence as clock bias adjustments. These phase rate changes were inconsistent with orbital motion, indicating that the clock rate is purposefully adjusted.

Short-Term Frame Clock Stability Bound

To probe the stability limits of the v1.0 and v1.5 Starlink SV frame clocks, we performed an Allan deviation analysis of 18 smoothed frame time histories similar to the black trace in Fig. 16. The duration of these time histories ranged from 10 to 15 seconds. The data originate from 9 unique Block v1.0 and v1.5 Starlink SVs whose signals were captured during 2022 and 2023 and whose frame timing was derived from assigned beams, insofar as could be ensured, and manifested no anomalous excursions besides the 1-Hz clock correction discontinuities. Given the processing and data selection involved in creating the smoothed frame time histories on which the composite Allan deviation was based, which includes not only removal of orbital effects but also some low- frequency clock deviations, the Allan deviation analysis should be taken as a lower bound on the frame clock stability of Block v1.0 and v1.5 Starlink SVs. This Allan deviation analysis shows that the Starlink frame clock

possesses best-case behavior broadly consistent with a temperature-compensated crystal oscillator (TCXO). For example, at an averaging time of 1 second, the fractional frequency stability $\sigma_y(\tau) = 2.5 \times 10^{-9}$, which is what one would expect from an average-quality TCXO at $\tau = 1$ second. Thus, we may conclude that the frame clock stability of the Starlink v1.0 and v1.5 frame clocks is no better than a TXCO, though, as will be discussed below, it can episodically be much worse.

Fig. 18: Three single-set examples of high-frequency frame clock irregularities (top) and histogram of frame clock irregularity RMS values (bottom). First order trends have been removed per each 1-second interval between clock corrections in both the single-set examples and before RMS calculation.

Other Anomalous Short-Term Frame Clock Behaviors

While the broad stability bound analysis addresses the best-case behavior of the Starlink frame clock, case-by- case analysis shows episodic high-frequency frame clock anomalies that would severely degrade opportunistic pseudorange-based PNT solutions formed from frame clock measurements. Fig. 18 illustrates several examples of these anomalies as the high-frequency components of the clock deviation trace. To set a baseline for comparison, the first subplot of Fig. 18 shows the high-frequency component of a frame clock deviation trace under nominal behavior. The clock deviation remains consistently tight between ±5 ns over the course of a whole FAI and possesses an RMS value of 1.8 ns. The next two subplots show anomalous departures from the nominal behavior. The last subplot aggregates the RMS values from 50 high-deviation clock traces as a histogram.

To derive the high-frequency clock stability traces presented in Fig. 18, we begin with TOAs extracted using the dual-capture system and cross-correlation against known sequences, as previously discussed. To isolate high- frequency variations in clock deviation, we conduct two stages of polynomial trend removals. First, with the clock bias adjustment data now available, we first vertically shift the original TOA data and then remove the cubic polynomial trend. This achieves a fuller elimination of cubic polynomial trends compared to the black trace presented in the bottom plot of Fig. 16, which was produced with the discontinuous clock bias adjustments still present during cubic polynomial trend removal. Second, a linear polynomial was removed from each 1 s between-adjustment piecewise segment to further eliminate low-frequency trends. This two-stage trend removal resulted in high-frequency frame clock irregularity data, albeit with slight disjoints occurring where TOA data was removed due to clock adjustment events.

Remarkably, a majority of the captures we took exhibited consistently low amounts of high-frequency instability at RMS values of 2.5 ns or lower. Importantly, the Starlink frame clock operating with nominal stability could support precise PNT systems. Unfortunately, the stability would oftentimes suddenly and unpredictably degrade, with instability RMS values reaching up to nearly ten times those experienced under nominal stability. We identified two anomalistic categories: oscillatory and excursive.

1) *Degradation due to Frame Clock Oscillations*: When investigating high-frequency clock stability data that exhibited high RMS values, we noticed that in many cases the frame clock instabilities were self-correlated in time. In particular, high-deviation, regular-interval variations were common. One such example can be seen in the second subplot of Fig. 18. In contrast with the nominal behavior shown in the first subplot, the clock errors repeatedly vary between -13 ns and 28 ns at a frequency hovering around 13.4 Hz. We dubb these frame clock oscillations. Oscillations occurred much more frequently than excursions, which will be discussed later. Block v2.0-Mini SVs also tended to exhibit larger high-frequency clock oscillations than Block v1.0 or v1.5 SVs.
2) *Degradation due to Frame Clock Excursions*: Different from oscillations, frame clock excursions are drastic, momentary outliers in frame clock deviation despite otherwise nominal stability.

These deviations, unlike the oscillations presented in the previous section, tend to be short-lived and irregularly occurring. One such example is shown in the third subplot of Fig. 18, where the frame clock appears to be behaving nominally except for sudden bursts of clock delay, e.g., at the 2.6 s and 5.4 s marks. These excursive behaviors tend to manifest less frequently than oscillations; thus, it is difficult to show that weak or composite signals always yield clock excursions, or vice versa. However, it can be said that these behaviors may exist in simplex signals, according to the previously mentioned capture taken from STARLINK-3894.

Lastly, one of the most interesting observations we made was that for the same SV, frame clock stability behavior frequently switched between nominal and degraded modes between FAIs with little to no transient behavior. This suggests that the degraded frame clock performance should not be blamed on hardware deficiencies. If it were, one would expect consistent behavior between FAI transitions. Instead, the hardware seems to be perfectly capable of producing nominal behavior, but software configuration issues might be to blame. If in fact the degraded performance can be blamed on software issues, said issues could be resolved via over-the-air software updates.

Nonetheless, both degraded stability modes are cause for concern—if ingested without warning, oscillations and excursions would introduce significant phase trauma in a PNT receiver, likely causing loss of phase lock and degraded PNT accuracy.

Conclusion

We have developed and used a dual-capture system and accompanying processing technique that simultaneously captures both Starlink and GPS L1 C/A signals, providing high-precision estimation of the signal time of receipt. With the data provided, we conduct a relative frame timing analysis of the Starlink frame clock and provide characterizations of its clock adjustment and short-term stability. Further, we identify two patterns of clock instability of concern for developers of a Starlink-based PNT receiver.

References

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

[2] Jafarnia-Jahromi, A., Broumandan, A., Nielsen, J., and Lachapelle, G., "Review article: GPS vulnerability to spoofing threats and review of antispoofing techniques," International Journal of Navigation and Observation, 2012, pp. 1–16.

[3] Humphreys, T. E., "Statement on the vulnerability of civil unmanned aerial vehicles and other systems to civil GPS spoofing," United States House of Representatives Committee on Homeland Security: Subcommittee on Oversight, Investigations, and Management, July 2012.

[4] Scott, L., "Anti-spoofing and authenticated signal architectures for civil navigation systems," Proceedings of the ION GNSS Meeting, 2003, pp. 1542–1552.

[5] Humphreys, T. E., Ledvina, B. M., Psiaki, M. L., O'Hanlon, B. W., and Kintner, Jr., P. M., "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," Proceedings of the ION GNSS Meeting, Institute of Navigation, Savannah, GA, 2008.

[6] 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.

[7] Gross, J. N., Kilic, C., and Humphreys, T. E., "Maximum-likelihood power-distortion monitoring for GNSS-signal authentication,"IEEE Transactions on Aerospace and Electronic Systems, Vol. 55, No. 1, 2018, pp. 469–475.

[8] Psiaki, M. L., O'Hanlon, B. W., Powell, S. P., Bhatti, J. A., Wesson, K. D., Humphreys, T. E., and Schofield, A., "GNSS Spoofing Detection using Two-Antenna Differential Carrier Phase," Proceedings of the ION GNSS+ Meeting, Institute of Navigation, Tampa, FL, 2014.

[9] O'Hanlon, B., Psiaki, M., Bhatti, J., and Humphreys, T., "Real-Time Spoofing Detection Using Correlation Between two Civil GPS Receiver," Proceedings of the ION GNSS Meeting, Institute of Navigation, Nashville, Tennessee, 2012.

[10] Gross, J. and Humphreys, T. E., "GNSS Spoofing, Jamming, and Multipath Interference Classification using a Maximum-Likelihood Multi-Tap Multipath Estimator," Proceedings of the ION International Technical Meeting, Jan. 2017.

[11] O'Hanlon, B., Bhatti, J., Humphreys, T. E., and Psiaki, M., "Real-Time Spoofing Detection in a Narrow-Band Civil GPS Receiver," Proceedings of the ION GNSS Meeting, Institute of Navigation, Portland, Oregon, 2010.

[12] Clements, Z., Yoder, J. E., and Humphreys, T. E., "Carrier-phase and IMU based GNSS Spoofing Detection for Ground Vehicles,"Proceedings of the ION International Technical Meeting, Long Beach, CA, 2022, pp. 83–95.

[13] Clements, Z., Yoder, J. E., and Humphreys, T. E., “GNSS Spoofing Detection: An Approach for Ground Vehicles Using Carrier-Phase and Inertial Measurement Data,” GPS World, Vol. 34, No. 2, 2023, pp. 36–41.

[14] Tanil, C., Khanafseh, S., Joerger, M., and Pervan, B., “An INS Monitor to Detect GNSS Spoofers Capable of Tracking Vehicle Position,” taes, Vol. 54, No. 1, Feb. 2018, pp. 131–143.

[15] Jafarnia-Jahromi, A., Daneshmand, S., Broumandan, A., Nielsen, J., and Lachapelle, G., “PVT solution authentication based on monitoring the clock state for a moving GNSS receiver,” European navigation conference (ENC), Vol. 11, 2013.

[16] Humphreys, T. E., “Detection Strategy for Cryptographic GNSS Anti-Spoofing,” IEEE Transactions on Aerospace and Electronic Systems, Vol. 49, No. 2, 2013, pp. 1073–1090.

[17] Fernandez-Hernandez, I., Winkel, J., O’Driscoll, C., Cancela, S., Terris-Gallego, R., Lo´pez-Salcedo, J. A., Seco-Granados, G., Dalla Chiara, A., Sarto, C., Blonski, D., et al., “Semi-Assisted Signal Authentication for Galileo: Proof of Concept and Results,” IEEE Transactions on Aerospace and Electronic Systems, 2023.

[18] Streufert, D., “ADS-B Exchange: World’s largest source of unfiltered flight data,” https://globe.adsbexchange.com/?r.

[19] LaChapelle, D. M., Narula, L., and Humphreys, T. E., “Orbital War Driving: Assessing Transient GPS Interference from LEO,”Proceedings of the ION GNSS+ Meeting, St. Louis, MO, 2021.

[20] 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, Vol. 68, No. 4, 2021, pp. 673–685.

[21] Clements, Z., Ellis, P., Psiaki, M. L., and Humphreys, T. E., “Geolocation of Terrestrial GNSS Spoofing Signals from Low Earth Orbit,” Proceedings of the ION GNSS+ Meeting, Denver, CO, 2022, pp. 3418–3431.

[22] Clements, Z., Ellis, P., and Humphreys, T. E., “Dual-Satellite Geolocation of Terrestrial GNSS Jammers from Low Earth Orbit,”Proceedings of the IEEE/ION PLANS Meeting, Monterey, CA, 2023, pp. 458–469.

[23] Clements, Z., Ellis, P., and Humphreys, T. E., “Pinpointing GNSS Interference from Low Earth Orbit,” Inside GNSS, Vol. 18, No. 5, 2023, pp. 42–55.

[24] McKibben, A., McKnight, R., Peters, B. C., Arnett, Z., and Ugazio, S., “Interference Effects on a Multi-GNSS Receiver On-Board a CubeSat in LEO,” Proceedings of the ION GNSS+ Meeting, Denver, CO, 2023, pp. 1245–1258.

[25] Sidi, A. and Weiss, A., “Delay and Doppler Induced Direct Tracking by Particle Filter,” IEEE Transactions on Aerospace and Electronic Systems, Vol. 50, No. 1, January 2014, pp. 559–572.

[26] Musicki, D., Kaune, R., and Koch, W., “Mobile Emitter Geolocation and Tracking Using TDOA and FDOA Measurements,” IEEE Transactions on Signal Processing, Vol. 58, No. 3, March 2010, pp. 1863–1874.

[27] Ho, K. and Chan, Y., “Geolocation of a known altitude object from TDOA and FDOA measurements,” IEEE Transactions on Aerospace and Electronic Systems, Vol. 33, No. 3, July 1997, pp. 770–783.

[28] Ellis, P., Rheeden, D. V., and Dowla, F., “Use of Doppler and Doppler Rate for RF Geolocation Using a Single LEO Satellite,”IEEE Access, Vol. 8, 2020, pp. 12907–12920.

[29] Ellis, P. and Dowla, F., “Performance bounds of a single LEO satellite providing geolocation of an RF emitter,” 2018 9th Advanced Satellite Multimedia Systems Conference and the 15th Signal Processing for Space Communications Workshop (ASMS/SPSC), IEEE, 2018, pp. 1–5.

[30] Ellis, P. B. and Dowla, F., “Single Satellite Emitter Geolocation in the Presence of Oscillator and Ephemeris Errors,” 2020 IEEE Aerospace Conference, IEEE, 2020, pp. 1–7.

[31] Chen, X., Morton, Y., Yu, W.-X., and Truong, T.-K., “GNSS Spoofer Localization with Counterfeit Clock Bias Observables on a Mobile Platform,” IEEE Sensors Journal, 2023.

[32] Clements, Z., Iannucci, P. A., Humphreys, T. E., and Pany, T., “Optimized Bit-Packing for Bit-Wise Software-Defined GNSS Radio,”Proceedings of the ION GNSS+ Meeting, St. Louis, MO, 2021, pp. 3749–3771.

[33] Nichols, H. A., Murrian, M. J., and Humphreys, T. E., “Software-Defined GNSS is Ready for Launch,” Proceedings of the ION GNSS+ Meeting, Denver, CO, 2022.

[34] Clements, Z., Goodridge, I., Ellis, P., Murrian, M. J., and Humphreys, T. E., “Demonstration of Single-Satellite GNSS Spoofer Geolocation,” Proceedings of the ION International Technical Meeting, Long Beach, CA, 2024, pp. 361–373.

[35] Gu¨nther, C., “A Survey of Spoofing and Counter-Measures,” NAVIGATION, Vol. 61, No. 3, 2014, pp. 159–177.

[36] Teunissen, P. J. and Montenbruck, O., editors, Springer handbook of global navigation satellite systems, Springer, 2017.

[37] Humphreys, T. E., “Interference,” Springer Handbook of Global Navigation Satellite Systems, Springer International Publishing, 2017, pp. 469–503.

[38] Dolman, E. C., “New Frontiers, Old Realities,” Strategic Studies Quarterly, Vol. 6, No. 1, 2012, pp. 78–96.

[39] Iannucci, P. A. and Humphreys, T. E., “Fused Low-Earth-Orbit GNSS,” IEEE Transactions on Aerospace and Electronic Systems, 2022, pp. 1–1.

[40] Neinavaie, M., Khalife, J., and Kassas, Z. M., “Acquisition, Doppler Tracking, and Positioning With Starlink LEO Satellites: First Results,” IEEE Transactions on Aerospace and Electronic Systems, Vol. 58, No. 3, 2022, pp. 2606–2610.

[41] Psiaki, M. L., “Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids,” NAVIGATION, Vol. 68, No. 3, 2021, pp. 621–641.

[42] Humphreys, T. E., Iannucci, P. A., Komodromos, Z. M., and Graff, A. M., “Signal Structure of the Starlink Ku-Band Downlink,” IEEE Transactions on Aerospace and Electronic Systems, 2023, pp. 1–16.

[43] Teng, Y. and Wang, J., “A closed-form formula to calculate geometric dilution of precision (GDOP) for multi-GNSS constellations,” GPS Solutions, Vol. 20, No. 3, 2016, pp. 331–339.

[44] Nardin, A., Dovis, F., and Fraire, J. A., “Empowering the Tracking Performance of LEO PNT by Means of Meta-Signals,” 2020 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE), IEEE, 2020, pp. 153–158.

[45] Komodromos, Z. M., Qin, W., and Humphreys, T. E., “Signal Simulator for Starlink Ku-Band Downlink,” Proceedings of the ION GNSS+ Meeting, 2023, pp. 2798–2812.

[46] Bla´zquez-Garc´ıa, R., Cristallini, D., Ummenhofer, M., Seidel, V., Heckenbach, J., and O’Hagan, D., “Capabilities and challenges of passive radar systems based on broadband low-Earth orbit communication satellites,” IET Radar, Sonar & Navigation, Vol. n/a, No. n/a, 2023.

[47] Yang, C. and Soloviev, A., “Starlink Doppler and Doppler Rate Estimation via Coherent Combining of Multiple Tones for Opportunistic Positioning,” 2023 IEEE/ION Position, Location and Navigation Symposium (PLANS), IEEE, 2023, pp. 1143–1153.

[48] SpaceX, “SpaceX Non-Geostationary Satellite System, Technical Parameters,” https://licensing.fcc.gov/myibfs/download.do? attachment key=1877844, Aug. 2019, SAT-MOD-20190830-00087.

[49] Robotics, T., “WidowX Dual XM430 Pan & Tilt,” https://www.trossenrobotics.com/widowx-x-series-dual-servo-robot-turret.aspx, September 2023.

[50] Neinavaie, M., Khalife, J., and Kassas, Z. M., “Exploiting Starlink Signals for Navigation: First Results,” Proceedings of the ION GNSS+ Meeting, St. Louis, Missouri, Sept. 2021, pp. 2766–2773.

[51] NSI-MI, “Standard Gain Horns,” https://www.nsi-mi.com/products/antenna-products/standard-gain-horns, September 2023.

[52] EverythingRF, “1000XDF,” https://www.everythingrf.com/products/low-noise-blocks/norsat/698-512-1000xdf, September 2023.

[53] Vallado, D., Crawford, P., Hujsak, R., and Kelso, T., “Revisiting spacetrack report # 3,” AIAA/AAS Astrodynamics Specialist Conference and Exhibit, 2006, p. 6753.

[54] Robotis, “Dynamixel XM430-W350-R,” https://www.robotis.us/dynamixel-xm430-w350-r/, September 2023.

[55] Maybell, M. and Simon, P., “Pyramidal horn gain calculation with improved accuracy,” IEEE Transactions on Antennas and Propagation, Vol. 41, No. 7, 1993, pp. 884–889.

[56] Jull, E., “Errors in the predicted gain of pyramidal horns,” IEEE Transactions on Antennas and Propagation, Vol. 21, No. 1, 1973, pp. 25–31.

[57] Wang, K. and El-Mowafy, A., “LEO satellite clock analysis and prediction for positioning applications,” Geo-spatial Information Science, Vol. 25, No. 1, 2022, pp. 14–33.

[58] Ge, H., Wu, T., and Li, B., “Characteristics analysis and prediction of Low Earth Orbit (LEO) satellite clock corrections by using least-squares harmonic estimation,” GPS Solutions, Vol. 27, No. 1, 2023, pp. 38.

[59] Jiang, C., Luo, Z., Zhu, M., Guan, M., Zhu, H., and Gao, M., “Characteristics Analysis of Leo Satellite Clock Corrections and Prediction by Sliding Estimation,” Available at SSRN 4818891.

[60] Jekeli, C. and Montenbruck, O., Springer Handbook of Global Navigation Satellite Systems, chap. Time and Reference Systems, Springer, 2017, pp. 25–58.

[61] Thompson, A., Moran, J., and Swenson, G., Interferometry and Synthesis in Radio Astronomy, Wiley, 2001.

[62] Royston, P., “Approximating the Shapiro-Wilk W-test for non-normality,” Statistics and computing, Vol. 2, 1992, pp. 117–119.