

Technical Report Documentation Page

1. Report No.	2. Government Accession No. N/A	3. Recipient's Catalog No. N/A
4. Title and Subtitle PNT Protection through Interference Monitoring and Robust Backups to GNSS		5. Report Date January 30, 2026
		6. Performing Organization Code N/A
7. Author(s) Todd E. Humphreys https://orcid.org/0000-0003-0749-6988 Zacharias M. Komodromos Samuel C. Morgan Zachary L. Clements Wenkai Qin W. Jeremy Morrison		8. Performing Organization Report No. N/A
		10. Work Unit No. (TRAIS) N/A
9. Performing Organization Name and Address The University of Texas at Austin Austin, Texas, 78712		11. Contract or Grant No. 69A3552348327
		13. Type of Report and Period Covered Final (Aug '24 to Dec '25)
12. Sponsoring Agency Name and Address U.S. Department of Transportation (SPR)		14. Sponsoring Agency Code N/A
15. Supplementary Notes Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.		
16. Abstract This project report on PNT protection is best viewed as part of a comprehensive navigation and timing solution concept called deep-layered navigation (DLN) in which synergistic but independent systems are layered to increase PNT accuracy and robustness. DLN is the navigation analog of the “defense in depth” concept in information security, where multiple layers of security controls and checkpoints are emplaced throughout a system such that even when some layers are breached, security is maintained. The two primary thrusts of our DLN approach, reported here, are (1) fusion of traditional GNSS with signals from low-Earth orbit (LEO) to strengthen PNT against interference, and (2) monitoring GNSS interference from LEO to geolocate its sources.		

17. Key Words Resilient positioning, navigation, and timing. Protection of radio bands designated RNSS. Augmentations to GNSS for PNT.		18. Distribution Statement No restrictions. This document is available to the public through NTIS: National Technical Information Service Springfield, Virginia 22161	
19. Security Classif.(of this report) Unclassified	20. Security Classif.(of this page) Unclassified	21. No. of Pages 103 pages	22. Price N/A

Form DOT F 1700.7 (8-72)

Reproduction of completed page authorized



CARMEN+

Center for Automated Vehicles Research
with Multimodal Assured Navigation

USDOT University Transportation Centers Program



Final Report: PNT Protection through Interference Monitoring and Robust Backups to GNSS

P.I.	Project Info:
Todd E. Humphreys	Grant No. 69A3552348327
The University of Texas at Austin	DUNS: 832127323
Aerospace Engineering	EIN #: 31-6025986
	Project Effective: August 1, 2024 Project End: December 31, 2025 Submission: January 30, 2026

Consortium Members:



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's work on PNT protection is best viewed as part of a comprehensive navigation and timing solution concept called deep-layered navigation (DLN) in which synergistic but independent systems are layered to increase PNT accuracy and robustness. DLN is the navigation analog of the "defense in depth" concept in information security, where multiple layers of security controls and checkpoints are emplaced throughout a system such that even when some layers are breached, security is maintained. The two primary thrusts of our DLN approach, reported here, are (1) fusion of traditional GNSS with signals from low-Earth orbit (LEO) to strengthen PNT against interference, and (2) monitoring GNSS interference from LEO to geolocate its sources.

Acknowledgements

Research was supported by the U.S. Department of Transportation under Grant 69A3552348327 for the CARMEN+ University Transportation Center. Cost sharing was supported by affiliates of the 6G@UT center within the Wireless Networking and Communications Group at The University of Texas at Austin. In particular, several components of the work reported herein were supported by Amazon's Project Kuiper as an affiliate of the 6G@UT center within the Wireless Networking and Communications Group at The University of Texas at Austin.

Executive Summary

Central to civil infrastructure (e.g., transportation, communications) is the problem of robust and secure PNT. This is especially true of urban environments, which offer more challenges, such as interference and multipath, when compared to open-sky conditions. As the only positioning system that currently offers absolutely referenced meter-level accuracy with global coverage, GNSS will no doubt play a significant role in civil infrastructure going forward. For example, if strengthened against jamming and spoofing, carrier-phase-differential GNSS (CDGNSS), coupled with low-cost inertial sensing, will be nearly sufficient for position, velocity, and timing (PVT) needs of urban air mobility (UAM). But nearly sufficient is insufficient: it is not enough for a UAM PVT solution to offer decimeter-accurate positioning with 99% availability, or even 99.9% availability. UAM will demand that its navigation systems offer dm-accurate positioning with integrity risk on the order of $10e-7$ for a meter-level alert limit and availability with several more 9s than 99.9%

This report's work in PNT protection is best viewed as part of a comprehensive navigation and timing solution concept called deep-layered navigation (DLN) in which synergistic but independent navigation systems are layered to increase accuracy and robustness. DLN is the navigation analog of the "defense in depth" concept in information security, where multiple layers of security controls and checkpoints are emplaced throughout a system such that even when some layers are breached, security is maintained. Likewise, in the safety-of-life UAM navigation context, multiple layers of navigation and timing systems, all interoperable and mutually reinforcing but substantially independent, are an essential defense against the whims of Mother Nature and the foibles of human nature.

At DLN's core sits redundant inertial navigation, which is virtually impervious to radio frequency (RF) interference, poor weather, signal blockage, and data ambiguity. The outermost layer---the default navigation system and first line of defense---is a specialized variant of inertially-aided CDGNSS, recently developed in the UT Radionavigation Laboratory that has been substantially secured against spoofing and substantially hardened against the multipath and signal blockage conditions of the urban ground vehicle environment, which can be considered a worst-case realization of the urban air vehicle environment. But despite its coupling with inertial sensing, the technique developed in the UT RNL cannot tolerate extended GNSS outages. A secondary source of absolute PVT is required to bound the growth of position errors.

Fusion of GNSS and signals of opportunity (SoPs) has been explored by the CARMEN+ team with promising results. For a safety-of-life application like UAM, SoPs will be viewed less favorably than a dedicated system such as GNSS. But SoP availability, ample bandwidth, and strong received power nonetheless make them an appealing component of any deep-layered navigation and timing strategy. Compared to terrestrially-sources SoPs,

SoPs from LEO mega-constellations such as SpaceX's Starlink, Amazon's Kuiper, and OneWeb enjoy the advantage of vastly broader---indeed, nearly global---coverage. The downlink signals from these constellations are also carried on higher frequencies (Ku-band for Starlink and OneWeb; Ka-band for Kuiper), whose short wavelengths permit use of small phased-array receiver antennas capable of forming narrow (few degrees) beams toward their serving satellites. Narrow beamforming makes such systems naturally resilient to radio interference.

As reported herein, we have developed novel signal capture, processing, and estimation techniques that enable extraction of navigation and timing measurements from broadband mega-LEO signals of opportunity. We envision broadband LEO-based PNT as a novel and useful component in a comprehensive deep-layered navigation and timing system for civil transportation and infrastructure.

A second thrust of the work reported here concerns GNSS interference detection and geolocation. 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 demonstrated techniques whereby GNSS interference signals captured by a single satellite can be used to characterize and geolocate the signals' source.

Project Outputs Achieved:

1. A technique for locating sources of GNSS spoofing from a single LEO SV.
2. An analysis of the timing properties of the Starlink LEO mega constellation's downlink signals to assess their suitability for pseudorange-based PNT.
3. Analysis of the signal structure of the OneWeb LEO mega constellation downlink signals to support opportunistic LEO pseudorange-based PNT.
4. A technique and demonstration of GNSS-free entry to a LEO mega-constellation network to support both communications and PNT.
5. Conference papers and archival journal articles reporting on each of these outputs.

Table of Contents

This report is organized into seven sections, each one representing a separate paper that was published during the course of the project. The sections are organized according to the order of the **Project Outputs Achieved** listed in the foregoing **Executive Summary**. References are bound to each individual section of the report.

Section 1: Z. L. Clements, P. B. Ellis, M. J. Murrian, M. L. Psiaki, and T. E. Humphreys, “Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit,” NAVIGATION, 2025. To be published.

Section 2: W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, “Timing properties of the Starlink Ku-band downlink,” IEEE Transactions on Aerospace and Electronic Systems, pp. 1–19, 2025.

Section 3: Z. M. Komodromos and T. E. Humphreys, “Signal parameter estimation and demodulation of the OneWeb Ku-Band downlink,” npj Wireless Technology, 2025. To be published.

Section 4: S. C. Morgan, Z. M. Komodromos, W. Qin, Z. L. Clements, A. M. Graff, W. J. Morrison, and T. E. Humphreys, “A mock implementation of fused LEO GNSS,” in Proceedings of the IEEE/ION PLANS Meeting, (Salt Lake City, UT), 2025. **Overall best student paper award.**

Section 5: Z. M. Komodromos, S. C. Morgan, Z. L. Clements, W. Qin, W. J. Morrison, and T. E. Humphreys, “Network-aided pseudorange-based LEO PNT from OneWeb,” in Proceedings of the IEEE/ION PLANS Meeting, (Salt Lake City, UT), 2025.

Section 6: W. Qin, Z. M. Komodromos, S. C. Morgan, and T. E. Humphreys, “Maximum likelihood time of arrival and Doppler estimation for precise Starlink-based PNT,” in Proceedings of the IEEE/ION PLANS Meeting, (Salt Lake City, UT), 2025.

Section 7: S. Morgan and T. E. Humphreys, “HOOC-EM: Fast beam sweeping for LEO mega-constellation customer terminals,” in Proceedings of the ION GNSS+ Meeting, 2024.

Single-Satellite-Based Geolocation of Broadcast GNSS Spoofers from Low Earth Orbit

Zachary L. Clements*, Patrick B. Ellis[†], Iain Goodridge[‡], Matthew J. Murrian[†], Mark L. Psiaki[‡], and Todd E. Humphreys*

*Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin

[†]Spire Global

[‡]Department of Aerospace and Ocean Engineering, Virginia Tech

Abstract—This paper presents an analysis and experimental demonstration of single-satellite single-pass geolocation of a terrestrial broadcast Global Navigation Satellite System (GNSS) spoofer from Low Earth Orbit (LEO). The proliferation of LEO-based GNSS receivers offers the prospect of 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 for the emitter. This paper presents a technique crafted specifically for indiscriminate broadcast-type GNSS spoofing signals. Furthermore, it explores how unmodeled oscillator instability and worst-case spoofer-introduced signal variations degrade the geolocation estimate. The proposed geolocation technique is validated by a controlled experiment, in partnership with Spire Global, in which a LEO-based receiver captures broadcast GNSS spoofing signals transmitted from a known ground station on a non-GNSS frequency band.

Index Terms—GNSS spoofing; emitter geolocation; interference localization; spectrum monitoring.

I. 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. A GNSS spoofer emits an ensemble of false GNSS signals intending that the victim receiver(s) accept them as authentic GNSS signals, thereby inferring a false position fix and/or clock offset [1], [2]. A successful spoofing attack may lead to serious consequences.

The academic community has long warned the public about the threat of GNSS spoofing [3]–[5]. Within the past decade, significant progress has been made in GNSS spoofing detection and mitigation [1], [2], [6]–[8]. 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 [9]–[14]. Consistency checks between the estimated signal and onboard inertial sensors can provide quick and reliable spoofing detection [15]–[18]. Monitoring the clock state can also be used to detect spoofing [19]–[21]. Cryptographic authentication techniques are currently being developed and implemented to verify received signals [22]–[27].

Although the recent advances in GNSS spoofing detection have been inspiring, many older GNSS receivers in current operation are unable to incorporate such defenses, leaving them vulnerable to attacks. For example, the civilian maritime and airline industries are encountering GNSS jamming and

spoofing at an alarming rate [28]–[34]. Anomalous positioning information broadcast by ships in Automatic Identification System (AIS) messages, and airplanes in Automatic Dependent Surveillance-Broadcast (ADS-B) messages, indicate recent widespread jamming and spoofing. These civilian aircraft and ships ensnared by GNSS spoofing are likely unintended targets caught in the electronic warfare crossfire near ongoing conflict zones.

GNSS spoofing attacks can be sorted into two categories, targeted spoofing and broadcast spoofing. In targeted spoofing, an attacker transmits spoofing signals for a specific (possibly-moving) target it wishes to deceive. This type of attack involves the attacker tailoring a spoofing trajectory for its specific target, causing a gradual pull-off from the victim’s true trajectory, and compensating for the relative motion between the spoofer and the target to minimize target’s probability of detection [35]. Targeted spoofing is a sophisticated, expensive, and difficult-to-detect attack that requires the attacker to have the ability to precisely track the target and craft spoofing signals in accordance with the target’s motion, all in real-time. Due to its complexity and narrow scope, this form of spoofing is the least common. Other GNSS receivers besides the targeted victim can also be captured by these signals, but a non-targeted receiver can more easily detect such spoofing. Moreover, targeted spoofing may involve narrow beamforming, making reception by non-target receivers unlikely.

Broadcast spoofing is less expensive, less complex, and wider in geographic extent than targeted spoofing, and thus more common. In broadcast spoofing, an attacker transmits spoofing signals broadly with the intent to deceive all GNSS receivers within a wide area. Because broadcast spoofing is non-targeted, victim GNSS receivers typically see a sudden jump in position and/or timing, which is trivial to detect with basic spoofing detection checks. Yet despite being easy to detect, broadcast spoofing remains effective at denying GNSS access to victims lacking proper defenses. When a GNSS receiver cannot confidently differentiate between authentic and spoofing signals, it is rendered useless—or worse: hazardously misleading. The spoofers recently affecting the aviation and maritime industries appear to be of the broadcast type.

Given that many currently deployed GNSS receivers are unable to defend themselves even against easy-to-detect broadcast spoofing, GNSS users need to be warned of hazardous GNSS-challenged environments. The proliferation of LEO-based GNSS receivers provides the potential of unprecedented

spectrum awareness, enabling GNSS interference detection, classification, and geolocation with worldwide coverage [36]–[41]. Existing and proposed LEO constellations provide worldwide coverage with frequent revisit rates, allowing for an always-updating operating picture, a noted shortfall in current capabilities [42]. 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 [39], [40], [43]–[45]. Multiple time-synchronized receivers can exploit time- and frequency-difference-of-arrival (T/FDOA) measurements to estimate the emitter location. The authors of the current paper were able to geolocate over 30 GNSS interference sources across the Near East from a dual-satellite time-synchronized capture [39], [40]. However, planning simultaneous multi-satellite captures to enable T/FDOA-based geolocation can be difficult to coordinate and expensive, whereas single-satellite collects are straightforward and less costly. Accordingly, this paper focuses on single-satellite geolocation.

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. But if a signal’s carrier can be tracked, or Doppler can be otherwise measured, then accurate single-satellite-based emitter geolocation is possible from Doppler measurements alone, provided that the emitter’s carrier frequency is quasi-constant [37], [46]–[48]. But 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 specially, as they do not transmit at a constant carrier frequency: they add an unknown time-varying frequency component to each spoofing signal, imitating the range-rate between the corresponding spoofed GNSS satellite and the counterfeit spoofed location [35]. A key contribution of the current paper 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. A single-receiver spoofer geolocation technique based on counterfeit clock observables is also presented in [49]. However, [49] only considers the spoofed pseudorange measurements and depends on a stationary receiver initialization period, which is not possible in LEO.

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

This paper makes four primary contributions. First, it presents a single-satellite, single-pass GNSS spoofer geolocation technique that extracts a range-rate between a LEO-based receiver and a terrestrial broadcast spoofer from captured raw samples. Second, it offers an experimental demonstration of the technique with a truth solution. Third, it derives an analytic expression for how transmitter clock instability degrades the single-satellite geolocation solution. Fourth, it investigates the geolocation positioning errors as a function of worst-case spoofed clock behavior.

Preliminary conference versions of this paper were published in [50], [51]. The current version significantly extends these with contributions three and four mentioned above.

II. SIGNAL MODELS

A. 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 $\tilde{\mathbf{x}} = [\mathbf{r}_{\tilde{\tau}}^T, \delta t_{\tilde{\tau}}, \mathbf{v}_{\tilde{\tau}}^T, \delta \dot{t}_{\tilde{\tau}}]^T$, where $\mathbf{r}_{\tilde{\tau}}$ is the spoofed position in Earth-centered-Earth-fixed (ECEF) coordinates, $\delta t_{\tilde{\tau}}$ is the spoofed clock bias increment, $\mathbf{v}_{\tilde{\tau}}$ is the spoofed velocity, and $\delta \dot{t}_{\tilde{\tau}}$ is the spoofed clock drift increment. 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 now presented. The ensemble of spoofing signals transmitted by the spoofer, denoted

$$\mathbf{x}(t) = \sum_{n=1}^N \begin{pmatrix} s_n(t) \end{pmatrix} \quad (1)$$

contains N spoofing signals, where the n th spoofing signal is denoted $s_n(t)$ for $n = 1, 2, \dots, N$. The n th spoofing baseband signal takes the form

$$s_n(t) = A_n D_n [t - \tau_n(t)] C_n [t - \tau_n(t)] \exp[j2\pi\theta_n(t)] \quad (2)$$

where A_n is the carrier amplitude, $D_n(t)$ is the data bit stream, $C_n(t)$ is the spreading code, $\tau_n(t)$ is the code phase, and $\theta_n(t)$ is the negative beat carrier phase [1]. The Doppler of the n th spoofing signal is related to $\theta_n(t)$ by

$$\tilde{f}_n(t) = \frac{d}{dt} \theta_n(t) \quad (3)$$

The spoofer adds a 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




Fig. 1: Shown here are the Doppler components in single-satellite spoofer geolocation. The Doppler components corresponding to (5) are shown on the left. The Doppler components for each spoofing signal corresponding to (4) are shown in red to the right.

Doppler of the n th transmitted spoofing signal may be modeled as

$$\tilde{f}_n(t) = -\frac{1}{\lambda} \hat{\mathbf{r}}_{rn}^T(t) (\mathbf{v}_r(t) - \mathbf{v}_{sn}(t)) - \frac{c}{\lambda} (\delta \dot{t}_r(t) - \delta \dot{t}_{sn}(t)) \quad (4)$$

where λ is the carrier wavelength, c is the speed of light, $\hat{\mathbf{r}}_{rn}$ is the unit vector pointing from the n th spoofed navigation satellite to the spoofed position, both in ECEF coordinates, $\mathbf{v}_r$ is the spoofed receiver velocity, $\mathbf{v}_{sn}$ is the n th spoofed navigation satellite velocity, and $\delta \dot{t}_{sn}$ is the spoofed clock drift of the n th 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 in (4) that compensates for the relative motion between the victim and spoofer, but in the case of broadcast spoofing, this term is zero.

B. Received Doppler Model

First consider a scenario in which a moving receiver captures a transmitted signal having a constant carrier frequency. The received Doppler $f_D(t)$ at the moving receiver can be modeled as

$$f_D(t) = -\frac{1}{\lambda} \hat{\mathbf{r}}^T(t) (\mathbf{v}_r(t) - \mathbf{v}_t(t)) - \frac{c}{\lambda} (\delta \dot{t}_r(t) - \delta \dot{t}_t(t)) \quad (5)$$

where $\hat{\mathbf{r}}$ is the unit vector pointing from the transmitter to the receiver, $\mathbf{v}_r$ is the velocity of the receiver, $\mathbf{v}_t$ is the velocity of the transmitter, $\delta \dot{t}_r$ is the clock drift of the receiver, and $\delta \dot{t}_t$ is the clock drift of the transmitter. Note that this is a simplified Doppler model that neglects higher-order terms. A complete Doppler model is presented in [52]. For the purposes of this paper, the simplified model is adequate, as will be confirmed by the experimental results.

Now consider a scenario in which a moving receiver captures an ensemble of transmitted spoofing signals from a stationary terrestrial spoofer ($\mathbf{v}_t(t) = 0$), as shown in Fig. 1. An analysis of how spoofer motion affects the geolocation solution is given in a prior version of this paper [50]. 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 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 time indices for clarity, the observed Doppler of the n th spoofing signal at the moving receiver, f_n , may be written as

$$\begin{aligned} f_n &= f_D + \tilde{f}_n \\ &= -\frac{1}{\lambda} \hat{\mathbf{r}}^T \mathbf{v}_r - \frac{c}{\lambda} (\delta \dot{t}_r - \delta \dot{t}_t) \\ &\quad - \frac{1}{\lambda} \hat{\mathbf{r}}_{rn}^T (\mathbf{v}_r - \mathbf{v}_{sn}) - \frac{c}{\lambda} (\delta \dot{t}_r - \delta \dot{t}_{sn}) \end{aligned} \quad (6)$$

What makes single-satellite GNSS spoofer geolocation difficult is the $\tilde{f}_n$ term: it is typically unknown, time-varying, and different for each spoofing signal. In the case of the matched-code jammer discovered in [37], $\tilde{f}_n = 0$. One may suppose that the operator's intent in that case was not to deceive victim receivers into inferring false locations like a spoofer. When $\tilde{f}_n = 0$, 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 modeled as in (6) yields final position estimates that are biased because the spoofing signals contain the unmodeled $\tilde{f}_n(t)$ term. In the following section, a technique is presented that removes $\tilde{f}_n(t)$ and extracts $\hat{\mathbf{r}}^T(t) \mathbf{v}_r(t)$, the range-rate time history between transmitter and receiver, which can be exploited for geolocation.

III. CONCEPTUAL OVERVIEW OF BROADCAST GNSS SPOOFER GEOLOCATION

This section presents an overview of the technique originally presented in [50], [51] for spoofer geolocation. The common Doppler components across all spoofing signals from (6) are indicated below:

$$\begin{aligned} f_n &= -\frac{1}{\lambda} \hat{\mathbf{r}}^T \mathbf{v}_r - \frac{c}{\lambda} (\delta \dot{t}_r - \delta \dot{t}_t) \\ &\quad - \underbrace{\left(\frac{1}{\lambda} \hat{\mathbf{r}}_{rn}^T (\mathbf{v}_r - \mathbf{v}_{sn}) - \frac{c}{\lambda} (\delta \dot{t}_r - \delta \dot{t}_{sn}) \right)}_{\text{common}} \end{aligned} \quad (7)$$

All common Doppler terms can be lumped into a single term

$$\gamma(t) = \frac{1}{c} \hat{\mathbf{r}}^T(t) \mathbf{v}_r(t) + \delta \dot{t}_r(t) - \delta \dot{t}_t(t) + \delta \dot{t}_{\bar{r}}(t) \quad (8)$$

so that (6) may be written

$$f_n = -\frac{1}{\lambda} \hat{\mathbf{r}}_{rn}^T (\mathbf{v}_r - \mathbf{v}_{sn}) - \frac{c}{\lambda} (\gamma - \delta \dot{t}_{sn}) \quad (9)$$

Upon processing an ensemble of spoofing signals, a GNSS receiver's PVT estimator produces, at each navigation epoch, the state estimate

$$\hat{\mathbf{x}}(t) = [\hat{\mathbf{r}}^T(t), \hat{\xi}(t), \hat{\mathbf{v}}_r^T(t), \hat{\gamma}(t)]^T \quad (10)$$

which is composed of the estimated spoofed position, the estimated receiver clock bias $\hat{\xi}(t)$, the estimated spoofed velocity, and the estimated receiver clock drift $\hat{\gamma}(t)$ [53]. A brief review of PVT estimation from pseudorange and Doppler

measurements is provided in [51], [54]. Note that the estimated receiver clock bias $\hat{\xi}(t)$ will include $\delta t_{\bar{r}}$ as a component but will not in general equal $\delta t_{\bar{r}}$.

The estimated clock drift $\hat{\gamma}(t)$, on the other hand, will track $\gamma(t)$ closely provided that the PVT estimator is configured with a clock model whose process noise intensity is sufficient to accommodate the variations in $\gamma(t)$ due to spoofing. Expressed in s/s, $\hat{\gamma}(t)$ contains all common Doppler terms, since the PVT estimator attributes common-mode frequency deviations across received signals to the receiver's clock drift. Importantly, $\hat{\gamma}(t)$ is unaffected by the unknown non-common Doppler components from $\hat{f}_n(t)$ for all $n \in \{1, 2, \dots, N\}$.

The time history $\hat{\gamma}(t)$ is the key to spoofer 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 $\hat{\mathbf{r}}^T(t)\mathbf{v}_r(t)$, which for a LEO-based receiver is typically the dominant component in $\gamma(t)$. A nonlinear least-squares estimator based on $\hat{\gamma}(t)$ is developed in the next section to estimate the spoofer's position.

For a targeted spoofing attack in which the spoofer attempts to compensate for true spoofer-to-victim line-of-sight velocity, $\gamma(t)$ could contain an extra term. If this term were to vary rapidly with time, it would cause trouble for this paper's technique. Relatedly, if the targeted victim's position and velocity were somehow accurately known to the LEO-based receiver, this paper's technique could produce accurate results provided that the estimator presented in the next section were updated to account for the known victim motion. 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 $\delta \dot{t}_r(t)$, $\delta \dot{t}_i(t)$, and $\delta \dot{t}_{\bar{r}}(t)$, are nuisance terms that potentially degrade geolocation accuracy. Fortunately, their contributions are typically minor or can be estimated. Consider $\delta \dot{t}_r(t)$. If the satellite's GNSS receiver and the radio frequency (RF) front-end capturing spoofing signals are driven by the same oscillator, then $\delta \dot{t}_r(t)$ is automatically estimated by the onboard GNSS receiver, provided it is not significantly affected by the spoofing, and thus $\delta \dot{t}_r(t)$ 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 the scenario assumed in this paper, 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, as achieved in [37]. Furthermore, robustness is achieved if a zenith-facing antenna feeds the onboard GNSS receiver's RF front-end, as the gain directed towards Earth will be strongly attenuated. Finally, PVT can be trivially maintained by a multi-constellation receiver when only single-constellation spoofing signals are present.

The terms $\delta \dot{t}_i(t)$ and $\delta \dot{t}_{\bar{r}}(t)$ originate from the spoofer. Specifically, $\delta \dot{t}_i(t)$ originates from the spoofer's hardware, while $\delta \dot{t}_{\bar{r}}(t)$ originates from its spoofer's software. The former

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 [37]. The spoofed clock drift $\delta \dot{t}_{\bar{r}}(t)$ arises from the spoofer's attack configuration, and will manifest at the victim as an increment to the victim's clock drift. It can be troubling for geolocation, but a potential attacker would typically opt to keep $\delta \dot{t}_{\bar{r}}(t)$ near constant, because if $\delta \dot{t}_{\bar{r}}(t)$ 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 $\delta \dot{t}_i(t) + \delta \dot{t}_{\bar{r}}(t)$ and summarized as follows: if the spoofer allows extraordinary frequency instability in its own oscillator so that $\delta \dot{t}_i(t)$ changes too rapidly, or if it attempts to induce a quickly-varying spoofed clock drift so that $\delta \dot{t}_{\bar{r}}(t)$ changes too rapidly, geolocation accuracy is degraded but, on the other-hand, the spoofing attack becomes trivially detectable.

Section VI explores the consequences for geolocation of cases where $\delta \dot{t}_i(t)$ departs from a constant model. It also presents an analysis of how aggressively an attacker can ramp $\delta \dot{t}_{\bar{r}}(t)$ without being detected by an optimal spoofing detection strategy that monitors the receiver clock drift, and an analysis of how the rate of change in $\delta \dot{t}_i(t) + \delta \dot{t}_{\bar{r}}(t)$ translates to geolocation error.

IV. SPOOFER GEOLOCATION WITH $\gamma(t)$

This section presents the measurement model, derives the measurement noise covariance matrix, and presents the nonlinear least-squares estimator for single-satellite spoofer geolocation.

A. Measurement Model

When a GNSS receiver processes spoofing signals, it first generates spoofed GNSS observables. These GNSS observables are beset with errors, modeled as zero-mean additive white Gaussian noise (AWGN), arising from thermal noise, local electromagnetic interference, and other minor effects. At every navigation epoch, the noisy spoofed GNSS observables are fed to the receiver's PVT estimator to produce an optimal estimate of the spoofed PVT solution, including $\hat{\gamma}(t)$.

Let $\gamma[i] = \gamma(i\Delta t)$ and $\hat{\gamma}[i] = \hat{\gamma}(i\Delta t)$, where Δt is the constant PVT solution interval and $i \in \mathcal{I} = \{1, 2, \dots, I\}$ is the solution index within a given data capture interval. Let $z[i]$ denote the i th measurement to be used for spoofer geolocation, modeled as

$$z[i] = c\hat{\gamma}[i] = c\gamma[i] + w_a[i], \quad i \in \mathcal{I} \quad (11)$$

The velocity-equivalent estimation error $w_a[i]$, which has units of m/s, is a discrete-time noise process with $\mathbb{E}[w_a[i]] = 0$ and $\mathbb{E}[w_a[i]w_a[j]] = \sigma_a^2\delta_{ij}$, for all $i, j \in \mathcal{I}$. Section IV-B will justify this model's assumption that $w_a[i]$ is white (uncorrelated in time) for a sufficiently large Δt that is larger than the settling time of its phase lock loop (PLL) or frequency lock loop (FLL), and the settling time of any Kalman filter used for obtaining the spoofed fix.

As stated before, $\delta\dot{t}_r$ is assumed to be known and fully compensated; accordingly, it will be neglected hereafter. Additionally, $\delta\dot{t}_r$ is part of the spoofer's attack configuration and, for now, will be modeled as constant due to the constraints mentioned in the prior section.

A more comprehensive model is considered for $\delta\dot{t}_t(t)$. Let $\delta\dot{t}_t[i] = \delta\dot{t}_t(i\Delta t)$, $i \in \mathcal{I}$. Over a capture interval, $\delta\dot{t}_t[i]$ is modeled as

$$c\delta\dot{t}_t[i] = c\delta\dot{t}_t[0] + b[i], \quad i \in \mathcal{I} \quad (12)$$

where $\delta\dot{t}_t[0]$ represents the spoofer oscillator's constant frequency bias and $b[i]$ is a Gaussian random walk process expressed as

$$b[i] = \sum_{k=1}^i \phi[k], \quad i \in \mathcal{I} \quad (13)$$

where $v[k]$ is a discrete-time Gaussian random process with $\mathbb{E}[v(k)] = 0$, $\mathbb{E}[v[k]v[j]] = \sigma_v^2 \delta_{kj}$ and $\mathbb{E}[w_a[k]v[j]] = 0$ for all $k, j \in \mathcal{I}$, and $b[0] = 0$. Using the model in [55, Chap. 8], σ_v^2 can be characterized as

$$\sigma_v^2 = 2\pi^2 h_{-2} \Delta t c^2 \quad (14)$$

where h_{-2} is the first parameter of the standard clock model based on the fractional frequency error power spectrum [37]. Scaling by c^2 converts to units of $(\text{m/s})^2$.

Note that $\delta\dot{t}_t[0]$ and $\delta\dot{t}_r$ can be combined into a single measurement bias b_0 that is constant across the capture interval. Furthermore, the AWGN and Gaussian random walk can also be combined into a single noise term $w[i]$. Thus we have

$$b_0 = -c\delta\dot{t}_t[0] + c\delta\dot{t}_r \quad (15)$$

$$w[i] = w_a[i] + b[i], \quad i \in \mathcal{I} \quad (16)$$

Given all of this, (11) is rewritten so that the final measurement model takes the form

$$z[i] = \hat{\mathbf{r}}_i^\top \mathbf{v}_{r,i} + b_0 + w[i], \quad i \in \mathcal{I} \quad (17)$$

The associated measurement covariance matrix R for the process $w[i]$ is now derived. Clearly, $w[i]$ is zero-mean, but because it contains a Gaussian random walk term, it is correlated over time. The $[i, j]$ th element of its measurement covariance matrix is

$$\begin{aligned} R[i, j] &= \mathbb{E}[w[i]w[j]] \\ &= \mathbb{E}\left[\left(w_a[i] + \sum_{k=1}^i v[k]\right)\left(w_a[j] + \sum_{l=1}^j \phi[l]\right)\right] \\ &= \mathbb{E}[w_a[i]w_a[j]] + \mathbb{E}\left[\left(\sum_{k=1}^i v[k]\right)\sum_{l=1}^j \phi[l]\right] \\ &= \mathbb{E}[w_a[i]w_a[j]] + \sum_{k=1}^i \sum_{l=1}^j \mathbb{E}[v[k]v[l]] \\ &= \sigma_a^2 \delta_{ij} + \sigma_v^2 \min\{i, j\} \end{aligned} \quad (18)$$

From this result, the measurement covariance matrix containing the AWGN and Gaussian random walk can be written as

$$R = R_a + R_b \quad (19)$$

where

$$R_a = \sigma_a^2 \mathbb{I}_{I \times I} \quad (20)$$

$$R_b = \sigma_v^2 M_{I \times I} \quad (21)$$

where $\mathbb{I}_{I \times I}$ is the identity matrix and M is an $I \times I$ matrix with $M[i, j] = \min\{i, j\}$, $i \in \mathcal{I}$. Note that this covariance matrix is a general result that can be applied to any range-rate-based positioning technique where the transmitter clock state is unknown.

B. Effects of Estimated $\gamma[i]$

One might question the choice to model the estimation error process $w_a[i] = c(\hat{\gamma}[i] - \gamma[i])$ as white, since $\hat{\gamma}[i]$ is the product of a state estimator and it is well known that state estimation errors are correlated in time. At epoch i , let $\tilde{\mathbf{x}}[i]$ denote the sequential PVT estimator's full state estimation error, $W[i]$ its feedback gain, $F[i]$ its state transition matrix, and $P[i]$ its state covariance. The covariance between sequential state errors is given by [56, Chap. 5]

$$\mathbb{E}[\tilde{\mathbf{x}}[i+1]\tilde{\mathbf{x}}^\top[i]] \neq [\mathbb{I} - W[i+1]H[i+1]]F[i]P[i] \quad (22)$$

The correlation between $w_a[i+1]$ and $w_a[i]$ for $i \in \mathcal{I}$ can be determined by analysis of this equation since $w_a[i]$ is an element of $\tilde{\mathbf{x}}[i]$.

Consider a scenario where the spoofer induces a static location with a typical GPS satellite geometry. The state estimated by an affected receiver consists of the position, clock bias, and clock drift, as in (10). Assume the receiver's PVT estimator applies a dynamics model consistent with a static position and the clock process noise model from [55]. Furthermore, assume that measurement errors are independent, zero-mean, and Gaussian with standard deviations of 1 m and 0.5 m/s respectively for the spoofed pseudorange and Doppler measurements.

A key tuning parameter in this model is the process noise of the receiver clock drift, which is governed by the h_{-2} coefficient, as in (14). Fig. 2 shows the Pearson correlation coefficient for $w_a[i]$ between subsequent navigation epochs over various values of modeled h_{-2} as a function of the time between epochs. As the process noise and time between epochs is increased, the time correlation of sequential estimation errors is reduced. This type of analysis can be performed to help determine the measurement interval length beyond which errors in the sequential estimates $\hat{\gamma}[i]$ can be accurately approximated as AWGN. For example, Fig. 2 indicates that, for $h_{-2} \geq 3 \times 10^{-19}$, measurements spaced by 100 ms or more may be treated as independent.

If h_{-2} were increased even further, the navigation filter becomes a sequence of point solutions and, in effect, the white noise-model of $w_a[i]$ is undoubtedly correct. The selection of h_{-2} becomes a tuning parameter for the system designer. This analysis involving nominal h_{-2} values becomes relevant because currently deployed LEO-based GNSS receivers can perform this technique and may not have the flexibility to change their own process noise.




Fig. 2: The Pearson correlation coefficient between sequential estimation errors $w_a[i]$ as a function of time between estimation epochs for various values of h_{-2} . As the receiver's modeled process noise intensity increases, the time correlation between estimation errors decreases.

C. Range-rate Nonlinear Least-Squares

Now that the measurements and the measurement covariance have been defined, a batch nonlinear least-squares estimator may be developed to solve for the state $\mathbf{x}$

$$\mathbf{x} = \begin{bmatrix} \mathbf{r}_t \\ b_0 \end{bmatrix} \quad (23)$$

where $\mathbf{r}_t$ is the transmitter's ECEF position and b_0 is the unknown measurement bias. Let $\mathbf{z}$ represent the $I \times 1$ stacked measurement vector. The standard weighted nonlinear least-squares cost function is

$$J(\mathbf{x}) = \frac{1}{2} [\mathbf{z} - \mathbf{h}(\mathbf{x})]^T \mathbf{R}^{-1} [\mathbf{z} - \mathbf{h}(\mathbf{x})] \quad (24)$$

where $\mathbf{h}(\mathbf{x})$ is the nonlinear measurement model function. The optimal estimate of $\mathbf{x}$ minimizes the cost J .

The linearized measurement model H is an $I \times 4$ matrix that takes the form

$$H = \begin{bmatrix} \frac{dh_1}{d\mathbf{r}_t} & 1 \\ \vdots & \vdots \\ \frac{dh_I}{d\mathbf{r}_t} & 1 \end{bmatrix} \quad (25)$$

where

$$\frac{dh_i(\mathbf{x})}{d\mathbf{r}_t} = \mathbf{v}_i^T \left(\frac{\hat{\mathbf{r}}_i \hat{\mathbf{r}}_i^T - \mathbb{I}_{3 \times 3}}{\rho_i} \right) \quad (26)$$

is the 1×3 Jacobian of the i th range-rate measurement. The range between the receiver and the transmitter at the i th measurement is denoted ρ_i . This measurement model Jacobian is equivalent to columns 1, 2, 3, and 8 of the Jacobian presented in [52], up to a scale factor.

Enforcing an altitude constraint significantly improves the problem's observability. This can be incorporated as an additional pseudo-measurement of the transmitter's altitude with respect to the WGS-84 ellipsoid, modeled as

$$z[I+1] = h_{\text{alt}}(\mathbf{x}) + w_{\text{alt}} \quad (27)$$

where the measurement error $w_{\text{alt}} \sim \mathcal{N}(0, \sigma_{\text{alt}}^2)$ is assumed to be independent of those for $z[i]$, $i \in \mathcal{I}$. The measurement's 1×4 Jacobian is

$$H_{\text{alt}} = [\cos(\phi_{\text{lat}})\cos(\lambda_{\text{lon}}), \cos(\phi_{\text{lat}})\sin(\lambda_{\text{lon}}), \sin(\phi_{\text{lat}}), 0] \quad (28)$$

where ϕ_{lat} and λ_{lon} are the latitude and longitude of $\mathbf{r}_t$, respectively. The measurement vector $\mathbf{z}$, vector-valued function $\mathbf{h}(\mathbf{x})$, Jacobian H , and error covariance R are all augmented appropriately to include the altitude pseudo-measurement.

Finally, the estimation error's Cramér-Rao lower bound (CRLB) can be approximated as

$$P_{\mathbf{x}\mathbf{x}} = (\mathbf{H}^T \mathbf{R}^{-1} \mathbf{H})^{-1} \quad (29)$$

V. EXPERIMENTAL RESULTS

The single-satellite geolocation technique described above was verified in a joint demonstration between the University of Texas Radionavigation Laboratory (UT RNL) and Spire Global. In this experiment, an ensemble of self-consistent spoofing signals was transmitted from a ground station while an overhead LEO satellite performed a raw signal capture. This section details the experimental setup and results. Preliminary results were presented in [51], which contains a comprehensive description of the special adaptations made to deal with the spoofer's non-GNSS carrier frequency.

A. Experimental Design

The UT RNL provided a baseband binary file containing an ensemble of GNSS spoofing signals to be transmitted, a filtered and downsampled version of the "clean static" recording in the TEXBAT dataset [57]. The original recording was a high-quality 16-bit 25 Msps (complex) recording of authentic GNSS signals centered at GPS L1 from a stationary antenna on top of the former Aerospace Engineering building at UT Austin. The front-end in the original recording was driven by a 10-MHz oven-controlled crystal oscillator (OCXO). Lowpass filtering and downsampling of the original file was required to ensure the transmitted signal was contained within Spire's available bandwidth. Additionally, onboard the satellite, the S-band capture device and onboard GNSS receiver were driven by the same oscillator, allowing precise time-tagging and compensation.

The spoofing file was transmitted from a ground station located in Perth, Australia. The transmitter was driven by a temperature-controlled crystal oscillator (TCXO). The transmitted spoofing signals were centered at S-band to avoid interfering with the GNSS bands. While the ground station was transmitting the spoofing file, an overhead LEO satellite performed a raw signal capture over 20 seconds, centered at the S-band carrier and sampled at 5 Msps (complex). In practice, all processing would be done by an onboard receiver. The duration of the raw capture should be as long as a frame in the spoofed navigation message, or 30 seconds in the case of GPS L1/CA, to ensure that the entire spoofed satellite ephemeris for each spoofed satellite could be decoded. Fig. 3 shows locations relevant to the demonstration. In the context of this paper, the physical location of the transmitter (spoofer) is in Perth, Australia and the spoofed location sits atop the former Aerospace Engineering building in Austin, Texas. Note that this spoofer could also be characterized as a meacon with a long delay from reception to transmission. The goal is to geolocate the spoofer's position in Perth.



Fig. 3: Left: The spoofed location atop the former Aerospace Engineering building in Austin, Texas. Center: The actual spoofer location, a Spire Global ground station located in Perth, Australia. Right: The ground track of the Spire Global LEO satellite during the 20-second signal capture.




Fig. 4: Top: UT RNL's GRID receiver display when processing the spoofing signals. Bottom: A scatter of GRID-derived position solutions. The red dot is the spoofed position. The 3D bias is 45.9 m, mostly concentrated in the vertical direction. This error is attributed to the S-band carrier.

B. Experimental Spoofer Geolocation with $\gamma(t)$

The transmitted spoofing signals captured in LEO were processed with the UT RNL's GRID software-defined GNSS receiver [58]–[60]. Fig. 4 shows the PVT solution obtained by processing the pseudorange and Doppler measurements of the spoofing signals. The position solution is slightly biased due to the code-carrier divergence caused by shifting the original L1-centered signal to the S-band carrier [51]. On GRID's display, the 4,810 m/s clock drift (labeled $\delta tRdot$) is immediately noticeable. Of course, no oscillator on a GNSS receiver would experience a clock drift so extreme.

To coax GRID into properly processing the S-band spoofing signals, special modifications to the receiver's configuration and PVT estimator had to be made. Reconfiguring such parameters is trivial within GRID's software-defined architecture. The bandwidths of the receiver's delay lock loop (DLL) and PLL were increased to maintain lock despite the




Fig. 5: Measured Doppler time history of each received spoofing signal. Also shown is the Doppler-equivalent time history $\hat{\gamma}(t)$ (black trace), which is used for geolocation.

code-carrier divergence introduced by the S-band carrier. The bandwidth of the DLL was set to 1.7 Hz and the bandwidth of the PLL was set to 40 Hz, introducing more noise. To minimize spurious variations in $\hat{\gamma}(t)$, the receiver's dynamics model was set to 'static', consistent with an assumed static spoofed location. The receiver's innovations-based anomaly monitor was disabled to prevent rejection of the PVT solution due to the unusually high estimated clock drift-rate. Other considerations related to the S-band carrier are detailed in [51].

A Doppler-equivalent time-history $\hat{\gamma}(t)$ over 17.75 seconds is shown as the black trace in Fig. 5 along with the raw measured Doppler of each spoofing signal. The GNSS receiver allowed itself to be spoofed and the true range-rate between the LEO-based receiver and the terrestrial transmitter was lumped in the receiver's clock drift estimate as explained in Section III. The measured Doppler time history of each spoofing signal, as given in (6), follows the shape of $\hat{\gamma}(t)$ because the range-rate between the spoofer and LEO-based receiver is dominant in all traces. The deviation in the measured Doppler time history of each spoofing signal from $\hat{\gamma}(t)$ is $\tilde{f}_n(t)$, as presented earlier.

The time history of $\hat{\gamma}(t)$ was fed to the nonlinear least-squares estimator described in Section IV. The final position fix, shown in Fig. 6, was within 68 meters of the true location.




Fig. 6: Top two figures: Final spoofed position estimate (white) based on $\hat{\gamma}(t)$. Shown in red is the true spoofed location. The error of the final estimate is 68 m. The true emitter is contained within the 95% horizontal error ellipse, derived from (29), which has a semi-major axis of 6.7 km. Bottom: Post-fit residuals of $\hat{\gamma}(t)$ time history are unbiased and have a standard deviation of 0.12 m/s.

Importantly, the true emitter position lay within the estimate's horizontal 95% error ellipse. For the measurement covariance matrix, σ_a was set to 0.15 m/s, and σ_v was set to 0.0163 m/s, which is consistent with the transmitter's TCXO. The error ellipse's eccentricity is dictated by the receiver-transmitter geometry. Shown in Fig. 6 are the Doppler post-fit residuals, which are zero-mean with a standard deviation of 0.12 m/s. Such small and unbiased residuals indicate that the estimator's model for $\hat{\gamma}(t)$ is highly accurate. This experiment provides a validation of this paper's geolocation technique.

C. Experimental Spoofed Geolocation with GNSS Observables

This paper's advocated technique requires a means of getting ephemerides and clock models of the spoofed navigation satellites implied in the spoofing. But for cases in which the GNSS receiver onboard a LEO satellite cannot be configured to produce a PVT solution from the spoofed signals, yet does produce standard Doppler observables for each spoofed signal, traditional Doppler-based geolocation as in [37] can be applied to estimate the spoofed location. Of course, as shown earlier, this will yield a biased estimate of the spoofed position because the time-varying frequency term $\tilde{f}_n(t)$ is unmodeled. However, if the spoofing signals induce a static terrestrial




Fig. 7: Geolocation using the observed Doppler time history of each spoofed PRN. Each individual spoofed position estimate is biased due to the unmodeled frequency component.




Fig. 8: Top: Range-rate residuals with respect to the estimated spoofed position. Bottom: Range-rate residuals with respect to the true spoofed position. Note: the colors used in this figure to denote different PRNs are the same as those given in the legend for Fig. 7.

location, the position bias due to the nonzero $\tilde{f}_n(t)$ is small enough that the geolocation solution remains useful.

The position bias is relatively small because the Doppler time rate of change between a stationary receiver on the surface of the Earth and a GNSS satellite in medium Earth orbit is never more than 1 Hz/s, and typically much smaller. Thus the range-rate between the LEO-based receiver and the physical spoofed is the dominant term in $\tilde{f}_n(t)$. Shown in Fig. 7 are the biased position fixes and corresponding error ellipses when each $\tilde{f}_n(t)$ time history is fed as measurements to the nonlinear least-squares estimator as described in Section IV. Only two of the seven 95% error ellipses contain the true spoofed position. The spread of the spoofed position estimates is relatively tight, with the maximum error being 1.9 km. Depending on the desired accuracy requirements, this level of accuracy may be sufficient. Note that if the spoofed's induced

trajectory were dynamic rather than static, the spread of the geolocation estimates would be larger, as shown in [50].

Shown in Fig. 8 are the range-rate residuals with respect to the estimated spoofer positions (top panel) and the true spoofer position (bottom panel). In the range-rate residuals with respect to the true spoofer position, the time-varying frequency component is visible, especially for PRNs 13 and 23, which also yield the final spoofer position estimates with the largest amount of error.

VI. SPOOFER CLOCK INSTABILITY ERROR ANALYSIS

This section analyzes how transmitter clock instability translates to range-rate-based geolocation positioning error. It is important to characterize such errors as they manifest in real-world applications. In this section, assume that $\delta t_{\bar{r}} = 0$ so that the effects of actual—not induced—clock instability may be considered in isolation. The marginal contribution of transmitter clock instability to horizontal positioning error scales directly with the transmitter oscillator quality, specified by h_{-2} in (14). This general result applies to any clock quality and any capture geometry.

As an example consider the capture scenario in Section V for a 20-second capture over Perth. Table I shows the contribution of transmitter clock instability to the 95% horizontal error ellipse semi-major and semi-minor axes in the absence of all other error sources. The orientation of the error ellipse is determined by the capture geometry. In general, the semi-major axis lies in the cross-track direction of the satellite’s motion, while the semi-minor lies in the along-track direction. Table I shows that single-satellite range-rate-based geolocation is sensitive to the transmitter clock quality. Thus, a spoofer could in theory use a low-quality oscillator to degrade geolocation accuracy. But its spoofing signals would then more easily be detected by victim receivers, as will be discussed in the next section, rendering it a less-effective spoofer.

Clock Quality	h_{-2}	Semi-major [m]	Semi-minor [m]
Low-quality TCXO	3×10^{-19}	51,449	2,033
TCXO	3×10^{-21}	5,145	203
Low-quality OCXO	3×10^{-23}	514	20
OCXO	3×10^{-25}	51	2

TABLE I: Theoretical marginal contribution of transmitter clock instability to the 95% horizontal error ellipse for the capture scenario specified in Section V in the absence of all other error sources.

The importance of correctly modeling R is emphasized here using Monte Carlo trials to compare two key metrics in geolocation: root mean square error (RMSE) between the true and estimated spoofer position, and containment percentage.

For the RMSE comparison, the true range-rate time history for the 20 second capture scenario specified in Section V was computed. For each Monte Carlo trial, both a realization of Gaussian random walk consistent with a specified h_{-2} and AWGN with $\sigma_a = 0.1$ m/s were added to the true range-rate. The noisy range-rate measurements were served to the nonlinear least-squares estimator with the correct measurement




Fig. 9: Top: Monte Carlo sample RMSE as a function of h_{-2} with the capture geometry specified in Section V, for an estimator applying the correct (R) and incorrect (R_a) measurement covariance. Bottom: Percentage increase in sample RMSE when R_a is applied rather than R .

covariance R as specified in (19), and then with an incorrect measurement covariance equal to R_a (i.e., R_b in (19) was set to zero). After the 10,000 Monte Carlo trials, the sample RMSE was calculated for the sets of geolocation estimates corresponding to R and R_a . This was repeated with various h_{-2} values representative of a range of oscillators from low-quality TCXO to OCXO. The results are shown in Fig. 9.

One notes that the sample RMSE exhibited when using the correct measurement covariance R nearly achieves the CRLB. By contrast, erroneously modeling the measurement noise as AWGN, as is the case when only R_a is used, ignores the time correlation introduced by the transmitter clock instability, resulting in a greater-than 20% increase in RMSE when the transmitter is driven by a low-quality TCXO. To be sure, the degradation in RMSE is only noticeable for $h_{-2} > 3 \times 10^{-23}$, corresponding to a low-quality OCXO or worse. The increase in RMSE becomes more prominent when a low-quality oscillator drives the transmitter because in this case the unmodeled Gaussian random walk process is the dominant contributor to the measurement noise, increasing the correlation between measurements.

Although taking R_a alone as the measurement covariance is incorrect, an *unbiased* estimate is still achieved. Nonetheless, the associated estimated state error covariance becomes erroneously low. Using the correct measurement covariance produces an unbiased minimum-variance estimate with properly sized state error covariance.

In addition to yielding a worse RMSE, using R_a results in a significantly worse containment percentage within the corresponding theoretical 95% error ellipse. Containment percentage is the percentage of trials in which the true transmitter position lies within the theoretical 95% error ellipse centered at the estimated location.

A separate study of 10,000 Monte Carlo trials was con-




Fig. 10: Top: Monte Carlo containment percentage when the theoretical 95% error ellipse has been calculated with R_a alone [setting $R_b = 0$ in (19)], for various values of the underlying parameter σ_a . Bottom: Area of the corresponding theoretical 95% error ellipse as a function of σ_a . The horizontal line shows the area of the theoretical 95% error ellipse with the correct full measurement covariance R . The vertical lines in both plots indicate the true value of σ_a assumed in the Monte Carlo simulations.

ducted, again with the capture geometry specified in Section V. For each trial, both a realization of AWGN with $\sigma_a = 0.1$ m/s and a Gaussian random walk consistent with a TCXO with $h_{-2} = 3 \times 10^{-21}$ were added to the true range-rate. When the correct measurement covariance R was used, the corresponding theoretical 95% error ellipse contained the transmitter in 95.31% of trials, as expected by a properly modeled estimator. The area of this 95% error ellipse was 3.47 km².

By contrast, when R_b was neglected and only R_a was used, there was significant degradation in the containment percentage. For a case with R_a based on $\sigma_a = 0.1$ m/s, the containment percentage fell to 1.38%. Fig. 10 shows the containment percentage for identical cases except with various different values of modeled σ_a . As one would expect, increasing the modeled σ_a improves containment percentage. If σ_a were increased to 1.7 m/s, a 95% containment percentage with R_a is achieved. But this artificial inflation of σ_a comes at the cost of having a larger 95% error ellipse. Fig. 10 also shows the area of the theoretical 95% error ellipse for various values of σ_a . The area of the 95% error ellipse for $\sigma_a = 1.7$ m/s is 5.90 km², which is a 70% increase in the 95% error ellipse area when compared using to the correct measurement covariance. If σ_a were set to maintain the same 95% error ellipse area as the correct measurement covariance, a containment percentage of only 84.8% is achieved.

Properly modeling transmitter instability is thus essential in range-rate-based geolocation so that the minimum-variance estimate is calculated and the theoretical containment percentage is maintained.

VII. CONTROLLING SPOOFING DETECTION WHILE DEGRADING GEOLOCATION ACCURACY

Researchers have developed formidable defenses against spoofing based on receiver clock state monitoring [19]–[21]. A would-be spoofer has little flexibility to meddle with the spoofed clock drift $\delta\dot{t}_r(t)$ if intending to avoid detection by such defenses. It follows that a stealthy spoofer is scarcely able to purposefully degrade geolocation accuracy.

But consider a conspicuous spoofer—one willing to accept a potentially high spoofing detection rate among affected receivers performing optimal time-based spoofing detection. In this case, the spoofer is allowed more flexibility to manipulate $\delta\dot{t}_r(t)$ with the aim of either (1) inflating victim receivers' timing error, or (2) confounding geolocation based on this paper's technique. This section derives and analyzes the attack configuration that maximally increases geolocation error while maintaining a specified detection rate among affected receivers implementing an optimal receiver clock drift monitoring spoofing detection strategy.

A. Optimal Spoofing Detection via Clock Drift Monitoring

An optimal spoofing detection technique via receiver clock drift monitoring is presented here. Consider a time interval that spans $k \in \mathcal{K} = \{1, 2, \dots, K\}$ uniformly sampled navigation epochs. At the k th epoch, the distribution of a GNSS receiver's measured clock drift $\delta\dot{t}_u$ is modeled as

$$c\delta\dot{t}_u[k] \sim \mathcal{N}(c\delta\dot{t}_u[k-1], \sigma_u^2) \quad (30)$$

where

$$\sigma_u^2 = \sigma_m^2 + q \quad (31)$$

is the steady-state measurement variance. Here, σ_m^2 is the component of the variance due to the measurement noise and clock dynamics function, and q is the process noise for $c\delta\dot{t}_u$, which is related to the time between navigation epochs Δt and the GNSS receiver clock parameter h_{-2}^u by [55]

$$q = 2\pi^2 h_{-2}^u \Delta t c^2 \quad (32)$$

Let

$$\eta_k = \frac{c\delta\dot{t}_u[k] - c\delta\dot{t}_u[k-1]}{\sigma_u} \sim \mathcal{N}(0, 1) \quad (33)$$

be the normalized increment in measured receiver clock drift at the k th epoch. Assume that increments are independent so that $\mathbb{E}[\eta_k \eta_j] = \delta_{ij}$ for all $k, j \in \mathcal{K}$.

Optimal spoofing detection amounts to a hypothesis test that attempts to distinguish the null hypothesis H_0 (receiver unaffected by spoofing) from the alternative hypothesis H_1 (receiver captured by spoofing). Note that this section focuses solely on $\delta\dot{t}_r[k]$, the spoofed clock drift increment, while assuming that the spoofer's transmitter clock drift $\delta\dot{t}_t(t) = 0$, which is opposite the preceding section's assumption. Additionally, this analysis assumes a static GNSS receiver performing detection so that the focus is on time-based spoofing detection. Let

$$\mu_k = \frac{c\delta\dot{t}_r[k] - c\delta\dot{t}_r[k-1]}{\sigma_u} \quad (34)$$

be the normalized spoofed clock drift increment across one inter-epoch interval, with initialization value $c\delta\dot{t}_T[0] = 0$ at $k = 0$, the moment when the spoofer captures the receiver.

Let θ_k represent the receiver's estimated clock drift increment at the k th epoch under either hypothesis. With the foregoing setup, this can be modeled as

$$H_0 : \theta_k = \eta_k, \quad k \in \mathcal{K} \quad (35)$$

$$H_1 : \theta_k = \eta_k + \mu_k, \quad \mu_k \neq 0, \quad k \in \mathcal{K} \quad (36)$$

A two-sided locally most powerful spoofing detection hypothesis test is applied because the receiver will not know the value of μ_k under H_1 . For a single epoch, the detection statistic $\Lambda^*(\theta_k)$ is

$$\Lambda^*(\theta_k) = \theta_k^2 \quad (37)$$

and has the following distributions under H_0 and H_1

$$H_0 : \Lambda^*(\theta_k) \sim \chi_1^2 \quad (38)$$

$$H_1 : \Lambda^*(\theta_k) \sim \chi_1^2(\lambda), \quad \lambda = \mu_k^2 \quad (39)$$

where χ_n^2 and $\chi_n^2(\lambda)$ denote, respectively, the chi-squared and noncentral chi-squared distributions with n degrees of freedom and noncentrality parameter λ .

Consider detection based on data taken over a time interval that spans K navigation epochs. Let $\boldsymbol{\theta} = [\theta_1, \theta_2, \dots, \theta_K]^T \in \mathbb{R}^K$ and $\boldsymbol{\mu} = [\mu_1, \mu_2, \dots, \mu_K]^T \in \mathbb{R}^K$. The joint test statistic then becomes

$$\Lambda^*(\boldsymbol{\theta}) = \sum_{k \in \mathcal{K}} \theta_k^2 = \boldsymbol{\theta}^T \boldsymbol{\theta} \quad (40)$$

with the following distributions under H_0 and H_1 :

$$H_0 : \Lambda^*(\boldsymbol{\theta}) \sim \chi_K^2 \quad (41)$$

$$H_1 : \Lambda^*(\boldsymbol{\theta}) \sim \chi_K^2(\lambda), \quad \lambda = \sum_{k \in \mathcal{K}} \mu_k^2 = \boldsymbol{\mu}^T \boldsymbol{\mu} \quad (42)$$

An optimal-decision constant false alarm rate threshold ν^* for a fixed probability of false alarm P_F can be calculated from

$$P_F = P(\Lambda^*(\boldsymbol{\theta}) > \nu^* | H_0) = 1 - F(\nu^*; K) \quad (43)$$

where $F(\nu^*; K)$ is the cumulative distribution function of χ_K^2 evaluated at the detection threshold ν^* . The probability of detection is

$$P_D(\boldsymbol{\mu}) = P(\Lambda^*(\boldsymbol{\theta}) > \nu^* | H_1) = 1 - F(\nu^*; K, \lambda) \quad (44)$$

$$= Q_{K/2}(\sqrt{\lambda}, \sqrt{\nu^*}) \quad (45)$$

where $F(\nu^*; K, \lambda)$ is the cumulative distribution function of $\chi_K^2(\lambda)$, and $Q_m(\alpha, \beta)$ is the Marcum Q function with $m = K/2$. The hypothesis test becomes

$$\Lambda^*(\boldsymbol{\theta}) \underset{H_0}{\overset{H_1}{\geq}} \nu^* \quad (46)$$

The spoofer must optimize its attack configuration against this optimal spoofing detection strategy.

B. Expression for Geolocation Error

One of the assumptions made when developing the estimator presented in Section IV was that $\delta\dot{t}_T$ is constant. If instead $\delta\dot{t}_T(t)$ is time-varying, the measurements $\hat{\gamma}[i]$ for all $i \in \mathcal{I}$ used for geolocation become be perturbed, increasing geolocation error. Let $\epsilon[i]$ represent the unmodeled time-varying $c\delta\dot{t}_T[i]$ for all $i \in \mathcal{I}$. Then at the i th measurement epoch, $c\hat{\gamma}[i] = c\gamma[i] + \epsilon[i]$. Let the vector of measurement perturbations over the capture interval be represented as $\boldsymbol{\epsilon} = [\epsilon[1], \epsilon[2], \dots, \epsilon[I]]^T \in \mathbb{R}^I$, and let $\tilde{\mathbf{x}} = [\tilde{e}, \tilde{n}, \tilde{b}]^T$ denote the geolocation estimation error in the east direction, north direction, and frequency bias, where $\tilde{e}$ and $\tilde{n}$ are defined in the East-North-Up (ENU) frame centered at the true spoofer position. Let $\tilde{H} \in \mathbb{R}^{I \times 3}$ denote the measurement Jacobian with respect to $\tilde{\mathbf{x}}$. The error $\tilde{\mathbf{x}}$ can be calculated as

$$\tilde{\mathbf{x}} = \left(\tilde{H}^T R^{-1} \tilde{H} \right)^{-1} \tilde{H}^T R^{-1} \boldsymbol{\epsilon} = B\boldsymbol{\epsilon} \quad (47)$$

The horizontal position error vector $\mathbf{e}_h$ is defined as

$$\mathbf{e}_h = [\tilde{e}, \tilde{n}]^T \quad (48)$$

Let $\tilde{B}$ be the first two rows of B , and define the matrix $A \in \mathbb{R}^{I \times I}$ as

$$A = \tilde{B}^T \tilde{B} \quad (49)$$

The absolute horizontal positioning error e_h due to the perturbation $\boldsymbol{\epsilon}$ can then be computed as

$$e_h = \sqrt{\mathbf{e}_h^T \mathbf{e}_h} = \sqrt{\boldsymbol{\epsilon}^T A \boldsymbol{\epsilon}} \quad (50)$$

Thus, the squared horizontal geolocation error e_h^2 is related to the perturbation $\boldsymbol{\epsilon}$ by the quadratic form $\boldsymbol{\epsilon}^T A \boldsymbol{\epsilon}$.

The spoofer seeks the perturbation $\boldsymbol{\epsilon}$ that maximizes e_h so that it can maximally degrade the accuracy of geolocation by a single sensor platform performing range-rate-based geolocation via this paper's technique. Suppose that $\boldsymbol{\epsilon}$ is subject to the constraint $\|\boldsymbol{\epsilon}\| \leq \zeta$, which will be defined in the next section. The optimization problem then becomes

$$\boldsymbol{\epsilon}^* = \underset{\|\boldsymbol{\epsilon}\| \leq \zeta}{\operatorname{argmax}} \boldsymbol{\epsilon}^T A \boldsymbol{\epsilon} \quad (51)$$

To solve this problem, A is factorized as $A = Q D Q^T$, where Q is orthogonal and $D = \operatorname{diag}(d_1, d_2, \dots, d_I)$ is a diagonal matrix composed of eigenvalues of A , which are all positive. Assume that the columns of Q contain the unitary eigenvectors corresponding to eigenvalues ordered such that $d_1 \geq d_2 \geq \dots \geq d_I$. Then

$$\boldsymbol{\epsilon}^T A \boldsymbol{\epsilon} = \boldsymbol{\epsilon}^T Q D Q^T \boldsymbol{\epsilon} = \mathbf{y}^T D \mathbf{y} \quad (52)$$

where $\|\boldsymbol{\epsilon}\| = \|Q^T \boldsymbol{\epsilon}\| = \|\mathbf{y}\|$. The value of $\mathbf{y}$ that respects $\|\mathbf{y}\| \leq \zeta$ and maximizes $\mathbf{y}^T D \mathbf{y}$ is given by $\mathbf{y}^* = [\zeta, 0, 0, \dots, 0]^T$. Let $\mathbf{v}^* \in \mathbb{R}^I$ denote the unitary eigenvector corresponding to the largest eigenvalue of A . The optimal $\boldsymbol{\epsilon}$ for this optimization problem is then

$$\boldsymbol{\epsilon}^* = \pm Q \mathbf{y}^* = \pm \zeta \mathbf{v}^* \quad (53)$$




Fig. 11: The attack trajectory $\mathbf{v}[i]^* - \mathbf{v}^*[1]$ that maximizes e_h for the LEO-based receiver geometry shown in Fig. 3.

C. Jointly Optimized Spoofing Clock Drift Selection

Now that an optimal spoofing detector based on receiver clock drift has been presented, and a perturbation ϵ^* that maximizes horizontal geolocation error subject to the constraint $\|\epsilon^*\| < \zeta$ has been defined, a spoofer can develop an attack configuration for $c\delta\dot{\mathbf{t}}_{\bar{\tau}}(t)$ that maximizes e_h while maintaining a specified probability of detection. It is assumed that the spoofer has perfect knowledge of the LEO-based receiver's position and velocity, which is representative of a worst-case scenario.

Let $c\delta\dot{\mathbf{t}}_{\bar{\tau}} = c[\delta\dot{t}_{\bar{\tau}}[1], \delta\dot{t}_{\bar{\tau}}[2], \dots, \delta\dot{t}_{\bar{\tau}}[I]]^T \in \mathbb{R}^I$ represent the spoofer's discretized time-varying attack configuration for $\delta\dot{\mathbf{t}}_{\bar{\tau}}(t)$. Suppose the spoofer sets $c\delta\dot{\mathbf{t}}_{\bar{\tau}} = \epsilon^*$. Then the vector of spoofed clock drift increments over $K = I - 1$ navigation epochs is equivalent to

$$\boldsymbol{\mu} = \frac{\zeta}{\sigma_u} C \mathbf{v}^* \in \mathbb{R}^K \quad (54)$$

where

$$C = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix} \in \mathbb{R}^{K \times I} \quad (55)$$

The only task remaining for the spoofer is to determine the value of ζ so that ϵ^* can be scaled appropriately. Suppose the spoofer is willing to allow a detection probability $\bar{P}_D$ for the detection test in (46). Based on the parameters σ_u , I , and P_F , the parameter ζ can be chosen to maintain an expected probability of detection $\bar{P}_D$. Given the functional form of the probability of detection in (45), ζ must satisfy the equation

$$Q_{K/2} \left(\frac{\zeta}{\sigma_u} \|C \mathbf{v}^*\|, \sqrt{\nu^*} \right) = \bar{P}_D \quad (56)$$

Following this, the spoofed clock drift trajectory $c\delta\dot{\mathbf{t}}_{\bar{\tau}}^*$ that maximizes the geolocation error while maintaining a specified probability of detection can be represented as

$$c\delta\dot{\mathbf{t}}_{\bar{\tau}}^* = \pm \zeta (\mathbf{v}^* - \mathbf{1} \mathbf{v}^*[1]) \quad (57)$$

where $\mathbf{1}$ is the appropriately sized vector of all ones and $\mathbf{v}^*[1]$ is the first element of $\mathbf{v}^*$. Note that subtracting $\mathbf{1} \mathbf{v}^*[1]$ ensures $c\delta\dot{\mathbf{t}}_{\bar{\tau}}[1] = 0$, consistent with initialization of the spoofing attack. This subtraction does not change the optimization




Fig. 12: Worst-case geolocation error for a spoofer that optimizes $c\delta\dot{\mathbf{t}}_{\bar{\tau}}^*$ for receivers performing 1-Hz spoofing detection tests with $\sigma_m = 0.05$ m/s and for the LEO-based receiver geometry shown in Fig. 3. The geolocation error is shown over a range of $\bar{P}_D$ for two representative victim receiver clock quality levels and three representative values of P_F .

processes, it only affects the estimated frequency bias b_0 , which is merely a nuisance parameter.

To illustrate the application of this analysis, consider the following example. Suppose a spoofer wishes to choose $\delta\dot{\mathbf{t}}_{\bar{\tau}}^*$ to maximally degrade geolocation by a LEO-based receiver capturing its signals over 21 seconds with the geometry shown in Fig. 3. Further suppose the LEO-based receiver computes measurements at 1 Hz, so that $I = 21$, and sets R with σ_v consistent with a TCXO and $\sigma_a = 0.1$ m/s. The attack trajectory $\mathbf{v}^* - \mathbf{1} \mathbf{v}^*[1]$ that maximizes horizontal geolocation error is shown in Fig. 11. It is interesting to note that the spoofer allocates the greatest detection risk (largest increments) at the beginning and end of the 21-second capture, while maintaining lower risk (smaller increments) in the interim.

Now assume that spoofing-affected receivers are performing navigation solutions once per second with $\sigma_m = 0.05$ m/s. Shown in Fig. 12 is the maximum horizontal geolocation error given a triad of $\bar{P}_D$, P_F , and affected receiver clock quality. For example, if the spoofer accepts a detection rate of $\bar{P}_D = 0.5$ by receivers equipped with a TCXO having their spoofing detector set with $P_F = 10^{-3}$, the maximum e_h due to $c\delta\dot{\mathbf{t}}_{\bar{\tau}}^*$ is 8.4 km.

To give the reader an idea of how capture geometry affects the maximum horizontal geolocation error, consider the same scenario, but with a 21-second detection-and-geolocation segment beginning 30 seconds earlier. This capture geometry is more favorable for geolocation. It results in a maximum horizontal geolocation error of 2.2 km. On the other hand, consider a 21-second segment beginning 30 seconds after the original. This capture geometry is worse for geolocation. It results in a maximum horizontal geolocation error of 23.5 km. It is important to note that this worst-case error is not a limitation of this paper's technique, but a limit of single-satellite range-rate-based geolocation of GNSS spoofers in general. And it should be remembered that the foregoing

analysis is for a worst-case situation in which the spoofer knows the LEO-based receiver's position and velocity time history.

VIII. CONCLUSION

This paper presented a single-satellite, single-pass technique for locating GNSS spoofers from LEO. The technique was validated in a controlled experiment in partnership with Spire Global in which a LEO-based receiver captured GNSS spoofing signals transmitted from a ground station. An analytic expression for how actual transmitter clock instability degrades the geolocation solution was derived. Finally, geolocation positioning error as a function of worst-case spoofed clock behavior subject to a constraint on probability of detection was investigated.

ACKNOWLEDGMENTS

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.

REFERENCES

- [1] M. L. Psiaki and T. E. Humphreys, "GNSS spoofing and detection," *Proceedings of the IEEE*, vol. 104, no. 6, pp. 1258–1270, 2016.
- [2] A. Jafarnia-Jahromi, A. Broumandan, J. Nielsen, and G. Lachapelle, "GPS vulnerability to spoofing threats and review of antispoofing techniques," *International Journal of Navigation and Observation*, pp. 1–16, 2012.
- [3] L. Scott, "Anti-spoofing and authenticated signal architectures for civil navigation systems," in *Proceedings of the ION GNSS Meeting*, pp. 1542–1552, 2003.
- [4] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, Jr., "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," in *Proceedings of the ION GNSS Meeting*, (Savannah, GA), Institute of Navigation, 2008.
- [5] T. E. Humphreys, "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.
- [6] T. E. Humphreys, "Interference," in *Springer Handbook of Global Navigation Satellite Systems*, pp. 469–503, Springer International Publishing, 2017.
- [7] M. L. Psiaki and T. E. Humphreys, *Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications*, vol. 1, ch. Civilian GNSS Spoofing, Detection, and Recovery, pp. 655–680. Wiley-IEEE, 2020.
- [8] K. Radoš, M. Brkić, and D. Begušić, "Recent advances on jamming and spoofing detection in GNSS," *Sensors*, vol. 24, no. 13, p. 4210, 2024.
- [9] K. D. Wesson, J. N. Gross, T. E. Humphreys, and B. L. Evans, "GNSS signal authentication via power and distortion monitoring," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, pp. 739–754, April 2018.
- [10] J. N. Gross, C. Kilic, and T. E. Humphreys, "Maximum-likelihood power-distortion monitoring for GNSS-signal authentication," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 1, pp. 469–475, 2018.
- [11] M. L. Psiaki, B. W. O'Hanlon, S. P. Powell, J. A. Bhatti, K. D. Wesson, T. E. Humphreys, and A. Schofield, "GNSS spoofing detection using two-antenna differential carrier phase," in *Proceedings of the ION GNSS+ Meeting*, (Tampa, FL), Institute of Navigation, 2014.
- [12] B. O'Hanlon, M. Psiaki, J. Bhatti, and T. Humphreys, "Real-time spoofing detection using correlation between two civil GPS receiver," in *Proceedings of the ION GNSS Meeting*, (Nashville, Tennessee), Institute of Navigation, 2012.
- [13] J. Gross and T. E. Humphreys, "GNSS spoofing, jamming, and multipath interference classification using a maximum-likelihood multi-tap multipath estimator," *Proceedings of the ION International Technical Meeting*, Jan. 2017.
- [14] B. O'Hanlon, J. Bhatti, T. E. Humphreys, and M. Psiaki, "Real-time spoofing detection in a narrow-band civil GPS receiver," in *Proceedings of the ION GNSS Meeting*, (Portland, Oregon), Institute of Navigation, 2010.
- [15] Z. Clements, J. E. Yoder, and T. E. Humphreys, "Carrier-phase and IMU based GNSS spoofing detection for ground vehicles," in *Proceedings of the ION International Technical Meeting*, (Long Beach, CA), pp. 83–95, 2022.
- [16] Z. Clements, J. E. Yoder, and T. E. Humphreys, "GNSS spoofing detection: An approach for ground vehicles using carrier-phase and inertial measurement data," *GPS World*, vol. 34, no. 2, pp. 36–41, 2023.
- [17] C. Tanil, S. Khanafseh, M. Joerger, and B. Pervan, "An INS monitor to detect GNSS spoofers capable of tracking vehicle position," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, pp. 131–143, Feb. 2018.
- [18] B. Kujur, S. Khanafseh, and B. Pervan, "Optimal INS monitor for GNSS spoofer tracking error detection," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 1, 2024.
- [19] A. Jafarnia-Jahromi, S. Daneshmand, A. Broumandan, J. Nielsen, and G. Lachapelle, "PVT solution authentication based on monitoring the clock state for a moving GNSS receiver," in *European navigation conference (ENC)*, vol. 11, 2013.
- [20] P. Y. Hwang and G. A. McGraw, "Receiver autonomous signal authentication (RASA) based on clock stability analysis," in *Proceedings of the IEEE/ION PLANS Meeting*, pp. 270–281, IEEE, 2014.
- [21] A. Khalajmehrabadi, N. Gatsis, D. Akopian, and A. F. Taha, "Real-time rejection and mitigation of time synchronization attacks on the Global Positioning System," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 8, pp. 6425–6435, 2018.
- [22] T. E. Humphreys, "Detection strategy for cryptographic GNSS anti-spoofing," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 2, pp. 1073–1090, 2013.
- [23] A. J. Kerns, K. D. Wesson, and T. E. Humphreys, "A blueprint for civil GPS navigation message authentication," in *Proceedings of the IEEE/ION PLANS Meeting*, May 2014.
- [24] I. Fernandez-Hernandez, J. Winkel, C. O'Driscoll, S. Cancela, R. Terris-Gallego, J. A. López-Salcedo, G. Seco-Granados, A. Dalla Chiara, C. Sarto, D. Blonski, *et al.*, "Semi-assisted signal authentication for Galileo: Proof of concept and results," *IEEE Transactions on Aerospace and Electronic Systems*, 2023.
- [25] T. Mina, A. Kanhere, A. Shetty, and G. Gao, "GPS spoofing-resilient filtering using self-contained sensors and chimera signal enhancement," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 2, 2024.
- [26] J. Anderson, S. Lo, A. Neish, and T. Walter, "Authentication of satellite-based augmentation systems with over-the-air rekeying schemes," *NAVIGATION: Journal of the Institute of Navigation*, vol. 70, no. 3, 2023.
- [27] J. Anderson, S. Lo, and T. Walter, "Authentication security of combinatorial watermarking for GNSS signal authentication," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 3, 2024.
- [28] C4ADS, "Above us only stars: Exposing GPS spoofing in Russia and Syria," April 2019. <https://c4ads.org/reports>.
- [29] S. Gebrekidan, "Electronic warfare confounds civilian pilots, far from any battlefield," *The New York Times*, Nov. 2023.
- [30] J. Arraf, "Israel fakes GPS locations to deter attacks, but it also throws off planes and ships," NPR, April 2024.
- [31] A. Tangel and D. FitzGerald, "Electronic warfare spooks airlines, pilots and air-safety officials," *Wall Street Journal*, Sept. 2024.
- [32] M. Felux, P. Fol, B. Figuet, M. Waltert, and X. Olive, "Impacts of global navigation satellite system jamming on aviation," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 3, 2024.
- [33] G. S. Workgroup, "GPS spoofing: Final report of the GPS spoofing workgroup," tech. rep., OPSGROUP, 2024.
- [34] O. Osechas, F. Fohlmeister, T. Dautermann, and M. Felux, "Impact of GNSS-band radio interference on operational avionics," *NAVIGATION*, vol. 69, no. 2, 2022.
- [35] A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Unmanned aircraft capture and control via GPS spoofing," *Journal of Field Robotics*, vol. 31, no. 4, pp. 617–636, 2014.
- [36] D. M. LaChapelle, L. Narula, and T. E. Humphreys, "Orbital war driving: Assessing transient GPS interference from LEO," in *Proceedings of the ION GNSS+ Meeting*, (St. Louis, MO), 2021.

- [37] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, M. L. Psiaki, and T. E. Humphreys, "First results from three years of GNSS interference monitoring from low Earth orbit," *NAVIGATION*, vol. 68, no. 4, pp. 673–685, 2021.
- [38] A. McKibben, R. McKnight, B. C. Peters, Z. Arnett, and S. Ugazio, "Interference effects on a multi-GNSS receiver on-board a cubesat in LEO," in *Proceedings of the ION GNSS+ Meeting*, (Denver, CO), pp. 1245–1258, 2023.
- [39] Z. Clements, P. Ellis, and T. E. Humphreys, "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit," in *Proceedings of the IEEE/ION PLANS Meeting*, (Monterey, CA), pp. 458–469, 2023.
- [40] Z. Clements, P. Ellis, and T. E. Humphreys, "Pinpointing GNSS interference from low Earth orbit," *Inside GNSS*, vol. 18, no. 5, pp. 42–55, 2023.
- [41] C. Chew, T. M. Roberts, and S. Lowe, "Rfi mapped by spaceborne gnss-r data," *NAVIGATION: Journal of the Institute of Navigation*, vol. 70, no. 4, 2023.
- [42] M. J. Berkowitz, "America's asymmetric vulnerability to navigation warfare: Leadership and strategic direction needed to mitigate significant threats," *National Security Space Association*, July 2024. <https://nssaspace.org/wp-content/uploads/2024/07/NAVWAR-FINAL.pdf>.
- [43] A. Sidi and A. Weiss, "Delay and Doppler induced direct tracking by particle filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, pp. 559–572, January 2014.
- [44] D. Musicki, R. Kaune, and W. Koch, "Mobile emitter geolocation and tracking using TDOA and FDOA measurements," *IEEE Transactions on Signal Processing*, vol. 58, pp. 1863–1874, March 2010.
- [45] K. Ho and Y. Chan, "Geolocation of a known altitude object from TDOA and FDOA measurements," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 33, pp. 770–783, July 1997.
- [46] P. Ellis, D. V. Rheeden, and F. Dowla, "Use of Doppler and Doppler rate for RF geolocation using a single LEO satellite," *IEEE Access*, vol. 8, pp. 12907–12920, 2020.
- [47] P. Ellis and F. Dowla, "Performance bounds of a single LEO satellite providing geolocation of an RF emitter," in *2018 9th Advanced Satellite Multimedia Systems Conference and the 15th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, pp. 1–5, IEEE, 2018.
- [48] P. B. Ellis and F. Dowla, "Single satellite emitter geolocation in the presence of oscillator and ephemeris errors," in *2020 IEEE Aerospace Conference*, pp. 1–7, IEEE, 2020.
- [49] X. Chen, Y. Morton, W.-X. Yu, and T.-K. Truong, "GNSS spoofer localization with counterfeit clock bias observables on a mobile platform," *IEEE Sensors Journal*, 2023.
- [50] Z. Clements, P. Ellis, M. L. Psiaki, and T. E. Humphreys, "Geolocation of terrestrial GNSS spoofing signals from low Earth orbit," in *Proceedings of the ION GNSS+ Meeting*, (Denver, CO), pp. 3418–3431, 2022.
- [51] 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.
- [52] M. L. Psiaki, "Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids," *NAVIGATION*, vol. 68, no. 3, pp. 621–641, 2021.
- [53] C. Günther, "A survey of spoofing and counter-measures," *NAVIGATION*, vol. 61, no. 3, pp. 159–177, 2014.
- [54] D. Odijk, "Positioning model," in *Springer Handbook of Global Navigation Satellite Systems*, pp. 605–638, Springer International Publishing, 2017.
- [55] R. G. Brown and P. Y. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*. Wiley, 2012.
- [56] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley and Sons, 2001.
- [57] T. E. Humphreys, J. A. Bhatti, D. P. Shepard, and K. D. Wesson, "The Texas Spoofing Test Battery: Toward a standard for evaluating GNSS signal authentication techniques," in *Proceedings of the ION GNSS Meeting*, 2012.
- [58] Z. Clements, P. A. Iannucci, T. E. Humphreys, and T. Pany, "Optimized bit-packing for bit-wise software-defined GNSS radio," in *Proceedings of the ION GNSS+ Meeting*, (St. Louis, MO), pp. 3749–3771, 2021.
- [59] H. A. Nichols, M. J. Murrian, and T. E. Humphreys, "Software-defined GNSS is ready for launch," in *Proceedings of the ION GNSS+ Meeting*, (Denver, CO), 2022.
- [60] T. Pany, D. Akos, J. Arribas, M. Z. H. Bhuiyan, P. Closas, F. Dovis, I. Fernandez-Hernandez, C. Fernández-Prades, S. Gunawardena, T. Humphreys, Z. M. Kassas, J. A. L. Salcedo, M. Nicola, M. L. Psiaki, A. Rügamer, Y.-J. Song, and J.-H. Won, "GNSS software defined radio: History, current developments, and standardization efforts," *NAVIGATION*, vol. 71, no. 1, 2024.

Timing Properties of the Starlink Ku-Band Downlink

Wenkai Qin*, Andrew M. Graff[†], Zachary L. Clements*, Zacharias M. Komodromos[†], Todd E. Humphreys*

**Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin*

[†]*Department of Electrical and Computer Engineering, The University of Texas at Austin*

Abstract—We develop signal capture and analysis techniques for precisely extracting and characterizing the frame timing of the Starlink constellation’s Ku-band downlink transmissions. The aim of this work is to determine whether Starlink frame timing has sufficient short-term stability to support pseudorange-based opportunistic positioning, navigation, and timing (PNT). A second goal is to determine whether frame timing is disciplined to a common time scale such as GPS time. Our analysis reveals several timing characteristics not previously known that carry strong implications for PNT. On the favorable side, periods of ns-level jitter in frame arrival times across all satellite versions indicate that Starlink hardware is fundamentally capable of the short-term stability required to support GPS-like PNT. But there are several unfavorable characteristics that, if not addressed, will make GPS-like PNT impractical: (1) The v1.0 and v1.5 Starlink satellites exhibit once-per-second abrupt frame timing adjustments whose magnitude (as large as 100s of ns) and sign appear unpredictable. Similar discontinuities are also present in the v2.0-Mini frame timing, though smaller and irregularly spaced. (2) Episodic 15-s periods of high frame jitter routinely punctuate the nominal low-jitter frame arrival timing. (3) Starlink frame timing is disciplined to GPS time, but only loosely: to within a few ms by adjustments occurring every 15 s; otherwise exhibiting drift that can exceed 20 ppm. These unfavorable characteristics are essentially incompatible with accurate PNT. Fortunately, they appear to be a consequence of software design choices, not hardware limitations. Moreover, they could be compensated with third-party-provided corrections.

Index Terms—Starlink, signal characterization, positioning, time synchronization, low Earth orbit

I. INTRODUCTION

Global navigation satellite system technology (GNSS) is currently the most prevalent used for positioning, navigation, and timing (PNT). However, traditional GNSS is vulnerable to jamming and spoofing attacks that can leave users without the ability to navigate or synchronize time [1], and threats to traditional GNSS are multiplying dramatically [2]–[4]. According to OPSGROUP, an international association of air transport professionals, GNSS spoofing incidents increased by 500% from 2023 to 2024 [4]. To strengthen radionavigation, researchers have recently focused on augmenting traditional GNSS with large low Earth orbit (LEO) communications constellations [5], with some proposing a combined communications-PNT service for future constellations [6]. Because these constellations offer higher power and wider bandwidth, they are inherently resilient to adversarial interference. Further, the two-way high-rate connectivity afforded by broadband communications constellations enables desirable

features such as user authentication and near-zero age of ephemeris and clock models.

Researchers interested in a free-to-use radionavigation receiver have investigated opportunistic approaches to PNT, i.e., PNT extraction with no direct cooperation from the constellation operator and limited *a priori* knowledge regarding satellites’ ephemerides and signals. SpaceX’s Starlink constellation is of particular interest: it offers the widest signal availability, serving millions of subscribers worldwide with its 7,000+ satellites [7]. Opportunistic approaches using Starlink’s Ku-band signals (10.7–12.7 GHz) have already proven fruitful: researchers in several groups have independently demonstrated Doppler-based positioning with accuracy on the order of 10 m [8]–[14]. Unfortunately, Doppler-based techniques cannot approach the exquisite timing precision offered by traditional GNSS: even in the optimistic scenarios posed in [13], [14], timing accuracy is limited to no better than 0.1 ms. By contrast, pseudorange-based PNT from Starlink holds the potential for both meter-accurate positioning and ns-accurate timing [6], [15]. But whether Starlink signals could support precise pseudorange-based PNT remains an open question whose answer depends on the details of the broadcast signals, including modulation, timing, and spectral characteristics.

The authors of [16] uncovered key information regarding the Starlink downlink frame structure, synchronization sequences, and spectral characteristics. Follow-on work in [17] and [18] discovered other predictable elements of each frame, and [19] revealed that Starlink beam switching occurs at 15-s, approximately-GPS-aligned intervals. Other studies have shown simulated impacts of various clock types in LEO [20], or have developed methods for predicting LEO clock corrections, such as those developed for the GRACE mission [21], [22]. Nonetheless, no prior work has characterized the stability of the Starlink frame timing, nor investigated its precise relationship to an absolute time scale such as GPS time (GPST), despite these essential details being prerequisite to development of Starlink-based PNT, whether opportunistic or not.

Our paper closes this knowledge gap. We leverage the signal structure details revealed in [16] to conduct a focused study of Starlink timing properties. First, we evaluate the short-term frame clock stability, defined as the set of clock behaviors that manifest on the order of several seconds or less whose study indicates both the quality of crystal oscillator(s) onboard the Starlink satellite vehicles (SVs) and the predictability of any onboard clock corrections. Second, we investigate the absolute




Fig. 1: Block diagram of the simultaneous Starlink and GPS signal capture system.

frame timing characteristics; i.e., we determine the extent to which Starlink frame clocks are disciplined to GPST.

A preliminary conference version of this paper appeared as [23]. The current version extends the former with a deeper examination of short-term frame clock stability, a new section analyzing absolute frame timing, a presentation of theoretical measurement bounds, and stronger conclusions.

II. SIGNAL CAPTURE

While many particulars of our signal capture system match those presented in [16], we reiterate the major signal pathways and note some updates made for the purposes of timing analysis.

Fig. 1 shows a block diagram of the signal capture system. The Starlink signal chain begins with reception of Ku-band signals through a 40 dBi parabolic dish antenna followed by frequency downconversion to L-band and signal amplification in a low-noise block (LNB). Further downconversion, amplification, filtering, and digital sampling at rates ranging from 50.0 to 62.5 Msps occur within the radio-frequency signal analyzer (RFSA) 1. This path is equivalent to the narrowband capture chain in [16].

A parallel GPS L1 capture pathway comprises an active GNSS antenna connected to an RFSA identical to RFSA 1, within which are realized the same signal processing and sampling operations as for the Starlink signal pathway. This second RFSA is denoted RFSA 2 in Fig. 1.

A key feature of the parallel capture setup is the GPS-disciplined 10-MHz oven-controlled crystal oscillator (OCXO), which drives both Starlink and GPS L1 signal pathways, allowing for synchronous capture with a unified clock model. To verify this, we captured signals from the GNSS antenna simultaneously with both RFSA, splitting the signal through equal-length cables just before insertion into each. Both data captures were then independently processed via GRID, our laboratory’s science-grade software-defined GNSS receiver [24]–[26]. By this we found that the two RFSA’s sample trains are synchronous to better than 1 ns when configured with matching sampling rates, and offset by measurable deterministic amounts when the sampling rates were not matched but were integer multiples of one another.

To measure the system’s cable delay from the Ku-band LNB to RFSA 1, required for absolute timing analysis, we performed the same experiment but routed one branch of

the split signal in a loopback configuration to near the Ku-band antenna LNB and back before insertion into RFSA 1. Dividing the observed differential latency by two yielded the unknown cable delay. Finally, to measure the additional cable delay in the GNSS capture pathway beyond that of the Starlink pathway, we substituted a GNSS antenna for the Ku-band antenna and LNB and measured the differential signal timing via simultaneous dual-RFSA capture and post-processing through GRID.

With this setup, we can capture up to 62.5 MHz of real-time Starlink bandwidth through RFSA 1 with highly stable uniform sampling, and mark each sample with a GPST timestamp to ns accuracy. This enables us to precisely determine the arrival time of a Starlink frame in GPST and study the stability of arrival times from frame to frame.

III. FRAME STRUCTURE AND SYSTEM CONCEPTS

This section introduces the Starlink frame structure, drawn from [16], and additional concepts and terminology needed to understand Starlink operation in general and Starlink-based PNT in particular.

A. Frame Structure

As shown in Fig. 2, each Starlink downlink orthogonal frequency division multiplexing (OFDM) frame consists of 302 intervals of length $T_{\text{sym}} = 4.4 \mu\text{s}$ plus a frame guard interval T_{fg} , for a total frame period of $T_f = 1/F_f$ s, where $F_f = 750$ Hz is the frame rate. Each frame begins with the primary synchronization sequence (PSS), which is natively represented in the time domain, followed by the secondary synchronization sequence (SSS), which is formatted as a standard 4QAM OFDM symbol. The final occupied symbol interval in each frame is the coda symbol (CS), which is followed by the frame guard interval. Subsequent frames may be present or not, depending on user demand.

The time-domain modulation sequence of the PSS, and the frequency-domain symbol sequence of the SSS were revealed in [16]. These sequences allow a Starlink receiver to perform the channel estimation necessary to demodulate each OFDM symbol in the frame. In particular, the PSS allows the receiver to precisely identify the beginning of each frame, while the SSS allows it to perform equalization across all subcarriers. The PSS and SSS have a fixed phase relationship, which allows coherent integration across the combined PSS + SSS interval, permitting more accurate synchronization and carrier frequency offset (CFO) estimation, as will be further examined in Section V.

The role of the CS, called the CSS in [16], is not entirely clear. Its frequency-domain symbol sequence is highly predictable, but not entirely so, unlike the SSS [17]. Likewise, other OFDM symbols in each frame have varying degrees of predictability, whether due to pilot symbols meant to aid channel estimation or to default symbols transmitted when user data do not occupy all the payload-bearing symbols in a frame [17], [18]. For the timing analysis in this paper, the local signal replica used for correlation comprises only the coherent PSS + SSS combination so that, insofar as possible, results are uniform across all received frames.




Fig. 2: Frame layout for the Ku-band Starlink downlink along time-frequency dimensions, from [16]. Indices along the horizontal axis enumerate the 303 intervals that constitute a single frame. The quantity T is the useful (non-cyclic) OFDM symbol interval, T_g is the symbol guard interval (cyclic prefix), and F is the subcarrier spacing. Other quantities are defined in the text.

B. Assigned Beams

At any given time, a user within Starlink's $\pm 53^\circ$ latitude primary coverage area may have a direct line of sight to dozens of overhead SVs above 20° elevation [6]. Each Starlink SV is capable of simultaneously directing up to 48 downlink beams to terrestrial service cells, 16 for each of its three downlink phased arrays. A proprietary beam assignment procedure assigns each beam to a ~ 20 -km-diameter service cell [27]. A service cell may be illuminated by up to 16 beams simultaneously, two beams for each of the eight frequency channels identified in [16]. We call the N_{bi} beams deliberately directed to the i th service cell *assigned beams*.

The number of distinct SVs casting assigned beams onto the i th service cell, $N_{asi} \leq N_{bi} \leq 16$, may be less than N_{bi} because a given SV may project multiple beams onto the same cell, each on a different channel or with a different polarization (right-hand vs. left-hand circular polarization). As Starlink user density has increased over the past few years, N_{asi} and N_{bi} have generally increased. It was shown in [19] that a search procedure prioritizing high-elevation SVs was effective at finding multiple illuminating SVs.

C. Fixed Assignment Interval

Several publications have noted the existence of a 15-s Starlink network reconfiguration interval [19], [28]–[30]. Because all beam assignments remain fixed over the duration of this interval, we call it the *fixed assignment interval* (FAI). FAI boundaries appear to be approximately aligned with GPST [19]. As discussed later, noticeable transitions in Starlink frame timing behavior occur at FAI boundaries. Accordingly, we split captures along FAI boundaries when analyzing frame clock behaviors.

We write $N_{bi}(l)$ to indicate the number of assigned beams for the i th cell during the l th FAI, and $N_{asi}(l) \leq N_{bi}(l) \leq 16$ to indicate the number of distinct SVs casting assigned beams onto the i th cell during the l th FAI.

D. Side Beams

Assigned beams are not the only ones from which a receiver in a given service cell may extract PNT information. Consider Fig. 3, which shows the normalized cross-correlation of a short interval of received Starlink data against a local signal replica composed of a coherent PSS + SSS combination. The highest correlation peaks shown correspond to frames from an assigned beam, whereas the weaker peaks correspond to frames from *side beams*, or beams directed toward other service cells. The assigned beam's pre-correlation SNR over this interval is approximately 21 dB, whereas for the side beam corresponding to the lower circled correlation peaks the pre-correlation SNR is only 3 dB. Despite their lower SNR, it is possible to extract time of arrival (TOA) measurements from such side-beam peaks precisely enough to support GPS-like PNT.

The post-correlation SNR of side-beam correlation peaks depends on the side beam's pre-correlation SNR and on the processing gain afforded by correlation against known synchronization sequences or other predictable elements of a frame. Reference [31] shows that correlation against a coherent PSS + SSS replica can support reliable signal detection across a large blind Doppler and time offset search space at a low false alarm rate provided that frames have a pre-correlation SNR higher than -18 dB. If the receiver can somehow assemble a local replica of an entire frame, whether because a Starlink SV transmits a known default frame during a time of low user demand [17], or because a third party receives the frame through a high-gain antenna and provides its contents through a side channel to the receiver, the full-frame processing gain makes it possible to reduce the detection threshold for pre-correlation SNR to -40 dB [31]. A later section will examine frame TOA measurement precision as a function of pre-correlation SNR and processing gain. For now, suffice it to say that useful TOA measurements could be extracted from all the minor peaks visible in Fig. 3.

There are two possible scenarios under which side beams




Fig. 3: Normalized cross-correlation of received Starlink data against a local PSS + SSS replica yields sharp peaks at the beginning of each frame. The peaks’ primary lobe is approximately 20 ns wide for a 55-MHz-bandwidth capture. The nominal interval between frames is $T_f = 1/750$ s. The data shown are for STARLINK-30580, a v2.0-Mini SV, from signals captured in February 2024, but the frame correlation pattern is broadly representative of signals captured since 2022 for v1.0, v1.5, and v2.0-Mini SVs.

may be observed. First, the side lobes of a beam may be strong enough for the beam to be detectable by a receiver in a service cell different from the one to which the beam is assigned. Signals received in this way may be severely attenuated: for the Ku-band communications SVs studied in [32], the downlink phased array’s side lobes were attenuated by ~ 30 dB relative to the main lobe.

Second, spillover from a beam assigned to a cell near the one occupied by the receiver can yield a detectable side beam. A beam’s main lobe creates a ground coverage region whose shape can be modeled as the elliptical intersection of a cone aligned with the main lobe’s boresight and Earth’s surface [33]. Let R_{bi} represent, for a given processing gain and for a beam assigned to the i th service cell, the region within which a receiver may reliably extract precise TOA measurements, and let R_{ci} be the i th hexagonal Starlink service cell. For reliable broadband service, it must be the case that $R_{ci} \subset R_{bi}$. Moreover, recognizing that the SNR requirement for TOA measurement is much lower than that for broadband communications, one expects the spillover region $R_{soi} \triangleq R_{bi} \setminus R_{ci}$ to be large. This is especially so given that R_{bi} ’s elliptical eccentricity varies inversely with the assigned beam’s elevation angle viewed from its target service cell [33], and given that Starlink seeks to lower its minimum transmission elevation angle from 25° to 20° [34]. Thus, spillover-induced side beams are expected to be plentiful.

We denote by $N_{si}(l) \geq N_{asi}(l)$ the total number of SVs from which a receiver in the i th service cell may reliably extract precise TOA measurements during the l th FAI. Due to the

presence of side beams, $N_{si}(l)$ may be significantly larger than $N_{asi}(l)$. For Starlink-based PNT, larger $N_{si}(l)$ ensures solution robustness and tends to increase the geometric diversity of a navigation solution, which decreases the solution’s dilution of precision [35].

E. Signal Terminology

Captured signals may include simultaneous transmissions from both assigned and side beams. We define a *composite signal* as one with significant power contributions from multiple beams, whereas a *simplex signal* contains significant power only from a single beam. Within a composite signal capture, the strongest transmission from a single beam is referred to as the *dominant signal*, while weaker transmissions from other beams are called *secondary signals*. Typically, the dominant signal corresponds to an assigned beam.

We identify a particular sequence of correlation peaks during a given FAI as corresponding to a unique beam by a two-factor test: (1) the peaks have approximately the same magnitude, and (2) the peaks are spaced from one another by integer multiples of $T_f = 1/F_f = 1/750$ s according to the relevant frame clock. Given a sufficiently long capture within a single FAI, this test is adequate to associate all correlation sequences of significant magnitude with their respective beams. Thus, for example, the largest 9 peaks in Fig. 3 correspond to frames from a unique assigned beam, and the set of 12 lower peaks, including the four lower circled ones, correspond to frames from a unique side beam.

To classify single-beam signals as dominant or secondary, we estimate their SNR by comparing post-correlation peaks with a noise floor estimate, derived from the complex signal variance in intervals presumed to lack active transmissions. However, strictly speaking, it is not always possible to ensure the noise floor estimate is completely free of secondary signals.

Whether a given captured signal is composite or simplex depends on the processing gain afforded by the local signal replica. For a coherent PSS + SSS replica, as in Fig. 3, signals with pre-correlation SNR below -18 dB are effectively undetectable [31], rendering them insignificant as secondary signals.

IV. CLOCK MODELS

As with any analysis of PNT systems based on radio wave propagation, unambiguous models of the various clocks involved are key to understanding and characterizing the system.

A. Beam-Specific Clock Cascade Model

Fig. 4 presents a clock cascade model for a single Starlink downlink beam. The model is beam-specific because, as will be shown later on, frame timing behavior may differ from beam to beam over the same FAI for the same SV. This remarkable observation implies that Starlink signal timing is much different from that of traditional GNSS, in which a single clock governs the whole transmission cascade across all frequencies. To be sure, in traditional GNSS timing offsets




Fig. 4: Clock cascade model for a single Starlink downlink beam.

may be present between different spreading codes on the same carrier and across multiple carriers, resulting in so-called differential code biases [36]. But the code rates as transmitted are all a constant multiple of a single base clock's reference frequency. By contrast, the frame sequences on different beams from the same Starlink SV may differ in both time offset and rate, a fact with significant implications for Starlink-based PNT.

The root of each clock cascade is a base oscillator. As will be shown, this oscillator is shared across all beams from a given SV. An onboard GNSS receiver driven by the base oscillator produces a clock disciplining signal s_d that corrects the base oscillator towards GPST. Together, the base oscillator and the GNSS receiver form a closed feedback loop that produces the GNSS-disciplined timing signal t_d .

B. Frame and Carrier Clocks

The GNSS-disciplined timing signal t_d drives the frame and carrier clocks. In turn, the frame clock's signal t_f governs the timing of frames transmitted by the SV, and the carrier clock's signal t_c drives the carrier onto which the information-bearing frames are modulated. The frame clock is likely a software-based clock whose output t_f depends on when baseband frames are loaded into buffers for mixing to radio frequency (RF) and subsequent transmission. The carrier clock is likely transparent, meaning that $t_c = t_d$, but it is represented in a manner identical to the frame clock for full model generality.

Note that in traditional GNSS the clock driving each SV's spreading code (analogous to the frame clock) also drives the underlying carrier that the spreading code modulates [37]. Thus, the code and the carrier—as transmitted—are locked together such that the code chipping rate is a constant rational multiple of the carrier frequency. One of this paper's key findings, elaborated later, is that such is not the case for the Starlink Ku-band downlink signals. Instead, the frame and carrier clocks operate somewhat independently, which is why they are represented separately in Fig. 4.

We represent frame and carrier clock offsets from t_d as Δt_f and Δt_c . These are related to t_d , t_f , and t_c by

$$t_d(t) = t_f(t) - \Delta t_f(t) \quad (1)$$

$$t_d(t) = t_c(t) - \Delta t_c(t) \quad (2)$$

where t represents true time, or time according to an ideal clock, such as is closely realized by GPST [38]. In this paper,




Fig. 5: Frame sequence timing diagram showing the transition from the $(l-1)$ th FAI to the l th FAI.

true time and GPST are taken to be synonymous. The frame and carrier clocks are related to t by

$$t = t_f(t) - \delta t_f(t) \quad (3)$$

$$t = t_c(t) - \delta t_c(t) \quad (4)$$

where $\delta t_f(t)$ and $\delta t_c(t)$ are the frame and carrier clock offsets.

The time derivatives of $\delta t_f(t)$ and of $\delta t_c(t)$ are called the frame and carrier clock drift. They are equivalent to the instantaneous fractional frequency deviation, written generically as $y(t)$, on which clock stability analysis is based [39].

C. Discrete-Time Frame Clock

The frame sequence timing diagram in Fig. 5 offers further details about the frame clock. Each frame as transmitted has a duration of exactly T_f according to the frame clock. Within each FAI, the frame slot index increments from 0 to $N_a - 1$, with $N_a = 11250$ being the number of frame slots in one FAI. Each FAI starts at the beginning of frame slot 0 and lasts $N_a T_f = 15$ s. The interval of unoccupied frame slots at the beginning of the l th FAI, called the FAI guard interval $T_{ag}(l) = N_{ag}(l)T_f$, spans a variable number of frame slots $N_{ag}(l)$. Note that, for any FAI index l , frame slot $N_{ag}(l)$ is occupied by definition, but other frame slots may not be occupied.

Let $t_f(l, m)$ be the frame clock time at the instant when the frame in the m th frame slot of the l th FAI begins to pass through the phase center of the SV's downlink antenna, where l and m are zero-based indices. By definition, we take this to be

$$t_f(l, m) \triangleq 15l + mT_f \quad (5)$$

Let $t^*(l, m)$ and $\delta t_f(l, m)$ be the corresponding GPST and frame clock offset. Then

$$t^*(l, m) = t_f(l, m) - \delta t_f(l, m) \quad (6)$$

Another of this paper's key findings is that $\delta t_f(l, 0) \approx 0$. Stated differently, a Starlink SV's frame clock departure from GPST at the beginning of each FAI is small—typically less than a few ms.

D. Receiver Clock

Now consider the clock of a receiver tracking signals from the Starlink downlink. The receiver clock time t_r is related to true time t by

$$t = t_r(t) - \delta t_r(t_r) \quad (7)$$

The receiver clock offset $\delta t_r(t_r)$ is represented as a function of t_r because it is natively ordered in receiver time in the course of solving for a position and time solution. The time derivative of δt_r with respect to t , denoted $\dot{\delta t}_r(t_r(t))$, is called the receiver clock drift.

Let $t_r(l, m)$ be the time of reception, according to the receiver clock, of the frame transmitted at true time $t^*(l, m)$. Let $\delta t_r(l, m)$ be the corresponding receiver clock offset and $t_*(l, m)$ be the corresponding true time of reception. More precisely, $t_r(l, m)$ is the receiver clock time at which the frame transmitted at true time $t^*(l, m)$ from the satellite's downlink antenna's phase center first reached the receiver antenna's phase center. The receipt time $t_r(l, m)$ can be related to $t_*(l, m)$, $t^*(l, m)$, and $t_f(l, m)$ by

$$t_*(l, m) = t_r(l, m) - \delta t_r(l, m) \quad (8)$$

$$t^*(l, m) = t_r(l, m) - \delta t_r(l, m) - \delta t_{\text{tof}}(l, m) \quad (9)$$

$$t_f(l, m) = t_r(l, m) - \delta t_r(l, m) - \delta t_{\text{tof}}(l, m) + \delta t_f(l, m) \quad (10)$$

where $\delta t_{\text{tof}}(l, m)$ is the frame's time of flight from transmission to reception, expressed as an interval in true time.

V. MEASUREMENT ERROR BOUNDS

Here we provide theoretical bounds on the Doppler and TOA error variance based on channel estimation using the PSS + SSS. We do this both to illustrate the level of precision possible with Starlink-based PNT and to prepare for later analysis investigating the relationship between frame and carrier Doppler.

A. Doppler Error Variance

A lower bound on Doppler frequency error variance for a signal $s(t)$ with energy E , noise spectral density N_0 , and mean time epoch t_0 can be expressed in the form [40, Eq. 11.27]

$$\sigma_{f_b}^2 \geq \frac{1}{2\alpha^2 \frac{E}{N_0}} \quad (11)$$

where α^2 is the effective squared time duration, defined as

$$\alpha^2 \triangleq \frac{\int_{-\infty}^{\infty} (2\pi(t - t_0))^2 |s(t)|^2 dt}{\int_{-\infty}^{\infty} |s(t)|^2 dt} \quad (12)$$

The PSS signal given in [16, Eq. 34] has an effective time duration of $\alpha_{\text{PSS}} = 7.745 \mu\text{s}$, while the SSS signal given in [16, Eq. 37] has an effective time duration of $\alpha_{\text{SSS}} = 7.8428 \mu\text{s}$. The coherent combination of PSS and SSS in a single frame results in $\alpha_1 = 15.975 \mu\text{s}$, and the coherent combination across two frames results in $\alpha_2 = 4.2 \text{ ms}$.

Fig. 6 plots the Doppler frequency root-mean-square error (RMSE) bound in units of Hz for the PSS-only, SSS-only, single-frame PSS + SSS, and two-frame PSS + SSS signals (i.e., coherent correlation across the PSS + SSS combination in a pair of adjacent frames). One may conclude from Fig. 6 that a Doppler estimation based on a single coherent PSS + SSS interval yields an RMSE far worse than assumed in [13]. Clearly, for Doppler-based Starlink PNT, coherent processing across many more than two OFDM symbols will be required.




Fig. 6: Doppler frequency RMSE bounds for various combinations of coherent processing based on the PSS and SSS.

B. Time of Arrival Measurement Error Variance

Similar to the foregoing Doppler frequency bound, a lower bound on TOA error variance for a signal with Fourier transform $S(f)$, energy E , and noise spectral density N_0 can be expressed in the form [40, Eq. 11.15]

$$\sigma_\tau^2 \geq \frac{1}{2\gamma^2 \frac{E}{N_0}} \quad (13)$$

where γ^2 is the effective squared bandwidth defined as

$$\gamma^2 \triangleq \frac{\int_{-\infty}^{\infty} (2\pi f)^2 |S(f)|^2 df}{\int_{-\infty}^{\infty} |S(f)|^2 df}. \quad (14)$$

The PSS signal given in [16, Eq. 34] has a Fourier transform

$$S_{\text{PSS}}(f) \triangleq \frac{1}{F_s} \sum_{k=-N_g}^{N-1} \text{rect}\left(\frac{f}{F_s}\right) \exp\left(-j2\pi(k + N_g)\frac{f}{F_s}\right) p_k \quad (15)$$

and an effective bandwidth of $\gamma_{\text{PSS}} = 4.3482 \times 10^8 \text{ Hz}$. Similarly, the SSS signal given in [16, Eq. 1.37] has a Fourier transform

$$S_{\text{SSS}}(f) \triangleq \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} X_{m1k} \exp(-j2\pi F T_g k) G_s(f - Fk) \quad (16)$$

$$G_s(f) \triangleq \mathcal{F}[g_s(t)] = T_{\text{sym}} \text{sinc}(T_{\text{sym}} f) \exp(-j\pi T_{\text{sym}} f) \quad (17)$$

and an effective bandwidth of $\gamma_{\text{SSS}} = 4.3658 \times 10^8 \text{ Hz}$. The PSS and SSS signal combined in a single frame have a Fourier transform of $S_1(f) = S_{\text{PSS}}(f) + \exp(-j2\pi T_{\text{sym}} f) S_{\text{SSS}}(f)$. In the case of coherent processing of the PSS + SSS across two frames, this becomes $S_2(f) = S_1(f) + \exp(-j2\pi T_f f) S_1(f)$.

Fig. 7 plots the TOA RMSE bound in units of meters for the PSS-only, SSS-only, single-frame PSS + SSS, and two-frame PSS + SSS signals. Due to the large bandwidth of the synchronization sequences, sub-centimeter accuracy is theoretically achievable at moderate SNR.




Fig. 7: TOA RMSE bounds for various combinations of coherent processing based on the PSS and SSS.

VI. DATA PREPROCESSING

We capture Starlink signals through RFSA 1 as shown in Fig. 1 at a complex sampling rate ranging from 50.0 to 62.5 MHz. Captured signals are then upsampled to the Starlink information symbol rate $F_s = 240$ MHz so that all subsequent operations may proceed as if the full 240-MHz Starlink channel bandwidth had been captured. The upsampled signal is cross-correlated against a local signal replica consisting of the coherent PSS + SSS combination to produce correlation peaks like those shown in Fig. 3. For the m th frame, let

$$x_{m01}(t) \triangleq \begin{cases} x_{m0}(t) & 0 \leq t < T_{\text{sym}} \\ x_{m1}(t - T_{\text{sym}}) & T_{\text{sym}} \leq t < 2T_{\text{sym}} \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

be the coherent concatenation of the time-domain PSS and SSS functions from [16]. Because the PSS and SSS are present in all Starlink downlink frames and are identical for all frames $m \in \mathbb{Z}$ and all SVs, $x_{m01}(t)$ can be used for correlation against all captured data.

The full discrete-time local signal replica used for correlation is the product of $x_{m01}(t)$ and a complex exponential:

$$y_{m01}[k] \triangleq x_{m01}(kT_s(1 - \beta)) \exp(j2\pi [F_c(1 - \beta) - \bar{F}_c] kT_s) \quad (19)$$

Here, $T_s = 1/F_s$ is the sampling interval, F_c is the center frequency of the OFDM channel, β is the CFO parameter, and $\bar{F}_c \approx F_c$ is the center frequency to which the receiver is tuned. Correlation proceeds in blocks of between 30 and 60 frame intervals T_f . Over each block, a constant β is applied that maximizes correlation peak magnitudes within the block. This process amounts to batch sequential frequency tracking.

A high-precision TOA measurement is extracted from the correlation peak for each frame received. Unless otherwise noted, all frame timing measurements were extracted from dominant signals with pre-correlation SNR exceeding -1 dB, which, upon correlation against the replica signal in (19), yields a post-correlation SNR of at least 25 dB for a ≥ 50 -MHz captured bandwidth [31]. Let $\mathcal{M}_l$ be the set of occupied-

frame indices for the l th FAI. A sequence of frame TOA measurements corresponding to $\mathcal{M}_l$ is extracted from the cross-correlation function. These are modeled as

$$\tilde{t}_r(l, m) = t_r(l, m) + w_r(l, m), \quad m \in \mathcal{M}_l \quad (20)$$

where $w_r(l, m)$ is zero-mean Gaussian measurement error with variance $\sigma_w^2(l, m)$. The m th measurement $\tilde{t}_r(l, m)$ is the receiver time of the discrete sample instant at which the correlation peak for the m th frame is maximized. As can be appreciated by examining the zoomed inset in Fig. 3, the measurement errors $\{w_r(l, m) \mid m \in \mathcal{M}_l\}$ contain errors due to (1) nearest-sample quantization of the maximizing location, (2) thermal noise, and (3) peak rounding caused by filtering in the RFSA for the relatively narrow bandwidth captured (e.g., 55 MHz for Fig. 3 as compared to the full $F_s = 240$ -MHz bandwidth). Their empirical standard deviation is approximately $\sigma_w = 2$ ns, or 0.6 m in equivalent length, which is well above the single-frame PSS + SSS RMSE bound for $\text{SNR} \geq -1$ in Fig. 7. Nonetheless, this level of measurement precision is more than adequate to study the short- and long-term Starlink frame clock behavior because, as will be shown, variations in $\delta t_f(t)$ are typically much larger than 2 ns.

VII. SHORT-TERM FRAME CLOCK STABILITY

This section explores short term (within a single FAI) Starlink frame clock behavior. The emphasis here is on relative timing and high-frequency variations in the frame clock offset δt_f , as opposed to absolute timing and low-frequency variations. Rearranging (10) to isolate $\delta t_f(l, m)$ yields

$$\delta t_f(l, m) = t_r(l, m) - t_r(l, m) + \delta t_r(l, m) + \delta t_{\text{tof}}(l, m) \quad (21)$$

Let $\delta t'_f(l, m)$ be equivalent to $\delta t_f(l, m)$ but with a 3rd-order polynomial fit across all $m \in \mathcal{M}_l$ removed. Assume like notation for the other terms in (21). Analysis of $\delta t'_f(l, m)$ is sufficient to characterize the short-term properties of the frame clock because $\delta t'_f(l, m)$ retains the high-frequency variations present in $\delta t_f(l, m)$. The detrended $\delta t_f(l, m)$ can be modeled as

$$\begin{aligned} \delta t'_f(l, m) &= t'_f(l, m) - t'_r(l, m) + \delta t'_r(l, m) + \delta t'_{\text{tof}}(l, m) \\ &\approx -t'_r(l, m) \\ &\approx -\tilde{t}'_r(l, m) \end{aligned} \quad (22)$$

This approximation is explained as follows: For $m \in \mathcal{M}_l$, $t_f(l, m)$ is an affine function of m , which implies $t'_f(l, m) = 0$. Likewise, $\delta t'_r(l, m) \approx 0$ because the receiver clock is a GPS-disciplined OCXO with negligible frequency error. Finally, $\delta t'_{\text{tof}}(l, m) \approx 0$ because the time of flight to an SV in LEO can be modeled to better than 1 ns over a 15-s FAI as a 3rd-order function.

In summary, for purposes of a short-term frame clock stability analysis, $\delta t'_f(l, m)$ is a valid proxy for $\delta t_f(l, m)$, and $\tilde{t}'_r(l, m)$, the 3rd-order-polynomial-detrended version of the frame TOA measurement $\tilde{t}_r(l, m)$, is equivalent to $\delta t'_f(l, m)$ to within a sign reversal and the ns-level measurement error $w_r(l, m)$.




Fig. 8: Top: $\tilde{t}'_r$ vs. t_f for an example capture occurring within a single FAI. Bottom: $\tilde{t}'_r$ (blue), piecewise 2nd-order polynomial fits to the truncated inter-adjustment segments (green), and $\tilde{t}''_r$, truncated inter-adjustment segments with first-order discontinuity eliminated (black). The data shown are for STARLINK-5141, a v1.5 SV, from signals captured in November 2023.

Note that until near the end of this section, we limit our analysis to v1.0 and v1.5 SVs, as their frame clocks behave differently from those of v2.0-Mini SVs.

A. Periodic Frame Clock Adjustments

The top panel of Fig. 8 shows an example $\tilde{t}'_r$ trace for an 8-s interval within a single FAI. Interpreting $\tilde{t}'_r$ as a proxy for the detrended frame clock deviation $\delta t'_f$, it is clear that the Starlink frame clock exhibits abrupt adjustments at a regular 1-Hz cadence. We believe these are the result of coarse GNSS disciplining of the base oscillator. Similar adjustments at the same 1-Hz cadence were evident in every capture of signals from Starlink v1.0 and v1.5 SVs. If due to GNSS disciplining, such large and frequent adjustments—up to several hundred ns at 1 Hz—reflect a base oscillator with poor stability.

Abrupt and coarse adjustments to the frame clock are obviously undesirable for pseudorange-based PNT. Unless they can be modeled or eliminated by some differential scheme, such adjustments would cause large errors in pseudorange modeling, and thus in position and timing estimation. They act, in effect, like the clock dithering implemented to intentionally degrade GPS accuracy under the Selective Availability program discontinued in May 2000 [41].

To assess their predictability and other characteristics, we analyzed 281 adjustments associated with 12 unique Starlink v1.0 and v1.5 SVs and made the following observations.

1) *Cadence*: Frame clock adjustment opportunities occur at an almost perfectly regular 1-Hz cadence. Of the 281 adjustments studied, all but one arrived within a few ms of an integer second from the previous one, according to the frame clock t_f . The single outlier, from a STARLINK-5666 (v1.5) capture, arrived 100 ms earlier than expected. At each




Fig. 9: Top: Cumulative probability distribution for Starlink v1.0 and v1.5 frame clock adjustment amplitudes. Bottom: Histogram of adjustment amplitudes.

opportunity, an adjustment may occur or not—note the lack of adjustment at 6.5 s in Fig. 8.

2) *Alignment with Respect to FAI*: One might expect frame clock adjustment opportunities to be aligned to the FAI such that 15 inter-adjustment intervals fit neatly within one FAI. This is not the case; there appears to be no fixed relationship between FAI boundaries and the 1-Hz adjustment opportunities. On this observation and a supporting one presented in Section IX rests our decision to model the GNSS disciplining of the base oscillator as occurring prior to and independent of any further adjustments made by the frame clock, as illustrated in Fig. 4. Thus, the frame clock t_f exhibits the 1-Hz disciplining adjustments present in t_d , but the adjustment amplitude and timing are not coordinated with the frame timing.

3) *Probability Distribution of Adjustment Amplitudes*: If the 1-Hz frame clock adjustments were quantized (say, occurring only in 20-ns steps), this would offer hope of developing an adjustment compensation mechanism within a pseudorange-based position and time estimator. Alas, this is not the case. Instead, the adjustment amplitudes appear to be smoothly distributed with mean $\mu_a = 20$ ns and standard deviation $\sigma_a = 117$ ns, as shown in Fig. 9. A Shapiro-Wilk goodness-of-fit test for normality yielded $p = 0.3$, suggesting a reasonably close alignment with the Gaussian distribution [42].

With $\sigma_a = 117$ ns, the adjustment amplitudes are commensurate with the symbol guard interval (the cyclic prefix) for Starlink's OFDM symbols, which is $T_g = 133$ ns [16]. Thus, one might expect the adjustments to cause inter-symbol interference (ISI) and thus a degradation in communications throughput. But close inspection reveals that the adjustments are not instantaneous: they occur slowly relative to the OFDM symbol rate, with a rise time of about 25 ms and a settling time of 150 ms. Thus they pose no risk of increased ISI.

4) *Frequency Adjustments*: For further insight into adjustment characteristics, we smooth $\tilde{t}_r'$ by eliminating the time shifts due to the 1-Hz adjustments. To perform this smoothing, we fit a 2nd-order polynomial to the inter-adjustment segments in $\tilde{t}_r'$, truncated to exclude the leading 0.15 s and final 0.02 s of data to avoid the adjustments' transient effects. The truncated inter-adjustment polynomial fits are shown in green in the bottom panel of Fig. 8. They and their corresponding data segments are then shifted in such a way as to minimize first-order discontinuity. That is, for each inter-adjustment segment in the FAI, we minimize the difference between the final value of the segment's shifted polynomial fit and the initial value of the next segment's shifted fit.

The resulting time series, denoted $\tilde{t}_r''$ and exemplified by the black trace in the bottom panel of Fig. 8, reveals that the 1-Hz disciplining adjustments not only impose a time shift in the base oscillator's time t_d but also a frequency shift. Note in Fig. 8 the abrupt change in the slope of $\tilde{t}_r''$ at the adjustment opportunity just prior to $t_f = 6$ s. As with time shifts, frequency shifts were found to occur randomly at each adjustment opportunity and exhibit apparently random, unquantized amplitudes. The presence of frequency adjustments further complicates any effort at modeling δt_f to obtain accurate position and time estimates from pseudorange measurements.

B. Nominal Jitter

A further processing step allows us to assess the high frequency variations (jitter) in $\tilde{t}_r$. We subtract from $\tilde{t}_r'$ the 2nd-order polynomial fits (the green segments in Fig. 8) of each truncated inter-adjustment segment to flatten the time series. We denote this flattened time series as $\tilde{t}_r'''$, an example of which is shown in the top plot of Fig. 10. Under nominal conditions for v1.0 and v1.5 SVs, the RMS value of $\tilde{t}_r'''$ ranges between 1.7 and 2.5 ns. A large contributor to this jitter is nearest-sample quantization noise, which, for $y_{m01}[k]$ from (19) sampled at $F_s = 240$ MHz, has an RMS value of $1/\sqrt{12}F_s = 1.2$ ns [43]. Assuming this noise is independent of other sources of jitter, the RMS contribution of the other sources ranges from 1.5 to 2.2 ns. Some of this is due to thermal noise in the receiver, so 2.2 ns serves as an upper bound on the jitter in the SVs' frame clock deviation δt_f under nominal behavior. This implies that the frame clock of Starlink v1.0 and v1.5 SVs is capable of maintaining jitter at the ns level, which, setting aside the 1-Hz time and frequency adjustments and the lower-frequency variations in δt_f , is adequately low to support accurate pseudorange-based PNT.

C. 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 $\tilde{t}_r'''$ time histories, such as the black trace in Fig. 8, for 18 separate FAIs. The duration of these $\tilde{t}_r'''$ time histories ranged from 10 to 15 s. The data originate from 9 unique v1.0 and v1.5 Starlink SVs whose signals were captured during 2022 and 2023 and whose frame timing was derived from




Fig. 10: Top three panels: The flattened time series $\tilde{t}_r'''$ for the same Starlink SV during three contiguous FAIs, showing nominal, oscillatory, and excursive high-frequency frame clock behavior. The data shown are for STARLINK-2119, a v1.0 SV, from signals captured in November 2023. Bottom panel: Histogram of per-FAI $\tilde{t}_r'''$ RMS values for all v1.0 and v1.5 SVs studied. A single RMS measurement was extracted from the $\tilde{t}_r'''$ time history corresponding to the dominant signal of each FAI.

dominant signals with nominal RMS levels (2.5 ns or below) in the corresponding $\tilde{t}_r'''$ traces. The resulting overlapping Allan deviation, shown in Fig. 11, reveals the average short-term stability of v1.0 and v1.5 frame clocks. Because frame slot occupancy is based on user demand, and therefore stochastic, the Allan deviation is derived from irregularly spaced data, as in [44]. But for every averaging time τ shown in Fig. 11, there were at least 1500 samples from which to estimate the corresponding Allan variance $\sigma_y^2(\tau)$, enough to yield highly accurate variance estimates.

Given the processing and data selection involved in creating the $\tilde{t}_r'''$ time histories on which the composite Allan deviation plot shown in Fig. 11 is based, which tends only to remove variation, the plot should be taken as a lower bound on the frame clock stability of v1.0 and v1.5 Starlink SVs. This best-case stability is broadly consistent with a temperature-compensated crystal oscillator (TCXO). For example, at an averaging time of $\tau = 1$ s, the fractional frequency deviation is $\sigma_y(\tau) = 2.5 \times 10^{-9}$, which is what one would expect from an average-quality TCXO. Thus, we may conclude that the short-term stability of the Starlink v1.0 and v1.5 frame clocks is no better than that of a TCXO, though, as will be shown, it can episodically be much worse. One should bear




Fig. 11: Overlapping Allan deviation of a composite frame clock constructed from 18 $\tilde{t}_r''$ time histories originating from 9 distinct v1.0 and v1.5 Starlink SVs.

in mind, however, that a TCXO-quality base oscillator would not preclude highly accurate pseudorange-based PNT. It is true that all traditional GNSS SVs employ highly stable atomic oscillators [39], but a PNT service fused with a broadband LEO communications service could get by with much cheaper and less stable clocks by sending near-zero-age clock models to users through its high-capacity communications channels. For example, assuming the Allan deviation shown in Fig. 11, clock model updates at 1 Hz (corresponding to an averaging time of $\tau = 1$ s) would be sufficient to keep satellite clock modeling errors in the PNT solution below about 2.5 ns.

D. Anomalous Frame Clock Behavior

The foregoing jitter and stability results apply only for nominal behavior of the Starlink frame clock, which holds for approximately half of the FAIs studied. The other half manifest various modes of high-frequency frame clock instability. We do not know the underlying cause of these anomalous frame clock behavior modes, but we note here certain patterns of behavior and offer conjectures about their meaning.

Consider the top three panels of Fig. 10, which show the flattened time series $\tilde{t}_r'''$ for 10-s intervals within three separate FAIs. The top panel shows nominal behavior, with 1.8 ns RMS, and acts as a baseline for comparison. The second panel shows much higher-RMS variations with significant time correlation. In this example trace, the deviations are bounded between -13 ns and 28 ns and include strong frequency components at 7.2 and 13.7 Hz. We classify anomalous frame clock behavior as oscillatory, as in this example, when one or two spectral components dominate. The third panel shows a $\tilde{t}_r'''$ time history with sudden and irregularly spaced deviations, e.g., at the 2.6 s and 5.4 s marks. We classify such behavior as excursive.

The bottom panel of Fig. 10 shows a histogram of $\tilde{t}_r'''$ RMS values for all v1.0 and v1.5 SVs studied. Each sample contributing to the histogram is a scalar RMS measurement extracted from the $\tilde{t}_r'''$ time history corresponding to the dominant signal of a single FAI. Of the 48 FAIs studied,

- 22 exhibited nominal frame clock behavior, with RMS values below 2.5 ns;

- 19 exhibited oscillatory behavior, with RMS values above 2.5 ns and having a principal frequency component typically residing between 12 and 14 Hz;
- 3 exhibited excursive behavior;
- and 4 exhibited other behaviors, such as a mix of oscillatory and excursive modes, or elevated RMS values but without sudden excursions or dominant frequency components.

A clue to the origin of anomalous frame clock behavior may be found in the following remarkable observation: Whatever mode the frame clock manifests—whether nominal, oscillatory, excursive, or otherwise—invariably persists during a full FAI, but can switch to a different mode in the next FAI, even for the same SV. In fact, the traces in the top three panels of Fig. 10 are for the same Starlink SV during three contiguous FAIs. From this clue we conclude that anomalous behavior is connected to satellite hardware or software configuration changes that occur at FAI boundaries. For example, it may be that each of the three traces in Fig. 10 comes from an assigned beam cast by a different one of the serving Starlink SV’s downlink phase arrays, of which each SV has three. If the three phased arrays are driven by separate clocks, each with its own characteristic high-frequency behavior, this would explain the FAI-aligned frame clock mode switching. Alternatively, it may be that the baseband frame assembly process is governed by parameters set in software that remain fixed over each FAI, and that some parameter combinations lead to nominal frame clock behavior whereas others lead to anomalous modes.

One might suspect that apparent frame clock mode switching is actually unrelated to any clock but instead due to complications in the TOA measurement process. For example, oscillatory measurement errors might be caused by the interaction of assigned beams and side beams whose frame slots are closely aligned in composite signals. Relatedly, excursive measurement errors could be caused by the TOA measurement process occasionally misidentifying a side peak of the PSS + SSS autocorrelation function as the primary peak in low SNR conditions. Both of these possibilities have been investigated and discarded. In fact, anomalous frame clock modes can manifest even in simplex signals (no significant side beam peaks) with extremely high SNR. For example, an FAI captured from STARLINK-3894 (v1.5) showed a mix of oscillatory and excursive frame clock behavior with $\tilde{t}_r'''$ RMS equal to 15 ns despite being an extraordinarily clean simplex signal with post-correlation SNR of 41 dB. Conversely, composite signals (one or more significant side beam signals), even with relatively low SNR, routinely yield $\tilde{t}_r'''$ RMS below 2.5 ns. Clearly, the anomalous frame clock behavior is inherent in the frame clock deviation δt_f and not in any oddity of the TOA measurement process.

As with the 1-Hz clock adjustments, the anomalous frame clock variations discussed here would tend to degrade pseudorange-based PNT solutions formed from frame TOA measurements. While the root cause of anomalous frame clock behavior remains a mystery, we emphasize that nearly half of the FAIs studied showed nominal frame clock jitter, with RMS values below 2.5 ns. Clearly, the Starlink v1.0 and v1.5 baseband frame assembly process, and at least a large fraction




Fig. 12: As the top panel from Fig. 8 but for a v2.0-Mini SV. The data shown are for STARLINK-30178 from signals captured in February 2024.

of its base oscillators, have sufficient stability to support high-accuracy pseudorange-based PNT, provided the 1-Hz frame clock adjustments could somehow be modeled or eliminated.

E. Starlink v2.0 Frame Clock

We also generated $\tilde{t}_r'$ time histories for v2.0-Mini SVs, an example of which is shown in Fig 12. These revealed frame clock adjustment behavior differing from that of v1.0 and v1.5 SVs in two ways: (1) the adjustments' magnitudes were generally much smaller, and (2) the adjustments occurred at irregular intervals, as opposed to the regular 1-Hz intervals for v1.0 and v1.5 SVs. Apparently, the Starlink v2.0-Mini SVs employ a different mechanism for base oscillator GNSS disciplining.

For v2.0-Mini SVs, the RMS values of the flattened time series $\tilde{t}_r''$ are higher than for the v1.0 and v1.5 SVs, ranging from 2.4 to 13.3 ns. There are three possible explanations for this: as compared to the v1.0 and v1.5 SVs, the v2.0-Mini SVs could have (1) oscillators with worse high-frequency stability, (2) a higher-noise GNSS disciplining technique, or (3) a frame clock that adds more high-frequency noise to the base oscillator signal t_d . Even still, the high-frequency jitter in the v2.0-Mini δt_f would not preclude meter-accurate pseudorange-based positioning because a batch of consecutive frame TOAs could be averaged to mitigate the higher jitter.

VIII. ABSOLUTE FRAME TIMING

Evidence of base clock GNSS disciplining raises the possibility that Starlink frames are continuously steered toward alignment with GPST. If true, this would be a notable discovery, inviting immediate use of Starlink for opportunistic pseudorange-based PNT. One could assume $\delta t_f(t) = 0$ for both assigned and side beams, which would be accurate to within the 1-Hz adjustment amplitude standard deviation, $\sigma_a = 117$ ns, or about 35 m in equivalent distance. Combining this satellite clock error model with publicly accessible Starlink ephemerides that are accurate (at least periodically) to within about $\sigma_e = 10$ m [45], one could form pseudoranges from a sequence of measured frame TOAs and process these with standard GNSS algorithms to obtain a position and time solution. Single-epoch 95% horizontal solution accuracy in this context can be approximated in meters as $\sqrt{4 \cdot \text{HDOP} (c^2 \sigma_a^2 + \sigma_e^2)}$

[15], where HDOP is the horizontal dilution of precision factor and c is the speed of light in vacuum. Given multiple signals from unique Starlink SVs with favorable geometry, HDOP could be as low as 0.55 [15], yielding a 95% horizontal error less than 54 m, tight enough for many applications of practical interest.

But if the frame clock is not transparent ($t_f \neq t_d$ in Fig. 3), then it may introduce errors that drive frame timing far from alignment with GPST despite the base clock being GNSS disciplined. This could render errors based on the simplistic clock model $\delta t_f(t) = 0$ too large for any useful PNT solution. A more sophisticated clock error model would be required.

This section addresses the question of whether Starlink frames are aligned with GPST. We first describe the methods used to deduce the time of transmission (TOT) in GPST of the m th frame of the l th FAI, written $t^*(l, m)$, given the corresponding frame TOA in GPST, $t_*(l, m)$. We then solve for $\delta t_f(l, m)$ and draw conclusions about Starlink frame clock steering and FAI alignment to GPST. For this initial study, it will suffice to determine $t^*(l, m)$ to better than $\sigma_a = 117$ ns.

A. Time of Transmission Calculation

TOT calculation begins with the frame TOA measurement $\tilde{t}_r(l, m)$ modeled in (20). We subtract from this an estimate of the offset $\delta t_r(l, m)$ obtained from the simultaneously captured GNSS signals, accounting for all cable delays, as described in Section II. This process allows us to determine $t_*(l, m)$ to within a few ns, which, in turn, is related to $t^*(l, m)$ by

$$t^*(l, m) = t_*(l, m) - \delta t_{\text{tof}}(l, m) \quad (23)$$

What remains is to calculate the frame's time of flight, $\delta t_{\text{tof}}(l, m)$, modeled as

$$\delta t_{\text{tof}}(l, m) = \frac{1}{c} \cdot \|\mathbf{r}_r(t_*(l, m)) - \mathbf{r}_t(t^*(l, m))\| + \delta t_{\text{atm}} \quad (24)$$

where $\mathbf{r}_r(t_*(l, m))$ is the receiver's location at the TOA, $\mathbf{r}_t(t^*(l, m))$ is the transmitter's location at the TOT, and δt_{atm} is the atmospheric (neutral and ionospheric) delay experienced by the signal over its path from transmitter to receiver. Substituting (23) into (24) yields the implicit relationship

$$\delta t_{\text{tof}}(l, m) = \frac{1}{c} \cdot \|\mathbf{r}_r(t_*(l, m)) - \mathbf{r}_t(t_*(l, m) - \delta t_{\text{tof}}(l, m))\| + \delta t_{\text{atm}} \quad (25)$$

from which $\delta t_{\text{tof}}(l, m)$ can readily be calculated numerically, provided δt_{atm} and smooth transmitter and receiver location functions $\mathbf{r}_t(t)$ and $\mathbf{r}_r(t)$.

To support determination of $t^*(l, m)$ to within $\sigma_a = 117$ ns, errors in $\mathbf{r}_t(t)$, $\mathbf{r}_r(t)$, and δt_{atm} must be small relative to this amount. For δt_{atm} , a Saastamoinen [46] neutral atmospheric model with average surface parameters was applied with Niell wet and dry mapping functions [47]; ionospheric delays, which are minimal at Ku-band, were ignored. Errors in δt_{atm} for all elevation angles are expected to be less than 20 ns. The receiver location $\mathbf{r}_r(t)$ was approximated as constant, i.e., $\mathbf{r}_r(t) = \mathbf{r}_r$. Although the true location, the antenna feedhorn's phase center, moves by up to one meter as the dish assembly rotates to track SVs, omission of this movement




Fig. 13: Deviation of true frame TOT (top) and true frame TOA (bottom) from t_f for subintervals of four representative FAIs from three different Starlink v2.0-Mini SVs. For visual clarity, each trace has been shifted vertically to begin at an offset of zero. Colors indicate corresponding FAIs. Numerical labels indicate the clock drift in parts per million.

from $\delta t_{\text{tof}}(l, m)$ calculations introduces negligible timing error compared to σ_a .

Determination of $\mathbf{r}_t(t)$ required more careful treatment. While public providers of orbital data such as Celestrak and SpaceTrack offer regularly updated ephemerides for all operational Starlink SVs, neither source reliably publishes these data continuously at the required precision. Most open-access ephemeris providers, including Celestrak, publish data as two-line elements, which typically exhibit km-level uncertainty at epoch. This is insufficient: km-level position uncertainty corresponds to μs -level timing uncertainty. SpaceTrack ephemerides offer meter-level positioning for all Starlink SVs; however, such precision is only available for a short time after epoch, with $1\text{-}\sigma$ position uncertainty climbing to 10 meters in less than one hour. This would have sufficed if the data refresh epochs were predictable so that signal captures could be set to coincide with ephemeris epochs. But SpaceTrack refreshes Starlink ephemerides up to thrice a day with epochs that vary randomly by several hours. As such, use of SpaceTrack datasets for $\mathbf{r}_t(t)$ was deemed impractical for the purposes of this paper.

Instead, we developed our own ephemeris models based on raw observables produced by the onboard GNSS receivers of a small number of v2.0-Mini SVs. These observables, provided by SpaceX and processed through our GRID receiver's central estimator as configured for LEO dynamics and a TCXO clock model, yielded discrete position solutions with formal errors below 10 m ($1\text{-}\sigma$). Position solutions were then fit over 12-13-min intervals with a 9th-order Chebyshev polynomial to yield a smooth $\mathbf{r}_t(t)$. Fit residuals were below 0.5 m RMS per dimension.

B. Clock Steering

We define clock steering as the process by which a clock is adjusted to match a reference time such that the clock offset does not drift significantly over relatively long ($> 10\text{ s}$) durations.




Fig. 14: Cost (on a dB scale) over the location search grid for ideal TOA clock steering. The left plot shows minimum-drift grid search results generated for the black line in the top panel of Fig. 13, whereas the right plot shows the same for the blue line. Both datasets originate from a capture taken from STARLINK-30178 in February 2024. Each plot's black circle represents the true receiver location. The black line shows the satellite ground track over the trace's time interval. Hot (yellow) areas signify areas experiencing degraded frame clock pre-compensation, while cold (blue) areas signify areas experiencing improved frame clock pre-compensation.

The top panel in Fig. 13 presents four frame clock deviation traces $-\delta t_f(l, m) = t^*(l, m) - t_f(l, m)$. These traces exhibit clock drift magnitudes varying from 0.2 to 21.4 ppm, values characteristic of a TCXO at best and a basic quartz oscillator at worst. The bottom panel shows frame clock deviations calculated from the same FAIs but substituting TOA in the place of TOT. We observe that all $t_* - t_f$ traces exhibit less deviation than their corresponding $t^* - t_f$ traces in the top plot. This suggests that Starlink SVs may be implementing frame clock pre-compensation to account for $\delta t_{\text{tof}}(l, m)$ changes as an SV passes overhead, thus mitigating the amount of frame clock drift experienced by Starlink receivers. Nonetheless, with long-term drift magnitudes ranging from 0.2 to 7.4 ppm, even the pre-compensated frame clock drift is far from ideal for pseudorange-based PNT.

It is worthwhile to consider the receiver location's impact on the effectiveness of frame clock pre-compensation. Pre-compensation cannot apply perfectly for all receivers in a given service cell, since the evolution of $\delta t_{\text{tof}}(l, m)$ changes with $\mathbf{r}_r$. Possibly, a Starlink SV pre-compensates such that receivers at the center of a service cell experience ideal clock steering ($t_* - t_f = 0$). To explore this possibility, we conducted a least-squares location search on a grid spanning 10° each of latitude and longitude constrained to the surface of the WGS84 ellipsoid for two separate FAIs, those corresponding to the black and blue traces of Fig. 13. For each gridpoint, a trace was constructed as

$$t_*(l, m) - t_f(l, m) = t^*(l, m) + \delta t_{\text{tof}}(l, m) - t_f(l, m) \quad (26)$$

where $\delta t_{\text{tof}}(l, m)$ is the time of flight from the transmitter to the gridpoint's location (rather than our receiver's). The gridpoint was then assigned a corresponding squared sum cost $J = \sum_{m \in \mathcal{M}_l} [t_*(l, m) - t_f(l, m)]^2$, where $\mathcal{M}_l$ represents all available occupied frame indices of the l th FAI. The results




Fig. 15: Normalized cross-correlation of received Starlink data against a local PSS + SSS replica showing the boundary between subsequent FAIs.

are shown in Fig. 14. Examination of the minimum-drift valley indicates that even in the best cases, $t_*(l, m) - t_f(l, m)$ traces still exhibited 0.1 ppm of gross clock drift. While this is certainly an improvement, it remains difficult to predict where the minimum-drift location will be. Consider that the two minimum-drift search results are generated for the same SV but for different FAIs, yet the minimum-drift valleys are far apart and differ in shape. Moreover, they appear unconnected with the location of the service cell or the SV's ground track. We conclude that Starlink frame timing is not driven towards alignment with GPST in a way that is both precise and predictable.

C. Fixed Assignment Interval Timing

The foregoing subsection showed that, during each FAI, Starlink frame timing is not driven toward alignment with GPST, whether at the SV or the target service cell. But it turns out that the frame clock does attempt approximate alignment with GPST at the beginning of each FAI. Fig. 15 shows the boundary between two successive FAIs; call these the $(l-1)$ th and l th FAIs. Note that frame occupancy in the visible segment of the $(l-1)$ th FAI appears maximal. Based on this, we assume that the final frame slot of the $(l-1)$ th FAI, whose index is $N_a - 1$, is occupied.

The l th FAI begins with the variable-length guard interval $T_{ag}(l) = N_{ag}(l)T_f$. Analysis of a dozen such boundaries revealed that $16 \leq N_{ag}(l) \leq 26$. We then calculated the true TOT for the final frame of the $(l-1)$ th FAI, denoted $t^*(l-1, N_a(l)-1)$, and estimated the true start time of the l th FAI as $t^*(l, 0) = t^*(l-1, N_a(l)-1) + T_f$. For the data shown in Fig. 15, $t^*(l, 0) = 514919.996874594$ when expressed in GPS seconds of week. Note that this is within about 3 ms of 514920, which is evenly divisible by 15. A similar pattern was evident in all such FAI transitions examined, from which we conclude that $t^*(l, 0)$ ranges between 4 and 2 ms before a 15-s GPST boundary. This finding may be expressed symbolically as

$$\lceil t^*(l, 0) \rceil - t^*(l, 0) \in [2, 4] \text{ ms} \quad (27)$$

$$\lceil t^*(l, 0) \rceil \bmod 15 = 0 \quad (28)$$

where $\lceil \cdot \rceil$ denotes the ceiling function and $l \in \mathbb{N}$.




Fig. 16: Top: Normalized cross-correlation of received Starlink data against a local PSS + SSS replica over 100 frame intervals, showing peak trains from a dominant (probably assigned) beam (orange) and from two side beams (green and black). Bottom: Corresponding measured frame TOA time histories $\tilde{t}_f$ with a linear trend removed from each. The data shown are for STARLINK-4577, a v1.5 SV, from signals captured in November 2023.

We also examined the true TOA $t_*(l, 0)$ and found that, while closer to an integer second, the range of variation was also approximately 2 ms, so it does not appear that FAI boundary timing is being steered into precise alignment with GPST at the location of the target service cell.

IX. COMPOSITE SIGNAL ANALYSIS

All frame timing results presented so far have been for dominant signals, typically from assigned beams. But deeper insight into the Starlink timing architecture can be gained by examining all constituent signals—both primary and secondary—of composite signals. One might be curious to know, for example, how closely secondary signals' frames are aligned with those of primary signals, whether they experience the same 1-Hz adjustments and excursions, and whether primary and secondary signals manifest similar levels of jitter. The answers to such questions would indicate whether the clock cascade in Fig. 4 should be thought of as beam-specific, phased-array-specific, or otherwise.

The composite signal shown in the top panel of Fig. 16 is an informative example in this regard. It consists of a dominant signal, marked in orange, and two clear secondary signals, marked in green and black. In the interval shown, the dominant signal experiences a large (150-ns) excursion, evident in the orange trace of the lower panel. Notably, one of the secondary signals exhibits nearly the same excursion whereas the other does not. What is more, just 250 ms before the interval shown there occurred a 1-Hz adjustment of the type discussed in Section VII-A. Different from the excursion, this adjustment affected all three signals identically. All composite signals we studied followed the same pattern: excursions and oscillations

TABLE I: Signal Statistics for Fig. 16

Signal	Frames	SNR (dB)	Offset (ns)	RMS (ns)	$\beta_f \times 10^5$
Orange	17	28	0	37.2	-1.2 ± 0.0014
Green	80	23	4.3	31.2	-1.2 ± 0.0006
Black	95	20	2.6×10^5	1.8	-1.1 ± 0.0005

may differ among beams of the same SV, but adjustments do not. This pattern implies a clock cascade model in which a single base oscillator underlies all transmissions from a given SV, but separate frame clocks drive disjoint subsets of the SV's beams.

For each signal in Fig. 16, Table I gives the number of frames whose TOAs were measured, the post-correlation SNR, the frame phase offset relative to the dominant signal's phase, the RMS of the detrended traces in the bottom panel of Fig. 16, and the frame rate offset parameter β_f . The latter is a measure of how compressed or stretched the signal's frame TOAs are compared to an exact $F_f = 750$ Hz cadence. The values for β_f are derived from a least-squares fit to the TOA time history for each signal. For the signals marked in orange and green, the excursion was removed before performing the least-squares fit. The uncertainty range shown for β_f is $\pm\sigma_f$, where σ_f^2 is the least-squares error variance for β_f assuming TOA measurement noise deviation $\sigma_w = 2$ ns.

The quantity β_f is the analog of the CFO parameter β but for the frame rate rather than the carrier frequency. For all traditional GNSS, $\beta_f = \beta$ for received signals (to within minor ionosphere-induced code-carrier divergence) because the same clock drives both carrier and spreading code generation. The local PSS + SSS replica in (18) also assumes identical β for both the carrier and the modulation $x_{m01}(t)$. But Table I identifies β_f as distinct from β because we discovered that $\beta_f = \beta$ does not hold universally for Starlink signals.

Table I indicates that the orange and green signals, which both experience the excursion, are also closely tied in other ways: their frame intervals are nearly perfectly aligned (within 4.3 ms), and they have identical β_f values. Contrarily, the excursion-free signal has a frame phase far from the other two (amounting to an offset of nearly 20% of the frame interval T_f) and a significantly different β_f value. For all three signals, the CFO parameter $\beta = (1.06 \pm 0.025) \times 10^{-5}$, where the uncertainty range comes from Fig. 6 for -10 dB pre-correlation SNR, a conservative bound for the signals in Fig. 16. Thus, in this case as in many others studied, β_f and β are significantly different—even different in sign. More generally, we found that $\beta_f = \beta$ only holds for about half of the FAIs studied, and that the condition $\beta_f = \beta$ can change at FAI boundaries for the same SV. This supports a clock cascade model in which $\Delta t_f(t)$ and $\Delta t_c(t)$ in (1) and (2) are independent.

X. CONCLUSIONS

Based on our study of Starlink v1.0, v1.5, and v2.0-Mini frame timing, we draw the following conclusions (C) and make the following conjectures (G) about Starlink's timing architecture.

- C1 Beam-to-cell assignments remain static over 15-s fixed assignment intervals (FAIs). Each downlink beam carries a sequence of frames containing user data.
- C2 All beams cast by a given SV carry signals that are ultimately driven by the same base oscillator, but the signals may or may not share the same frame clock, which means that frame phase, rate, and short-term stability can differ significantly from beam to beam.
- C3 Periodic adjustments of the base oscillator cause large (~ 120 ns for v1.0 and v1.5 SVs, ~ 10 ns for v2.0-Mini SVs) abrupt deviations in frame timing that are common across all beams of the same SV.
- C4 Frame TOA jitter is nominally small (below 2.5 ns RMS), but frame timing can episodically manifest large unpredictable oscillations and excursions.
- C5 Frame TOA behavior can change on FAI boundaries, even for the same SV.
- C6 Frame short-term timing stability is ultimately limited by the quality of the base oscillator, which may be as good as, but no better than, a TCXO.
- C7 Frame timing is not driven toward alignment with GPST, whether at the SV or the target service cell. Rather, frames as transmitted can have a rate that is significantly different (e.g., 20 ppm) from the nominal $F_f = 750$ Hz rate.
- C8 FAIs are aligned within a few ms to a 15-s boundary of GPS seconds of week.
- C9 The carrier phase offset parameter β and frame rate offset parameter β_f may differ significantly, unlike with traditional GNSS signals, for which $\beta_f = \beta$ holds universally (to within minor ionosphere-induced code-carrier divergence).
- G1 The periodic adjustments seen in frame timing are the result of GNSS disciplining of the base oscillator.
- G2 The carrier signal is driven directly by the base oscillator, meaning that the carrier clock in Fig. 4 is transparent. By contrast, frames are assembled at baseband and then mixed to RF and modulated onto the carrier by a process that is subject to the variations (including those due to GNSS disciplining) of the base oscillator but that may impose additional frame timing and rate variations.
- G3 Signals having nearly identical frame rate offset parameter β_f and frame phase offset share the same frame clock.
- G4 Signals sharing the same frame clock are carried by beams emanating from the same downlink phased array, of which each Starlink SV has three.
- G5 Abnormal frame clock behaviors (oscillations, excursions) are not intentional. They are the result of frame packaging and modulation algorithms that satisfy communications requirements but take no special care to ensure frame timing stability.
- G6 The frame packaging and modulation algorithms are software-defined and could be updated to eliminate excursions and oscillations.

Three key implications (I) for pseudorange-based positioning and timing based on Starlink follow from the foregoing conclusions and conjectures:

- I1 Large and apparently unpredictable variations in frame timing, which can differ from beam to beam, and the lack of fine-grained frame phase steering toward a common time standard such as GPST, make Starlink currently unsuitable for purely opportunistic pseudorange-based PNT despite public availability of SV ephemerides.
- I2 Current Starlink SV hardware appears fundamentally capable of supporting pseudorange-based PNT with position and timing accuracy exceeding traditional GNSS.
- I3 There exist three options for making Starlink suitable for high-accuracy pseudorange-based PNT: (1) SpaceX could update Starlink's software-defined frame clock to eliminate abnormal oscillations and excursions and to steer frames into precise (ns-level) alignment with GPST at the transmitter; (2) SpaceX could broadcast a low-latency frame clock model $\delta t_f(t)$ that would allow users to compensate for frame clock variations; (3) a third party could establish a network of reference stations that would measure Starlink frame TOA and send a low-latency frame clock model $\delta t_f(t)$ to its subscribers.

We note that option (2) could be limited to Starlink subscribers, whereas (1) and (3) would allow non-subscribers to benefit from Starlink-based PNT. A subscriber-limited version of (2) would involve bi-directional communication (at least for subscriber authentication), whereas a subscription-free version would enable passive (radio silent) Starlink-powered PNT. Option (3) would require a reference network dense enough to obtain accurate models $\delta t_f(t)$ for all unique frame clocks. In the worst case, this would require one reference receiver per ~ 20 -km-diameter service cell.

ACKNOWLEDGMENTS

Research was supported by the U.S. Department of Transportation under Grant 69A3552348327 for the CARMEN+ University Transportation Center, by Sandia National Laboratories, and by affiliates of the 6G@UT center within the Wireless Networking and Communications Group at The University of Texas at Austin. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. The United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes.

REFERENCES

- [1] T. E. Humphreys, "Interference," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 469–503.
- [2] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, M. L. Psiaki, and T. E. Humphreys, "First results from three years of GNSS interference monitoring from low Earth orbit," *NAVIGATION*, vol. 68, no. 4, pp. 673–685, 2021.
- [3] Z. Clements, P. Ellis, and T. E. Humphreys, "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit," in *Proceedings of the IEEE/ION PLANS Meeting*, Monterey, CA, 2023, pp. 458–469.
- [4] G. S. Workgroup, "GPS spoofing: Final report of the GPS spoofing workgroup," OPSGROUP, Tech. Rep., 2024. [Online]. Available: <https://ops.group/blog/gps-spoofing-final-report>
- [5] W. Stock, R. T. Schwarz, C. A. Hofmann, and A. Knopp, "Survey on opportunistic PNT with signals from LEO communication satellites," *IEEE Communications Surveys & Tutorials*, 2024.
- [6] P. A. Iannucci and T. E. Humphreys, "Fused low-Earth-orbit GNSS," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2022.
- [7] P. Gomez-del Hoyo and P. Samczynski, "Starlink-based passive radar for Earth's surface imaging: first experimental results," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.
- [8] Z. M. Kassas, N. Khairallah, and S. Kozhaya, "Ad astra: Simultaneous tracking and navigation with megaconstellation LEO satellites," *IEEE Aerospace and Electronic Systems Magazine*, 2024.
- [9] M. Neinavaie and Z. M. Kassas, "Cognitive sensing and navigation with unknown OFDM signals with application to terrestrial 5G and Starlink LEO satellites," *IEEE Journal on Selected Areas in Communications*, 2023.
- [10] N. Jardak and R. Adam, "Practical use of Starlink downlink tones for positioning," *Sensors*, vol. 23, no. 6, p. 3234, 2023.
- [11] C. Yang and A. Soloviev, "Starlink Doppler and Doppler rate estimation via coherent combining of multiple tones for opportunistic positioning," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, 2023, pp. 1143–1153.
- [12] S. Kozhaya and Z. M. Kassas, "On the fundamental tracking performance and design considerations of radio navigation," *IEEE Journal on Selected Areas in Communications*, 2024.
- [13] M. L. Psiaki, "Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids," *NAVIGATION*, vol. 68, no. 3, pp. 621–641, 2021.
- [14] M. O. Moore, M. L. Psiaki, and R. M. Buehrer, "Time-diverse Doppler-only LEO PNT: Initial solution," in *Proceedings of the ION GNSS+ Meeting*, 2024, pp. 3532–3541.
- [15] P. A. Iannucci and T. E. Humphreys, "Economical fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, 2020.
- [16] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–16, 2023.
- [17] Z. M. Komodromos, W. Qin, and T. E. Humphreys, "Weak signal acquisition and tracking of the Starlink Ku-Band downlink to enable global PNT," in *Proceedings of the ION Joint Navigation Conference (JNC)*, 2024.
- [18] S. Kozhaya, J. Saroufim, and Z. M. Kassas, "Starlink for PNT: A trick or a treat?" in *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, 2024.
- [19] 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.
- [20] K. Wang and A. El-Mowafy, "LEO satellite clock analysis and prediction for positioning applications," *Geo-spatial Information Science*, vol. 25, no. 1, pp. 14–33, 2022. [Online]. Available: <https://doi.org/10.1080/10095020.2021.1917310>
- [21] H. Ge, T. Wu, and B. Li, "Characteristics analysis and prediction of low earth orbit (LEO) satellite clock corrections by using least-squares harmonic estimation," *GPS Solutions*, vol. 27, no. 1, p. 38, 2023.
- [22] C. Jiang, Z. Luo, M. Zhu, M. Guan, H. Zhu, and M. Gao, "Characteristics analysis of LEO satellite clock corrections and prediction by sliding estimation," *Available at SSRN 4818891*, 2024.
- [23] 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.
- [24] H. A. Nichols, M. J. Murrian, and T. E. Humphreys, "Software-defined GNSS is ready for launch," in *Proceedings of the ION GNSS+ Meeting*, Denver, CO, 2022.
- [25] Z. Clements, P. A. Iannucci, T. E. Humphreys, and T. Pany, "Optimized bit-packing for bit-wise software-defined GNSS radio," in *Proceedings of the ION GNSS+ Meeting*, St. Louis, MO, 2021, pp. 3749–3771.

- [26] T. Pany, D. Akos, J. Arribas, M. Z. H. Bhuiyan, P. Closas, F. Dovis, I. Fernandez-Hernandez, C. Fernández-Prades, S. Gunawardena, T. Humphreys, Z. M. Kassas, J. A. L. Salcedo, M. Nicola, M. L. Psiaki, A. Rügamer, Y.-J. Song, and J.-H. Won, "GNSS software defined radio: History, current developments, and standardization efforts," *NAVIGATION*, vol. 71, no. 1, 2024.
- [27] S. Pekhterev, "The bandwidth of the Starlink constellation and the assessment of its potential subscriber base in the USA," *SatMagazine*, vol. 11, pp. 54–57, 2021.
- [28] G. Huston, "A transport protocol's view of Starlink," May 2024, accessed: 2024-11-22. [Online]. Available: <https://blog.apnic.net/2024/05/17/a-transport-protocols-view-of-starlink/>
- [29] N. Mohan, A. E. Ferguson, H. Cech, R. Bose, P. R. Renatin, M. K. Marina, and J. Ott, "A multifaceted look at Starlink performance," in *Proceedings of the ACM on Web Conference 2024*, 2024, pp. 2723–2734.
- [30] E. Grayver, R. Nelson, E. McDonald, E. Sorensen, and S. Romano, "Position and navigation using Starlink," in *2024 IEEE Aerospace Conference*. IEEE, 2024, pp. 1–12.
- [31] Z. M. Komodromos, W. Qin, and T. E. Humphreys, "Signal simulator for Starlink Ku-Band downlink," in *Proceedings of the ION GNSS+ Meeting*, 2023, pp. 2798–2812.
- [32] A. Hills, J. M. Peha, J. Munk, and S. Pogorelc, "Controlling antenna sidelobe radiation to mitigate Ku-band LEO-to-GEO satellite interference," *IEEE Access*, vol. 11, pp. 71 154–71 163, 2023.
- [33] A. Hills, J. M. Peha, and J. Munk, "Feasibility of using beam steering to mitigate Ku-band LEO-to-GEO interference," *IEEE Access*, vol. 10, pp. 74 023–74 032, 2022.
- [34] J. Brodtkin. (2024, October) SpaceX claims Starlink can offer gigabit speeds if FCC approves new plan. [Online]. Available: <https://arstechnica.com/tech-policy/2024/10/spacex-claims-starlink-can-offer-gigabit-speeds-if-fcc-approves-new-plan/>
- [35] D. Odijk, "Positioning model," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 605–638.
- [36] O. Montenbruck, A. Hauschild, and P. Steigenberger, "Differential code bias estimation using multi-GNSS observations and global ionosphere maps," *NAVIGATION*, vol. 61, no. 3, pp. 191–201, 2014.
- [37] M. Meurer and F. Antreich, "Signals and modulation," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 91–119.
- [38] C. Jekeli and O. Montenbruck, "Time and reference systems," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 25–58.
- [39] R. Beard and K. Senior, "Clocks," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 121–164.
- [40] M. Skolnik, *Introduction to radar systems 2nd Edition*. McGraw-Hill, 1980.
- [41] D. Odijk and L. Wanninger, "Differential positioning," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 753–780.
- [42] P. Royston, "Approximating the Shapiro-Wilk W-test for non-normality," *Statistics and computing*, vol. 2, pp. 117–119, 1992.
- [43] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2325–2383, 1998.
- [44] C. Hackman and T. E. Parker, "Noise analysis of unevenly spaced time series data," *Metrologia*, vol. 33, no. 5, p. 457, 1996.
- [45] A. Liu, X. Xu, Y. Xiong, and S. Yu, "Maneuver strategies of Starlink satellite based on SpaceX-released ephemeris," *Advances in Space Research*, 2024.
- [46] J. Saastamoinen, "Atmospheric correction for the troposphere and stratosphere in radio ranging of satellites," in *Geophysical Monograph 15*, S. W. Henriksen, Ed. Washington, D.C.: American Geophysical Union, 1972, pp. 247–251.
- [47] A. E. Niell, "Global mapping functions for the atmosphere delay at radio wavelengths," *Journal of Geophysical Research*, vol. 101, pp. 3227–3246, 1996.

Signal Parameter Estimation and Demodulation of the OneWeb Ku-Band Downlink

Zacharias M. Komodromos*, Todd E. Humphreys†

**Department of Electrical and Computer Engineering, The University of Texas at Austin*

†*Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin*

Abstract—We present techniques for estimating key parameters of OneWeb’s Ku-band downlink signal (10.7–12.7 GHz) and reveal it as a single-carrier QPSK signal with a 230.4 MHz symbol rate. The techniques also estimate the signal’s roll-off factor and center frequencies. We further provide the first published account of OneWeb signal demodulation, revealing the basic frame structure of the downlink signal, including a synchronization sequence that repeats every millisecond and is common across all beams, channels, and satellites. Identifying this sequence enables making time-of-arrival measurements from OneWeb signals. These findings contribute to the growing body of research focused on repurposing low-Earth-orbit satellite communication signals for positioning, navigation, and timing.

Index Terms—OneWeb, signal identification, signal processing, positioning, low Earth orbit

I. INTRODUCTION

The rapid expansion of low-Earth-orbit (LEO) mega-constellations has enabled unprecedented global broadband coverage. These mega-constellations could also revolutionize the positioning, navigation, and timing (PNT) landscape by serving as a new source of signals for PNT [1]–[7]. LEO signals offer increased satellite visibility, enhanced geometric diversity, and improved robustness through a greater variety of signals. Compared to traditional Global Navigation Satellite System (GNSS) signals, LEO communication signals provide significantly higher bandwidths and received power. These advantages make repurposing LEO signals as a complement or backup to GNSS particularly appealing, especially amid rising threats of jamming and spoofing [8]–[12].

With over 600 LEO satellites, OneWeb is the second-largest LEO constellation after SpaceX’s Starlink [13]. Its polar orbits provide global coverage and are 15 times closer to Earth than GPS orbits. Moreover, its space-to-Earth link operates in the accessible 10.7–12.75 GHz band.

Both OneWeb and Starlink have been the focus of recent studies exploring their potential for Doppler and carrier-phase-based positioning [14]–[22]. But compared to pseudorange-based PNT techniques, Doppler-based techniques have worse timing accuracy by many orders of magnitude (milliseconds vs. nanoseconds), even under optimistic measurement noise and satellite clock offset rate assumptions [3], [23], [24]. Recognizing that many PNT applications of practical interest require accurate timing, we seek a characterization of OneWeb’s signal structure sufficient to enable pseudorange measurements, just as [5] provided for Starlink.

Prior research has identified key characteristics of OneWeb’s signals, including approximate center frequencies, channel bandwidths [25]–[28], and a 10-ms periodicity believed to be unique to each satellite beam [14], [25]. But beyond these basic observations, OneWeb’s signal remains undisclosed and unpublished, unlike Starlink’s signal, whose structure and parameters are documented in [5]. Indeed, the existing literature does not even identify OneWeb’s modulation scheme. In [25], the authors blindly detect the presence of modulation that repeats within the same satellite beam, but no demodulation is attempted. Furthermore, OneWeb’s Ku-band signal parameters have not been rigorously estimated, nor has any demodulated data been presented as in [5]. A detailed understanding of the signal structure, parameters, and synchronization sequences would enhance OneWeb’s potential as a PNT source and provide an analytical basis for assessing its limitations. This information is crucial to achieving our primary goal of harnessing OneWeb for positioning and timing through pseudorange-based methods.

We present a signal model for OneWeb’s Ku-band downlink, incorporating the identified modulation scheme and accounting for carrier frequency offset (CFO). Using established methods, we estimate the symbol rate and pulse shape roll-off, which are essential for demodulation. A typical Quadrature Phase Shift Keying (QPSK) demodulation processing chain is then applied to the signal. Analysis of the demodulated signal reveals that the 10-ms periodicity observed in prior studies may be merely a temporary consequence of the present low demand for OneWeb-provided data. We also identify a short synchronization sequence that repeats every 1 ms and is present on all satellites and beams.

To summarize, our paper offers three primary contributions. First, it provides a signal model for OneWeb’s Ku-band downlink and demonstrates how to estimate its key parameters. This process applies to all Gen 1 OneWeb satellites, and likely to future versions as long as the signal remains single-carrier. Second, it identifies a synchronization sequence obtained by demodulating data from multiple satellites. The demodulation process further reveals that the 10-ms periodicity observed in prior studies is unlikely to be a permanent feature, as the data transmitted every 10 ms appear to be repeated default data. Finally, the paper shows how a local replica of the repeating modulation can be used to generate time-of-arrival (TOA) and Doppler measurements.




Fig. 1: Block diagram of the OneWeb signal capture process.

II. SIGNAL CAPTURE

OneWeb’s roughly 650 satellites follow polar orbits at an altitude of 1200 km. Their downlink signal has approval for the 10.7–12.7 GHz space-to-Earth frequency band. A variety of commercially available user terminals (UTs) are available for OneWeb, with dual parabolic or flat panel designs. These would be ill-suited for our signal capture needs, as they do not provide access to the raw signal samples. Further, we would need to investigate the quality of the UT’s internal clock used for downconversion and sampling or risk it having it skew any subsequent signal time stability analysis. Nonetheless, we can glean useful information from public OneWeb UT specifications, such as their being designed to support QPSK, 8PSK, and 16APSK modulation with a 250-MHz-wide channel spacing [29].

In light of the challenges of using a commercial UT, we opted to employ our own steerable 90-cm offset parabolic dish with an approximately 3-degree beamwidth. Using publicly available ephemerides from North American Aerospace Defense Command (NORAD) in the form of Two-Line Elements (TLEs), we can steer the dish to track OneWeb satellites overhead. Our antenna’s narrow beamwidth limits captures to a single satellite at a time. As we show later on, the satellites have a fixed beam, where the beam footprint moves with the satellite, as opposed to a fixed-cell approach like Starlink’s, where the beam is fixed to a cell on the ground.

Fig. 1 outlines the hardware used to capture the raw IQ samples. Our parabolic dish is equipped with a feedhorn connected to a low-noise block (LNB) with a conversion gain of 60 dB and a noise figure of 0.8 dB. The LNB downconverts 10.7–11.7 GHz signals to 950–1950 MHz, or 11.7–12.75 GHz to 1100–2150 MHz. The antenna’s nominal gain is 40 dBi at 12.5 GHz, but suffers at least 4–5 dB of losses due to lack of circular-to-linear polarizer and feedhorn misalignment. The antenna is located on the campus of The University of Texas at Austin, with a clear view of the sky.

The signal capture system allows selection between two capture modes: wide and narrow. The wide capture is a fixed 500 Msps, while the narrow capture is variable from 250 Msps and below. The downstream hardware then performs

additional downmixing, bandpass filtering, and 16-bit complex sampling. Both the LNB and downstream downconversion and sampling hardware are locked to the same GPS-disciplined oven-controlled crystal oscillator (OCXO).

The usable bandwidth of the narrow capture varies depending on the sample rate, with a maximum of 200 MHz. The usable bandwidth of the wide capture is roughly 400 MHz. The capture system is capable of capturing on two channels at once, with the limitation that the sampling rate must be identical for the sampling to begin simultaneously. Such a setup is referred to as dual capture.

For single-carrier signals like OneWeb’s, the narrow captures are challenging or impossible to use for data recovery, but are useful for observing patterns in the signal structure. The wide captures produce roughly 2 GB/s of data, and are prone to overflow the host or device-side data buffers resulting in dropped samples. As such, our hardware permits less than 10 seconds of continuous capture in that mode. For narrow captures, the device is capable of capturing for the duration a satellite is overhead.

III. SIGNAL MODEL

One might expect OneWeb to use orthogonal frequency division multiplexing (OFDM) given its proven success in space-to-Earth communications by Starlink, and its domination in wireless communications. Various sources, including a test report for one of the UTs [29], suggest OneWeb instead uses a single-carrier signal. Assuming the 250 MHz channel spacing specified in [29] for the Ku band, OneWeb could presumably have 8 channels over the 2 GHz Ku-band, as assumed in [25].

We claim, and later provide evidence, that OneWeb employs a Multi-Frequency Time Division Multiple Access (MF-TDMA) Single Carrier (SC) scheme, on each of 8 separate frequency channels. The sections below outline the signal model for a SC QPSK signal, and subsequently the model of the received OneWeb signal.

A. QPSK Signal Model

Fundamentally, a SC signal is a train of pulses, each modulated by a phase and amplitude shift corresponding to

the data symbol being transmitted on that pulse. The resulting continuous stream of symbols can be characterized by a simple signal model. The baseband signal model for a SC QPSK signal is given by

$$s(t) = \sum_m \exp(j\pi a_m/2) p\left(\frac{t - mT_{\text{sym}}}{T_{\text{sym}}}\right) \quad (1)$$

where statistically independent symbol phases $a_m \in \mathbb{N}_4 \triangleq \{0, 1, 2, 3\}$ for $m \in \mathbb{Z}$ can be encoded to represent two bits per symbol, and T_{sym} is the symbol period. The symbol phase values $a_m \in \mathbb{N}_4$ will be referred to as symbols hereafter.

The pulse shaping function $p(t)$ is left unspecified in (1). Our later signal parameter estimation shows that a square root raised cosine (SRRC) pulse nicely fits the OneWeb signals captured. Pulse shaping is essential to limit the excess bandwidth of the signal, for regulatory and interference purposes. Proper filtering of a pulse-shaped signal within the receiver can also reduce inter-symbol interference (ISI). Due to its popularity, we tested against the SRRC. The only parameter of note for the SRRC is its rolloff factor β_r , which is a measure of its excess bandwidth.

In a single-carrier MF-TDMA protocol, subsequences of symbols are structured in a hierarchy of slots, frames, blocks, etc. We will use the term slot to describe the smallest grouping of symbols that is self-contained in the sense that it includes one or more predictable symbol sequences that mark the beginning of a slot and allow synchronization to it. A slot contains data destined for a small number of users—typically a single user. A contiguous set of slots whose data appear to be correlated in some way will be called a frame.

Different standards outline how frames are constructed. DVB-S2, a popular standard for satellite digital broadcasting, defines base band (BB) frames constructed from a number of input streams time-multiplexed into a physical layer (PL) frame. If a given stream has no data, a default physical-layer frame is inserted, or if a BB frame is incomplete, it is padded. Going from a BB frame to a PL frame also involves appending parity bits and modulating the data based on some modulation and coding scheme. While this processing may provide some structure to the symbol stream, the signal as observed by a receiver will nevertheless follow the model in (1).

B. Received Signal Model

The transmitted signal passes through the LEO-to-Earth channel and through the receiver front-end and discretization process. It is affected by multipath fading, Doppler, delay, filtering, digitization, and noise. SATCOM studies conducted in the Ku-band with measurements filtered to 80 MHz indicate that delay spread is minor for receivers experiencing only light shadowing and not near other objects [30]. We will treat the effect of delay spread as negligible, since in our case the open sky view provides a strong line-of-sight (LOS) component, with few if any multipath components entering our antenna's narrow beam. Another study on dispersive delays in the Ku-band with a 200-MHz bandwidth receiver attributes a majority of the delay to atmospheric dispersion, and shows sub-millimeter delay [31]. Due to these findings, we adopt a

simple additive white Gaussian noise (AWGN) model for the received signal's noise.

The most pressing phenomenon to model is the Doppler effect. For appropriately low bandwidths and time durations, one often assumes a narrowband model, which treats the Doppler effect as a simple frequency shift [32]. Using definitions from [5], let v_{los} be the magnitude of the LOS velocity between the satellite and a receiver, and $\beta \triangleq v_{\text{los}}/c$ be the carrier frequency offset (CFO) parameter, where c is the speed of light. The narrowband model requires that $\beta F_{\text{sym}} T_{\text{synch}} \ll 1$ for T_{synch} some interval over which we expect to maintain time synchronization, and for the symbol rate $F_{\text{sym}} = 1/T_{\text{sym}}$. For a LEO satellite with a large bandwidth like Starlink or OneWeb, the requirement is violated [5], [33], [34]. Thus, a wideband model that accounts for Doppler compression/dilation of the modulation, manifesting as time scaling, is needed in addition to frequency shifting. Given this, our received baseband analog model is

$$y_a(t) = s((t - \tau_0)(1 - \beta)) \times \exp(j2\pi[F_c(1 - \beta) - F_{\text{cr}}](t - \tau_0)) + w(t) \quad (2)$$

where F_c is the center frequency of the OneWeb channel, F_{cr} is the center frequency to which the receiver is tuned, τ_0 is the delay experienced by the signal along the least-time path from transmitter to receiver, and $w(t)$ is complex-valued zero-mean additive white Gaussian noise (AWGN) whose in-phase and quadrature components each have (two-sided) spectral density $N_0/2$. It is important to note here that since β depends on the LOS velocity, it is time varying. Also, the satellite clock frequency error causes the same effect as LOS motion on CFO. To keep the model simple, we will lump the effects of satellite clock frequency error into the CFO parameter β unless otherwise indicated.

The final stage is discretization. The received signal passed through a low-pass filter $h(t)$ with bandwidth F_h and sampled at a rate $F_{\text{sr}} > F_h$. As mentioned earlier, the useful bandwidth (3 dB) of our captured signals is roughly $F_h = 200$ MHz for the narrow capture, and $F_h = 400$ MHz for the wide capture. The signal is then quantized to 16 bits, and the resulting discrete-time signal is

$$y(n) = \int_{-\infty}^{\infty} h(n/F_{\text{sr}} - \tau) y_a(\tau) d\tau, \quad n \in \mathbb{Z} \quad (3)$$

At some stages of our processing, we further digitally low-pass filter the signal closer to the symbol rate F_{sym} for the wideband captures.

IV. SIGNAL PARAMETER ESTIMATION

Three parameters must be known to reliably demodulate the signal. The symbol rate F_{sym} is most important, since any error in its estimate would strain a symbol timing recovery loop, leading to possible data loss.

Another important parameter is the rolloff factor β_r of the pulse shaping function. A receiver incorporates a matched filter to maximize the SNR, where the filter's impulse response is the time-reversed pulse shape. If the rolloff factor is not accurately estimated, the matched filter will be suboptimal,

leading to increased ISI. In practice, our low-multipath environment and frequency-flat (non-dispersive) channel render the effect of ISI negligible, thus errors in β_r will not significantly affect our demodulation.

The final parameter to estimate is the center frequency F_c of each channel. It is also not as consequential as F_{sym} since, from a receiver's perspective, there are various techniques to blindly estimate a QPSK signal's frequency offset. Estimators often lump together all the frequency shifting effects into a single value, corresponding to the exponent in (2), thereby compensating for Doppler before demodulation.

A. Exploiting Signal Cyclostationarity

To estimate the symbol rate, we exploit the cyclostationarity of the QPSK signal. There are numerous examples in the literature of exploiting cyclostationarity to identify periodicity in signals [5], [35], [36]. Specifically, the cyclic autocorrelation (CA) function of a QPSK signal reveals its symbol rate [37]–[39].

The goal of CA analysis is to identify the hidden second-order periodicity in the signal. As an example, consider a simple sinusoid centered at f_c with additive noise. If the noise power is sufficiently high, it can be difficult in the time domain to identify whether the signal is present, let alone its periodicity. Yet observing the sinusoid's spectrum using a Fourier transform can easily reveal the hidden first-order periodicity as impulses at $\pm f_c$. Likewise, the CA function reveals the hidden second-order periodicity in an SC QPSK signal.

Let $y(n)$ be the received signal described in (3). For now, assume $\beta = 0$ and $F_{\text{cr}} = F_c$. The autocorrelation function is defined as

$$R_y(n, k) = \mathbb{E}[y(n+k)y^*(n)] \quad (4)$$

where $n, k \in \mathbb{Z}$, $\mathbb{E}[\cdot]$ is the expectation operation, and $y^*(n)$ is the complex conjugate of $y(n)$. Since $s(t)$ in (1) is a pulse train with statistically independent phase shifts a_m , $m \in \mathbb{Z}$, it follows that $R_y(n, k)$ is periodic in n for certain values of k with a period equal to $T_{\text{sym}}F_{\text{sr}}$. This makes $y(n)$ cyclostationary and allows for its autocorrelation function to be expanded in a Fourier series as

$$R_y(n, k) = \sum_{\alpha \in \mathcal{A}(\xi)} R_y^\alpha(k) \exp(j2\pi\alpha n) \quad (5)$$

where $\mathcal{A}(\xi) = \{q/\xi : q \in \mathbb{Z}\}$. The particular set $\mathcal{A}(T_{\text{sym}}F_{\text{sr}})$ contains the so-called cyclic frequencies. The Fourier coefficient $R_y^\alpha(k)$ constitutes the signal's CA function.

Following the framework from [40] for discrete-time signals, we approximate $R_y^\alpha(k)$ as

$$\hat{R}_y^\alpha(k) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} y(n+k)y^*(n) \exp(-j2\pi\alpha n) \quad (6)$$

When operating on a finite-length signal, N is assumed to be much larger than the cyclic period of the signal, and the signal is assumed to be long enough that $y(N-1+k)$ is defined for all k values of interest.




Fig. 2: Cyclic autocorrelation function of four captures with $k = 0$. The horizontal axis coordinate is $\tilde{\alpha} = \alpha F_{\text{sr}}$. Peaks range from 230.399875 to 230.400180 MHz.

B. Estimation of F_{sym}

We can now define an estimator for the symbol rate as

$$\hat{F}_{\text{sym}} = \arg\max_{\tilde{\alpha} \in \mathcal{S}} \tilde{R}_y^\alpha(0) \quad (7)$$

where $\mathcal{S}$ is some constrained set of cyclic frequencies around the expected symbol rate, and $\tilde{\alpha} = \alpha F_{\text{sr}}$ is scaled such that the search space is in Hertz. From *a priori* knowledge (e.g., observation of the signal's spectrum and perusal of regulatory or commercial documentation), we can devise a narrow range of possible symbol rates over which to search.

In practice, our simplifying assumptions $F_{\text{cr}} = F_c$ and $\beta = 0$ are not accurate, and this must be addressed. Note from the received baseband analog model in (2) that the exponential can be non-unitary when $\beta \neq 0$ or $F_{\text{cr}} \neq F_c$, but this will not affect the periodicity in the modulation and so will not affect $\tilde{R}_y^\alpha(k)$ nor $\hat{F}_{\text{sym}}$. On the other hand, for a nonzero β the modulation $s(t - t_0)$ gets compressed or stretched as $s((t - t_0)(1 - \beta))$. This CFO time-scaling effect does change the periodicity—and therefore the observed symbol rate—of the received modulation.

To address this, we correct for the largest component of time scaling—the one due to LOS motion between the satellite and receiver—before performing symbol rate estimation. From the TLE satellite ephemerides and the known receiver location we calculate the LOS velocity v_{los} , from which we estimate the motion-induced component of the CFO parameter as $\hat{\beta} = v_{\text{los}}/c$. We then resample the received signal at $F_{\text{sr}}(1 - \hat{\beta})$ to undo the motion-induced time scaling. This results in at most a few Hertz of error in the Doppler shift, and similarly negligible error in the time scaling. Note that this process does not compensate for CFO introduced by transmitter clock frequency error, which, as will be shown, leads to slight residual errors in the symbol rate estimate.

After correcting for time scaling as described, we estimated F_{sym} based on four captures from March and June 2024. Fig. 2 shows the CA function (6), with N equal to the number of samples in 1 ms of data and $k = 0$, for each of the four captures. Peaks range from 230.399875 to 230.400180 MHz. Using only 1 ms of data instead of an infinitely long signal causes the peaks to present with finite width. The variance in the location of the peaks is due primarily to transmitter clock frequency error. The peaks are centered around 230.4 MHz to within about 200 Hz, from which we conclude that $F_{\text{sym}} = 230.4$ MHz.

C. Resampling

Resampling wide-capture signals at an integer multiple of F_{sym} facilitates symbol synchronization and demodulation. Recall that for wide-capture signals, the original signal is lowpass filtered to $F_h < F_{\text{sr}}$ and sampled at $F_{\text{sr}} > F_{\text{sym}}$. Resampling to some $F_s < F_{\text{sr}}$ proceeds by first applying a polyphase anti-aliasing filter at a new lower $F_h \geq F_{\text{sym}}$ followed by uniform resampling at F_s . The useful frequency content of the signal is reduced due to the filtering, but what is removed is either noise or adjacent channels which are not of interest. For what follows, we resample wide captures at twice the symbol rate, going from $F_{\text{sr}} = 500$ Msp/s to $F_s = 2F_{\text{sym}} = 460.8$ Msp/s.

Narrow-capture signals are only used in this paper for long-duration correlation analysis for which no resampling was required.

D. Symbol and Carrier Frequency Synchronization

With F_{sym} known, one can proceed with carrier frequency and symbol synchronization. This is done with the resampled wide-capture signal. Carrier frequency synchronization proceeds in two steps, coarse CFO compensation followed by phase tracking. Coarse compensation can adopt the method from Sec. IV-B in which TLEs are used to wipe off the expected motion-induced frequency shift. But this method does not compensate for any offset $F_{\text{cr}} - F_c$ from the receiver's center frequency to the true channel center frequency, and ignoring this offset may strain the carrier phase tracking loop. Thus, a method such as [41] is preferable for coarse CFO compensation. If, as demand increases, OneWeb adopts M -PSK with $M > 4$, one could extend the work in [42], or rely on the method introduced in [43].

After coarse CFO compensation, inspection of the samples rendered on the complex plane reveals a QPSK-like constellation, but with samples between the canonical symbol locations due to the non-unitary number of samples per symbol, and with residual constellation rotation over time due to imperfect frequency synchronization. To refine the frequency synchronization and thereby arrest the rotation, we apply a second-order phase-locked loop.

Having achieved frequency synchronization, we proceed to symbol synchronization. For this, we pass the signal through a symbol timing recovery loop to align to the pulse apex, then resample to one sample per symbol. We adopted decision-directed approach to symbol synchronization using the Mueller-Muller timing error detection method [44, Chapter 8]. A preliminary decision on the m th symbol can then be made on each pulse-centered sample. The entire process is summarized by the block diagram in Fig. 3.

Symbol demodulation is preliminary at this stage because it neglects matched filtering to the pulse $p(t)$, which cannot yet be applied because the rolloff factor β_r remains unknown. Once the rolloff factor is estimated using the preliminary symbol estimates, as described in the following subsection, a second pass of the process in Fig. 3 is carried out, this time with symbol synchronization and detection aided by matched




Fig. 3: Block diagram of signal processing chain, where $\hat{a}_m \in \mathbb{N}_4 \triangleq \{0, 1, 2, 3\}$ is the hard-decision estimated value of the m th symbol's QPSK phase shift.




Fig. 4: Measured QPSK symbol values after Doppler correction and symbol synchronization with matched filtering to $p(t)$ after determination of β_r . The signal's pre-correlation SNR is 10.35 dB.

filtering, which reduces ISI and therefore improves symbol detection accuracy.

The process leading to symbol demodulation requires an entire channel to be present in the data; thus only our wide captures are suitable. Fig. 4 shows the successful result of symbol demodulation with matched filtering to $p(t)$ after determination of β_r for 10 μ s of data.

E. Estimation of β_r

The rolloff factor β_r can be estimated based on the preliminary symbol estimates produced without matched filtering. Only the wide captures are suitable for this, as the narrow captures distort the pulse shape. Let the n th sample after resampling and coarse frequency compensation be denoted $y_s(n)$, $n \in \mathbb{Z}$. Suppose we isolate a subset of samples $\{y_s(n) : n \in \mathcal{N}\}$ for a given set $\mathcal{N}$ of contiguous indices. Following the steps depicted in Fig. 3, we obtain the preliminary hard-decision demodulated symbols $\{\tilde{a}_m \in \mathbb{N}_4 : m \in \mathcal{M}\}$, where $\mathcal{M}$ is the set of symbol indices corresponding to the sample indices in $\mathcal{N}$. We can then generate a local replica $l(n; \beta_r)$, $n \in \mathcal{N}$ of the resampled-and-coarse-frequency-compensated signal by a discrete version of (1) with $\tilde{a}_m$ substituted for a_m , $m \in \mathbb{Z}$, as a function of a candidate rolloff factor β_r for $p(t)$. Defining a search space $\mathcal{B}$ of possible rolloff factors, we can determine which local replica best matches the original signal. Our estimator for β_r thus becomes

$$\hat{\beta}_r = \underset{\beta_r \in \mathcal{B}}{\operatorname{argmax}} \sum_{n \in \mathcal{N}} y_s(n) l^*(n; \beta_r) \quad (8)$$




Fig. 5: The objective function from (8) vs. β_r with correlation based on 1 ms of data.

Fig. 5 shows the normalized objective function over $\beta_r \in [0, 0.25]$. The function is maximized for $\beta_r \approx 0.1$. Setting $\mathcal{B} = \{0.25, 0.2, 0.15, 0.1, 0.05\}$, which are the values on this range allowed by the DVB-S2X standard [45], an extension of DVB-S2, we conclude that the OneWeb downlink rolloff factor is $\beta_r = 0.1$.

F. Estimation of Channel Center Frequencies

OneWeb is authorized to downlink in the 10.7–12.7 GHz Ku band [46]. Taking into account the 250-MHz channel spacing mentioned in [29] and existing literature showing channelization [25], we suppose there are eight channels spaced 250 MHz apart within the 10.7–12.7 GHz band.

For channel center frequency estimation, we follow the same procedure outlined in [5]. For this, we exploit a synchronization sequence in the OneWeb signal that will be described later on. The sequence as transmitted repeats once per ms. We estimate the effect of time scaling on the received signal from a series of synchronization sequence time of arrival (TOA) measurements. From (2), one notes that time scaling is only a function of the CFO. From this we obtain an estimate $\hat{\beta}$ of the CFO parameter.

Next we estimate the same CFO parameter from the exponent in (2) given an *a priori* estimate of the channel center frequency, denoted $\bar{F}_{ci}$. If the prior channel center frequency estimate is correct, then the offset $\bar{F}_{ci} - F_{ci}$ can be compensated, leaving the CFO as the sole effect on the frequency shift, expressed as F_d . We can estimate the CFO parameter from the frequency shift as $\bar{\beta} = -F_d/\bar{F}_{ci}$. With the two estimates $\hat{\beta}$ and $\bar{\beta}$, we can then estimate the channel center in MHz as

$$\hat{F}_{ci} = \left\lfloor \frac{\bar{F}_{ci}}{1 + \bar{\beta} - \hat{\beta}} \right\rfloor, \quad i = 1, 2, \dots, 8 \quad (9)$$

where rounding to the nearest MHz is justified for the same reasons given in [5].

From this we conclude that the OneWeb downlink channels are centered at $F_{ci} = (10.7 + 0.25(i - 0.5))$ GHz, $i = 1, 2, \dots, 8$. These are roughly aligned with Starlink's channel centers as reported in [5]. We only observe activity on channels 2, 3, and 4 over Austin, Texas. The baseband signal model in (1) represents a single channel of the MF-TDMA signal from the transmitter.

Table I summarizes the parameter values for the OneWeb Ku-band downlink signal as found by applying the foregoing estimators.

TABLE I: OneWeb Downlink Signal Parameter Values

Name	Parameter	Value
Symbol rate	F_{sym}	230.4 MHz
Rolloff factor	β_r	0.1
i th channel center frequency	F_{ci}	$10.7 + 0.25(i - 0.5)$ GHz




Fig. 6: Overlaid power spectral density plots from wide captures centered at F_{c2} , F_{c3} , and F_{c4} . Data are from ONEWEB-0696, ONEWEB-0352, and ONEWEB-0644 in July 2024.

V. FURTHER SYSTEM ANALYSIS

We can exploit our knowledge of the OneWeb downlink signal parameters and our ability to demodulate the signal to further probe OneWeb's system behavior and data structure. In particular, we seek answers to the following questions: How are channels mapped to beams? What explains the 10-ms periodicity in OneWeb data noted in earlier studies? What can be revealed about the frame substructure? Do there exist any sequences that are invariant across frames, beams, channels, and satellites, as with Starlink's PSS and SSS?

A. Channel Activity

Fig. 6 shows a composite power spectral density plot of three separate wide captures centered at F_{c2} , F_{c3} and F_{c4} , showing activity on the three channels. We observe the channels regularly fade in and out of activity over ~ 20 -second windows in such a way that no channel is continuously active. Activity on channels 2 and 4 is synchronous, whereas channel 3 is only active when the other two are not. This behavior is consistent with the spectrogram shown as Fig. 5 in [28]. (Note that the channel indexing in [28] starts with what we designate as channel 2.)

B. Beam Patterns

OneWeb downlink signals within a given channel exhibit some patterns that are observable even without knowledge of the signal parameters. One such pattern is the signals' time-varying power, as shown in the top plot of Fig. 7. This variation in power is a result of the receiver being illuminated by different transmitter beams as the satellite passes overhead. OneWeb satellites can reportedly produce 32 user beams in the Ku-band, half of which are currently active for downlink [47]. These beams emanate from the satellite antenna array in a fixed pattern that projects a total footprint of roughly 1080×1080 km² on Earth's surface. Assuming each of the 16 downlink beams covers approximately the same area, the footprint of each is roughly 68×1080 km².




Fig. 7: Top: Normalized signal power for a single OneWeb channel. Colors indicate the intervals during which the receiver was within the footprint of an active beam. Bottom Left: Satellite ground track relative to the receiver location (lower vertex of the black triangle). Each arrow shows the movement of the sub-satellite point over the time interval with the corresponding color in the top plot. Thus, when the first recorded beam was passing over the receiver during the interval marked in red in the top plot, the sub-satellite point was entering Oklahoma. Bottom Center: Beam footprints (modeled as rectangular) at the time of the first peak, assuming the receiver is centered latitudinally within the 3rd of 16 beams. Colors indicate correspondence with the signal power time history. Bottom Right: Beam footprints at the time of the second peak, at which point the receiver is centered within the 5th beam. Data are for ONEWEB-0114 from a capture centered at F_{c2} taken in March 2024.

Consider the power time history shown in Fig. 7, which lasts about 190 seconds. As the satellite moves overhead, the receiver is illuminated by successive beams, resulting in a power variation pattern similar to that observed in Fig. 5 of [25]. For further analysis, we generate satellite positions using TLE ephemerides and manually align the center of the third beam's footprint to our receiver location at the moment of the first peak in power (the capture began after the first two beams passed over the receiver's location). Tracking the beam footprints for subsequent peaks makes it clear that every other beam of the 16 is inactive on a given channel. A reasonable explanation for this observation, and for the alternating channel activity noted above, is that OneWeb's frequency reuse strategy activates successive beams on different channels. Thus, for the satellite whose data are shown in Fig. 7, one may reasonably assume that channels 2 and 4 are active on beams 1, 3, $\dots$, 15, whereas channel 3 is active on beams 2, 4, $\dots$, 16.

C. Data Patterns across Beams

For frames captured within approximately 20 seconds of each other, we find a repeating pattern of nearly identical demodulated symbols with a 10-ms period. This pattern is consistent with the 10-ms-spaced peaks in the signal's autocorrelation function noted in [14], [25]. Such repetition implies a signal with low information content, which is clearly incompatible with a high-rate communication system. We suspect that near our receiver's location in Austin, Texas, OneWeb currently has few customers, if any, and that, like the DVB-S2 standard, OneWeb's protocol inserts default physical-layer frames when there is low demand for downlink data.

The 10-ms periodicity is connected to OneWeb's fixed beams. This can be demonstrated by a narrow capture whose time span is long enough for data from at least two active beams to be present. First, we use TLE ephemerides to eliminate motion-induced CFO effects over the capture, as described in Section IV-B. Next we choose a 10-ms segment from somewhere within the capture. Call this a reference segment. We normalize the reference segment to have unit energy and cross-correlate it against the full capture. The resulting correlation profile reveals that the reference segment is only strongly correlated with an approximately 20-second window of the capture. By selectively choosing the reference segment to maximize the correlation within its window, and repeating this process for various reference segments, we can produce a composite plot like the blue, red, and violet traces shown in Fig. 8. The pattern that emerges matches the power profile over the capture, confirming that each active beam has a unique data pattern, as first reported in [25].

To appreciate OneWeb's frequency reuse strategy, consider a reference segment chosen at the moment when the receiver is equidistant from the latitudinal centers of two active beam footprints. If there is significant overlap in the active beam footprints, then this reference segment will be strongly correlated with the unique data pattern associated with both beams, albeit at a lower power for each beam than if it had been chosen at the center of that beam. This is in fact exactly what we observe, as shown by the gold trace in Fig. 8. Clearly, there is considerable overlap between active beams even when only every other beam is active on a given channel. This explains why OneWeb does not allow adjacent beams to be active on




Fig. 8: Cross-correlation power of four different unit-energy-normalized 10-ms reference segments, extracted at the times noted by the respective circles, against the full un-normalized CFO-compensated capture. The correlation peaks for reference segments taken at the maximum of each beam are approximately equivalent in magnitude, indicating uniform power across beams. The correlation peak of the gold trace is lower because the overall power of the captured data is reduced at the seam between beams. Data are from ONEWEB-0114 in June 2024.

the same channel despite such a policy reducing the system's area spectral efficiency. Each beam pattern would have to be significantly sharper in its power rolloff to allow each beam to carry all channels without significant inter-beam interference.

D. Frame Structure

OneWeb's simultaneous operation of single-carrier signals in multiple channels constitutes a Multi-Frequency Time-Division Multiple Access (MF-TDMA) scheme. In such schemes, it is typical for a controller to devise a terminal burst time plan (TBTP) to efficiently transmit data to each UT, possibly in response to requests from the UTs on the uplink side. The downlink stream is divided into frames within which each user is assigned one or more time slots to receive data, per the TBTP. Guard intervals filled with predictable data are typically inserted as a preamble within each slot to facilitate slot identification and continual time and frequency synchronization.

The signal demodulation process depicted in Fig. 3, including matched filtering to the SRRC pulse $p(t)$, produces a sequence of symbol estimates. Let $\hat{a}_m \in \mathbb{N}_4$ denote the m th symbol estimate for a given satellite, and let $\mathbf{a} = (\hat{a}_m : m \in \mathcal{M})$ denote the vector of symbol estimates with sequential indices for some set $\mathcal{M} \subset \mathbb{Z}$ of contiguous indices. By examining such sequences for different $\mathcal{M}$ and different satellites, we can identify patterns that reveal the fundamental structure of the OneWeb signal, including repeating sequences.

In view of the strong 10-ms periodicity observed in the estimated symbols across all satellites, we define the frame period T_f to be 10 ms, or $N_f = 2304000$ symbols. Key questions are the following: (1) Within the same beam, what fraction of the N_f symbols repeat from frame to frame? (2) Are there frame sub-segments that are common across all beams of the same satellite? (3) Are there such sub-segments common across all satellites?

Our approach to addressing these questions proceeds as follows. For a given satellite and beam, we process N frames'

worth of wide-capture samples collected near the latitudinal center of the beam through the decoding pipeline in Fig. 3 to produce a sequence of symbol estimates. One can think of these data as having been collected near one of the peaks shown in Fig. 8. Initially, there is no basis by which to define the beginning of a frame, so we arbitrarily select the first N_f symbols and declare these to be the first frame, etc. Let $\mathbf{a}_n = (\hat{a}_m : m \in \mathcal{M}_n) \in \mathbb{N}_4^{N_f}$ denote the vector of estimated symbols over the n th frame, with $\mathcal{M}_n = \{(n-1)N_f+1, (n-1)N_f+2, \dots, nN_f\}$, and let $\mathbf{a}_n(m)$ denote the m th element of $\mathbf{a}_n$, with $m \in [1, N_f] = \mathcal{M}_1$.

Taking advantage of the strong frame-to-frame correlation evident in Fig. 8, we correlate $\mathbf{a}_n$, $n \in \{2, 3, \dots, N\}$ against $\mathbf{a}_1$ to verify that the frames are mutually aligned with no extraneous or missing symbols in any frame. We then compute the consensus frame $\bar{\mathbf{a}} \in \mathbb{N}_4^{N_f}$ for the capture, whose m th element $\bar{\mathbf{a}}(m)$ is the mode of the set $\{\mathbf{a}_1(m), \mathbf{a}_2(m), \dots, \mathbf{a}_N(m)\}$.

At this point, we can compare each frame symbol-by-symbol against $\bar{\mathbf{a}}$. We define the agreement ratio R_m as

$$R_m = \frac{1}{N} \sum_{n=1}^N \mathbf{1}(\mathbf{a}(m) = \bar{\mathbf{a}}_n(m)), \quad m \in \mathcal{M}_1 \quad (10)$$

where $\mathbf{1}(x)$ is the indicator function, equal to 1 when x is true and otherwise 0. We assume that frames may contain synchronization sequences or default payload data having perfect frame-to-frame agreement. We denote the index set for these by $\mathcal{S} \subset \mathcal{M}_1$. We further assume that frames may contain payload data, identified by the index set $\mathcal{P} \subset \mathcal{M}_1$, whose values are independent and uniformly distributed over the domain of a_m (which is $\mathbb{N}_4$ for QPSK). Finally, we assume that frames may contain intermediate-type data, identified by the index set $\mathcal{I} \subset \mathcal{M}_1$, that are neither perfectly constant nor perfectly random from frame to frame. We assume these three sets are exhaustive so that $\mathcal{M}_1 = \mathcal{S} \cup \mathcal{P} \cup \mathcal{I}$.

In a noise-free scenario, $R_m = 1$ for $m \in \mathcal{S}$, and $R_m \approx 0.25$ for $m \in \mathcal{P}$, assuming QPSK modulation and large N . Under the AWGN signal model (2), and assuming large N and perfect phase and symbol synchronization (conditions of coherent detection), $R_m \approx 1 - P_e$ for $m \in \mathcal{S}$, where P_e is the symbol error rate, which can be calculated based on SNR and F_{sym} [48].

Another useful analytical metric is the frame mismatch rate M_n , or the fraction of symbols in $\mathbf{a}_n$ that fail to match the corresponding symbols in $\bar{\mathbf{a}}$:

$$M_n = \frac{1}{N_f} \sum_{m=1}^{N_f} \mathbf{1}(\mathbf{a}(m) \neq \bar{\mathbf{a}}_n(m)) \quad (11)$$

Fig. 9 presents exemplary results for a capture with SNR = 10 dB and $N = 100$. The top plot indicates near-unity R_m , $m \in \mathcal{M}_1$ except for short bursts spaced by 1 ms during which R_m falls to near 0.25. Each burst lasts approximately 25 μs and contains a 2- μs subinterval over which R_m is again approximately unity. The frame mismatch rate M_n , shown in the bottom plot, hovers around 0.02, meaning that about 98% of the symbols in a contiguous set of frames agree with the consensus frame. This may be compared with $P_e = 5.8 \times 10^{-3}$, the symbol error rate corresponding to SNR = 10 dB under the




Fig. 9: Top: Agreement ratio R_m for $m \in \mathcal{M}_1$ over a 10-ms capture. The ratio has been smoothed by a 10-sample moving average for visual clarity. Bottom: The frame mismatch rate M_n for $n \in \{1, 2, \dots, 80\}$. Data are for a capture with SNR = 10 dB from ONEWEB-0394 in June 2024.

AWGN model. Clearly, contiguous frames are highly similar, but there exists some systematic disagreement between them that is not due merely to AWGN. Similar results were obtained for other captures analyzed.

Based on these results, we can develop a basic model of OneWeb's frame structure. According to this model, shown in Fig. 10, a 10-ms frame spans 10 1-ms slots, each consisting of a header and a payload. The header, which corresponds to the 1-ms-spaced intervals of mostly low agreement ratio shown in the top plot of Fig. 9, contains a synchronization sequence and other information unique to each slot. The synchronization sequence occurs at approximately the midpoint of the header and corresponds to the brief 2- μ s subinterval over which R_m is approximately unity. The synchronization sequence will be detailed in a later section. The payload typically contains unique downlink data destined to one or more users. But when no active users are present, it is filled with data from a sequence of default payloads.

For further discussion of header and payload properties, it will be convenient to introduce some additional notation. Let $\mathbf{h}_{sfbc_i}$ and $\mathbf{p}_{sfbc_i}$ respectively be vectors containing the header and payload symbols for slot s of frame f of beam b on channel c of satellite i . Also let $\mathcal{F}_{bci}(k)$ be the set of frame indices transmitted by the b th beam on the c th channel of the i th satellite for the k th default payload constancy window, defined as the time interval over which default payload remains constant. The k th default payload $\mathbf{d}_{bci}(k) = (\mathbf{p}_{1fbci}(k), \mathbf{p}_{2fbci}(k), \dots, \mathbf{p}_{10fbci}(k))$ spans a full frame and is identical for all $f \in \mathcal{F}_{bci}(k)$. However, the slot-level payloads are different from one another ($\mathbf{p}_{sfbc_i}(k) \neq \mathbf{p}_{lfbc_i}(k)$ for $s \neq l$). Moreover, the default payload is different from




Fig. 10: Model of OneWeb's downlink frame structure.

beam to beam ($\mathbf{d}_{bci}(k) \neq \mathbf{d}_{lci}(k)$ for $b \neq l$), different from channel to channel ($\mathbf{d}_{bci}(k) \neq \mathbf{d}_{bli}(k)$ for $c \neq l$), and different from one constancy window to another ($\mathbf{d}_{bci}(k) \neq \mathbf{d}_{bli}(l)$ for $k \neq l$).

Our observation has been that the default payload constancy window lasts approximately one hour. During this time, the default payload for beam b on channel c of satellite i is identical to that for all other satellites on the same beam and channel ($\mathbf{d}_{bci}(k) = \mathbf{d}_{bcl}(k) \triangleq \mathbf{d}_{bc}(k)$ for $i \neq l$). Our estimate of the default payload constancy window duration was obtained by capturing data from corresponding beams of many satellites over a two-hour interval. Although our antenna can only point to a single satellite at one time, and requires some time to switch between them, we conjecture that each $\mathbf{d}_{bc}(k)$ remains constant because it was found to be so for each satellite tracked during the k th window. Three constancy windows were observed during the 2-hour interval, two partial and one complete.

E. Synchronization Sequence

Discovery of the default payload $\mathbf{d}_{bci}(k)$ is significant for opportunistic use of OneWeb signals. When known, it enables long-duration correlation and thus high processing gain. But it must be noted that opportunistic use of the default payload is complicated by the need to re-estimate $\mathbf{d}_{bci}(k)$ for each beam b , channel c , satellite i , and constancy window k . Moreover, $\mathbf{d}_{bci}(k)$ is only continually present when the OneWeb network is unburdened. In environments where users are highly active, $\mathbf{d}_{bci}(k)$ will be only intermittently present or absent altogether.

Thus, it would be of great value to discover a symbol sequence within the OneWeb data that is invariant across slots, frames, beams, channels, and satellites, as is true for Starlink's PSS and SSS [5]. Such invariance would make the sequence especially useful for correlation against OneWeb signals in cold-start conditions, and in high-data-rate-demand environments where the default payload is not present. Moreover, it would constitute an unambiguous feature in the OneWeb data stream relative to which an evolving frame structure model could be referenced.

To avoid confusion with the default payload, finding such a sequence requires correlation of frame-length estimated symbol vectors across beams, channels, and satellites. Let $\bar{\mathbf{a}}_{bci}$ be a length- N_f consensus frame for beam b on channel c of

satellite i . Circularly correlating this against similar consensus frames from different beams, channels, and satellites reveals peaks spaced by $N_f/10$ samples, or one slot length. Aligning a large number of diverse consensus frames to the nearest slot and then calculating an agreement ratio R_m , $m \in \mathcal{M}_1$, what emerges are short 400-symbol ($\sim 2\text{-}\mu\text{s}$) bursts of near-perfect agreement spaced by 1 ms. We call this sequence of symbols the OneWeb synchronization sequence (SS). The SS resides at approximately the midpoint of each slot's header. In Fig. 9, it is visible in the top panel's inset as the short subinterval over which $R_m \approx 1$.

To verify the SS's invariance, we correlated it against Doppler-corrected narrow captures that were not involved in its initial identification. The verification data encompassed narrow captures from diverse beams, channels, and satellites across various days in June 2024. In all cases we found strong correlation against the candidate SS and so consider it verified.

We provide the SS in the form of an 800-bit hexadecimal number q_{ss} in which each pair of bits represents a symbol:

```
qss = B5D0 CDB5 66F9 5A93 F90B 0060 834E 073C
      9EC3 EAAA D425 C677 93B0 EE1F 993C 5CF5
      2FFE 5839 CC7E 5170 FE09 31EF 33CD 3E13
      16F4 3E9E 2A17 5D4B 2D9B E629 2E62 6386
      B994 6849 7811 5074 5930 417E 3338 E497
      3A3A 5B05 CFBD 5A8F 669D 9D31 EEB8 B48C
      B7E2 2DBA
```

Let $m \in \{0, 1, 2, \dots, N_{ss} - 1\}$ be the symbol index within the SS, where $N_{ss} = 400$ is the number of symbols. The decimal value a_m obtained from the hexadecimal sequence is the symbol phase used in (1) to generate a QPSK signal.

$$a_m = \left\lfloor \frac{q_{ss}}{4^m} \right\rfloor \bmod 4 \quad (12)$$

The formula in (12) extracts the m th symbol by first dividing q_{ss} by 4^m so that the desired bit pair occupies the two least significant bits after flooring. Taking the modulo 4 isolates this pair from all others. To ensure proper interpretation of (12), we provide the first eight values of a_m :

$$(a_0, \dots, a_7) = (2, 2, 3, 2, 1, 3, 2, 0)$$

The duration of the SS is $T_{ss} = N_{ss}/F_{sym} = 1.73 \mu\text{s}$ and its period is 1 ms. Curiously, we find that symbols at indices $m = 10, 25$, and 42 at times vary from SS to SS, whereas all others are invariant.

VI. CONCLUSIONS

We have presented and applied a blind signal identification process to uncover key parameters in OneWeb's Ku-band downlink signal. We further identified a synchronization sequence that can be used to passively exploit OneWeb signals for pseudorange-based positioning, navigation, and timing (PNT), and have proposed a model for the OneWeb frame structure. Moreover, we have discovered that when no active users are present, the OneWeb symbol stream is identical across all satellites for corresponding beams and channels for a period of approximately one hour. The results in this paper illuminate a path to use of OneWeb signals as a backup to traditional GNSS for PNT. Further studies are needed to probe

OneWeb's timing stability and relationship to an absolute time scale such as GPS time.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Z. Kassas, in *Navigation from low earth orbit — Part 2: Models, implementation, and performance in Position, Navigation, and Timing Technologies in the 21st Century*, 2021, vol. 2.
- [2] N. Jarak and Q. Jault, "The potential of LEO satellite-based opportunistic navigation for high dynamic applications," *Sensors*, vol. 22, no. 7, p. 2541, 2022.
- [3] M. L. Psiaki, "Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids," *NAVIGATION*, vol. 68, no. 3, pp. 621–641, 2021.
- [4] P. A. Iannucci and T. E. Humphreys, "Fused low-Earth-orbit GNSS," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2022.
- [5] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–16, 2023.
- [6] Z. M. Komodromos, S. C. Morgan, Z. L. Clements, W. Qin, W. J. Morrison, and T. E. Humphreys, "Network-aided pseudorange-based LEO PNT from OneWeb," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.
- [7] S. C. Morgan, Z. M. Komodromos, W. Qin, Z. L. Clements, A. M. Graff, W. J. Morrison, and T. E. Humphreys, "A mock implementation of fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.
- [8] T. E. Humphreys, "Interference," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 469–503.
- [9] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, M. L. Psiaki, and T. E. Humphreys, "First results from three years of GNSS interference monitoring from low Earth orbit," *NAVIGATION*, vol. 68, no. 4, pp. 673–685, 2021.
- [10] Z. Clements, P. Ellis, and T. E. Humphreys, "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit," in *Proceedings of the IEEE/ION PLANS Meeting*, Monterey, CA, 2023, pp. 458–469.
- [11] G. S. Workgroup, "GPS spoofing: Final report of the GPS spoofing workgroup," OPSGROUP, Tech. Rep., 2024. [Online]. Available: <https://ops.group/blog/gps-spoofing-final-report>
- [12] Z. L. Clements, P. B. Ellis, M. J. Murrian, M. L. Psiaki, and T. E. Humphreys, "Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit," *NAVIGATION*, 2025, submitted for review.
- [13] W. S. Limited, "Amendment to modification application for U.S. Market Access Grant for the OneWeb Ku- and Ka-Band system," https://licensing.fcc.gov/myibfs/download.do?attachment_key=3495551, Jan. 2021, SAT-APL-20210112-00007.
- [14] S. Kozhaya, H. Kanj, and Z. M. Kassas, "Multi-constellation blind beacon estimation, Doppler tracking, and opportunistic positioning with OneWeb, Starlink, Iridium NEXT, and Orbcomm LEO satellites," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2023, pp. 1184–1195.
- [15] M. Neinavaie, J. Khalife, and Z. M. Kassas, "Exploiting Starlink signals for navigation: First results," in *Proceedings of the ION GNSS+ Meeting*, St. Louis, Missouri, Sept. 2021, pp. 2766–2773.
- [16] —, "Acquisition, Doppler tracking, and positioning with Starlink LEO satellites: First results," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 3, pp. 2606–2610, 2022.
- [17] J. Khalife, M. Neinavaie, and Z. M. Kassas, "The first carrier phase tracking and positioning results with Starlink LEO satellite signals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 2, pp. 1487–1491, 2022.

- [18] Z. M. Kassas, S. Kozhaya, H. Kanj, J. Saroufim, S. W. Hayek, M. Neinavaie, N. Khairallah, and J. Khalife, "Navigation with multi-constellation LEO satellite signals of opportunity: Starlink, OneWeb, Orbcomm, and Iridium," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2023, pp. 338–343.
- [19] N. Jurdak, R. Adam, and Q. Jault, "Leveraging multi-LEO satellite signals for opportunistic positioning," *IEEE Access*, 2024.
- [20] J. Saroufim and Z. M. Kassas, "Ephemeris and timing error disambiguation enabling precise LEO PNT," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.
- [21] S. Shahcheraghi, J. Saroufim, and Z. M. Kassas, "Acquisition, Doppler tracking, and differential LEO-aided IMU navigation with uncooperative satellites," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.
- [22] H. Sallouha, S. Saleh, S. De Bast, Z. Cui, S. Pollin, and H. Wymeersch, "On the ground and in the sky: A tutorial on radio localization in ground-air-space networks," *IEEE Communications Surveys and Tutorials*, vol. 27, no. 1, pp. 218–258, 2025.
- [23] B. McLemore and M. L. Psiaki, "Navigation using Doppler shift from LEO constellations and INS data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4295–4314, 2022.
- [24] A. Baron, P. Gurfil, and H. Rotstein, "Implementation and accuracy of Doppler navigation with LEO satellites," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 2, 2024.
- [25] S. Kozhaya and Z. M. Kassas, "A first look at the OneWeb LEO constellation: Beacons, beams, and positioning," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–7, 2024.
- [26] R. Blázquez-García, D. Cristallini, M. Ummenhofer, V. Seidel, J. Heckenbach, and D. O'Hagan, "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. [Online]. Available: <https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/rsn2.12446>
- [27] R. Blázquez-García, D. Cristallini, M. Ummenhofer, V. Seidel, J. Heckenbach, and D. O'Hagan, "Experimental comparison of Starlink and OneWeb signals for passive radar," in *2023 IEEE Radar Conference (RadarConf23)*, 2023, pp. 1–6.
- [28] R. Blázquez-García, T. Hauschild, P. Markiton, M. Ummenhofer, V. Seidel, and D. Cristallini, "Passive radar imaging based on multistatic combination of Starlink and OneWeb illumination," in *2024 IEEE Radar Conference (RadarConf24)*, 2024, pp. 1–6.
- [29] Radio Communications & EMC, "OneWeb ow70l UT test report FCC id xxx-intow70ldac," <https://fcc.report/FCC-ID/XXX-INTOW70LDAC/5364479.pdf>, 2021, XXX-INTOW70LDAC.
- [30] E. L. Cid, M. G. Sanchez, and A. V. Alejos, "Wideband analysis of the satellite communication channel at Ku- and X-bands," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 4, pp. 2787–2790, 2015.
- [31] T. Hobiger, D. Piester, and P. Baron, "A correction model of dispersive troposphere delays for the ACES microwave link," *Radio Science*, vol. 48, no. 2, pp. 131–142, 2013.
- [32] B. Lathi and Z. Ding, *Modern Digital and Analog Communication Systems*, ser. Oxford series in electrical and computer engineering. Oxford University Press, 2019.
- [33] T. Zhao and T. Huang, "Cramer-Rao lower bounds for the joint delay-Doppler estimation of an extended target," *IEEE transactions on signal processing*, vol. 64, no. 6, pp. 1562–1573, 2016.
- [34] M. Neinavaie, J. Khalife, and Z. M. Kassas, "Doppler stretch estimation with application to tracking globalstar satellite signals," in *MILCOM 2021 - 2021 IEEE Military Communications Conference (MILCOM)*, 2021, pp. 647–651.
- [35] O. A. Dobre, "Signal identification for emerging intelligent radios: Classical problems and new challenges," *IEEE Instrumentation & Measurement Magazine*, vol. 18, no. 2, pp. 11–18, 2015.
- [36] H. Sun, S. Zhong, J. Tang, and J. Yuan, "A blind estimation method of QPSK/QQPSK symbol rate," in *2022 4th International Conference on Communications, Information System and Computer Engineering (CISCE)*. IEEE, 2022, pp. 138–141.
- [37] W. Gardner, "Exploitation of spectral redundancy in cyclostationary signals," *IEEE Signal Processing Magazine*, vol. 8, no. 2, pp. 14–36, 1991.
- [38] L. Mazet and P. Loubaton, "Cyclic correlation based symbol rate estimation," in *Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers (Cat. No.CH37020)*, vol. 2. IEEE, 1999, pp. 1008–1012 vol.2.
- [39] Y. Jin and H. Ji, "Robust symbol rate estimation of PSK signals under the cyclostationary framework," *Circuits, systems, and signal processing*, vol. 33, no. 2, pp. 599–612, 2014.
- [40] W. A. Gardner, *Statistical spectral analysis : a nonprobabilistic theory / by William A. Gardner*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [41] T. Nakagawa, M. Matsui, T. Kobayashi, K. Ishihara, R. Kudo, M. Mizoguchi, and Y. Miyamoto, "Non-data-aided wide-range frequency offset estimator for QAM optical coherent receivers," in *2011 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference*, 2011, pp. 1–3.
- [42] Y. Wang, K. Shi, and E. Serpedin, "Non-data-aided feedforward carrier frequency offset estimators for QAM constellations: A nonlinear least-squares approach," *EURASIP journal on advances in signal processing*, vol. 2004, no. 13, pp. 856 139–856 139, 2004.
- [43] P. Ciblat and M. Ghogho, "Blind NLLS carrier frequency-offset estimation for QAM, PSK, and PAM modulations: performance at low SNR," *IEEE Transactions on Communications*, vol. 54, no. 10, pp. 1725–1730, 2006.
- [44] M. Rice, *Digital Communications: A Discrete-time Approach*. Pearson/Prentice Hall, 2009.
- [45] ETSI, "Digital Video Broadcasting (DVB); Second Generation Framing Structure, Channel Coding and Modulation Systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications; Part 2: DVB-S2 Extensions (DVB-S2X)," European Telecommunications Standards Institute (ETSI), ETSI Standard EN 302 307-2 V1.1.1, October 2014. [Online]. Available: https://www.etsi.org/deliver/etsi_en/302300_302399/30230702/01.01.01_20/en_30230702v010101a.pdf
- [46] Federal Communications Commission, "ONEWEB NON-GEOSTATIONARY SATELLITE SYSTEM technical information to supplement Schedule S," <https://fcc.report/IBFS/SAT-LOI-20160428-00041/1134939.pdf>, June 2017, SAT-APL-20210112-00007.
- [47] —, "ONEWEB NON-GEOSTATIONARY SATELLITE SYSTEM PHASE 2: MODIFICATION TO AUTHORIZED SYSTEM technical information to supplement Schedule S," <https://fcc.report/IBFS/SAT-MPL-20200526-00062/2379706.pdf>, May 2020, SAT-MPL-20200526-00062.
- [48] M. K. Simon and M. Alouini, *Digital Communication over Fading Channels*, 2nd ed. New York: Wiley, 2005.

A Mock Implementation of Fused LEO GNSS

Samuel C. Morgan*, Zacharias M. Komodromos[†], Wenkai Qin[†], Zachary L. Clements*,
Andrew Graff[‡], W. Jeremy Morrison*, Todd E. Humphreys*

**Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin*

[†]Department of Electrical and Computer Engineering, The University of Texas at Austin

[‡]CesiumAstro

Abstract—This paper presents the first experimental demonstration of fused low Earth orbit (LEO) GNSS using pseudorange measurements, achieved through a mock implementation that distributes orbit and clock models as if supplied by the constellation operator. Obtaining resilient and accurate position, navigation, and timing (PNT) from broadband LEO constellations is a topic of vigorous ongoing research. Such constellations offer key advantages over traditional GNSS: proliferated satellite coverage, high power, wideband signals, and the potential for near-zero age-of-data clock and orbit models. SpaceX’s expanding Starlink constellation and known frame structure make it an especially compelling candidate for LEO PNT. But recent analysis of its downlink frame timing and its published orbits indicates that stand-alone opportunistic use of Starlink signals (i.e., without cooperation from Starlink nor access to accurate third-party models) could not provide pseudorange-based PNT competitive in performance with traditional GNSS. By contrast, a fused LEO GNSS approach in which Starlink supplies real-time clock and orbit corrections to subscribers, offers a promising pathway to accurate PNT despite frame timing irregularities. This study develops downlink-beam-specific parameterized orbit and clock models tailored to Starlink’s frame clock variations. Then, as if from the constellation operator in a hypothetical fused LEO GNSS service, these models are distributed to users who combine them with live-signal Starlink observables to produce navigation and timing solutions with 10-meter-level accuracy. These results demonstrate that fused LEO GNSS is a viable alternative or complement to traditional GNSS for moderately accurate and resilient PNT.

Index Terms—Starlink, positioning, low Earth orbit, fused LEO GNSS

I. INTRODUCTION

High data throughput has quickly become the standard for enabling high-quality video streaming, efficient remote work, and realistic low-latency extended and augmented reality, among other applications. Until recently, these capabilities were limited to areas with existing infrastructure or significant financial investment in new infrastructure. The advent of low Earth orbit (LEO) mega-constellation broadband networks has expanded these capabilities globally.

To achieve high data throughput and serve multiple users, large signal bandwidths are required. Interestingly, these large bandwidths inherently enable highly accurate time of arrival (TOA) measurements, which are critical for positioning, navigation, and timing (PNT) services. However, current LEO broadband networks such as SpaceX’s Starlink and Eutelsat’s OneWeb have yet to fully capitalize on this advantage by offering PNT services to their users. While Starlink offers a Starlink positioning mode for their user terminals, it is far from mature, often taking several minutes to achieve a

low-resolution position solution [1]. Moreover, it is not a service offered directly to users, but is instead only an internal mechanism to support a user terminal’s initial network entry.

The use of LEO constellations for PNT is already well-studied in the literature [2], [3]. They offer an attractive complement or backup to GNSS amid rising threats of GNSS jamming and spoofing [4]–[8]. Broadly, four primary approaches have been identified. Table I summarizes each approach in terms of deployment costs, availability, accuracy, time to fix, and dependency on traditional GNSS. References [9] and [10] give an overview of several dedicated LEO PNT projects, while the authors in [11] develop one of the first receivers for such a constellation—Xona Space Systems’ Pulsar. The fused LEO GNSS concept of operations is explored in [12] and [13], and methods for extracting high-quality TOA measurements using existing communications infrastructure within a fused LEO GNSS paradigm are presented in [14] and [15].

Two recent examples of network-aided opportunistic LEO PNT based on Doppler measurements are presented in [17] and [18], with promising positioning results. In [17], a mobile rover is aided by a reference station with known position. This approach additionally requires GNSS timing for initialization. The authors of [18] performed simulations along with two field experiments. In their first experiment, they eliminated satellite vehicle (SV) clock and ephemeris errors by using the same receiver for both the reference and the rover. In their second experiment, a 635-km baseline separated the reference and rover, implying they received signals from different Starlink beams. As reported in [26], significant frame and frame-rate offsets exist from beam to beam, but whatever carrier Doppler offsets were present appear not to have prevented a useful positioning solution.

Note that Table I also classifies the opportunistic approach presented in [16] as network-aided because the receiver being tracked is aided by GNSS signals during an initialization phase, which allows estimation of LEO satellite ephemeris, clock, and clock rate errors just as occurs with network aiding. In effect, the receiver acts as its own reference station during the initialization phase, holding over key estimated values during the GNSS-denied phase.

The authors in [19] present another network-aided opportunistic approach using the Iridium constellation. They set up nearly-zero and short baselines between their reference and rover receivers, then perform positioning using time difference of arrival and frequency difference of arrival. But their PNT solutions attain only modest accuracy due to Iridium’s narrow bandwidth, which limits TOA measurement resolution.

TABLE I: LEO PNT Trade Space

	Dedicated [9]–[11]	Fused [12]–[15]	Network-Aided Opportunistic [16]–[20]	Stand-Alone Opportunistic [21]–[25]
Description	LEO constellation or hosted payloads dedicated solely to PNT.	Fuses a secondary PNT service with a primary communications service.	Like stand-alone opportunistic, but a network of reference stations provides corrections.	Exploits unmodified signals from communications SVs. Public ephemerides. No network aiding.
Marginal deployment cost	high: constellation of SVs	low: uses communications hardware and signals	high: network of reference stations	very low
Potential availability	mid-term	near-term	near-term for local coverage	immediate
Potential accuracy [m]	< 1	< 1	< 1	< 100
Time to fix [s]	< 10	< 10	< 10	< 1000
Dependency on traditional GNSS	@SVs in near-term	@SVs in near-term	@reference stations in near-term	@SVs in some cases

In [20] the authors implement a similar approach using the OneWeb constellation. They envision a third-party network-aided opportunistic approach in which the reference stations provide clock and orbit models as a form of differential corrections. Due to OneWeb’s wide bandwidth signal, they achieve good PNT accuracy despite poor-quality ephemerides.

Notably, the overwhelming majority of opportunistic LEO PNT methods presented in the literature are Doppler-based [27]. But compared to pseudorange-based PNT techniques, Doppler-based techniques have worse timing accuracy by many orders of magnitude (milliseconds vs. nanoseconds), even under optimistic measurement noise and satellite clock offset rate assumptions [22], [24], [28]. Recognizing that many PNT applications of practical interest require accurate timing, we focus on pseudorange-based PNT, or PNT based on TOA measurements.

Among current LEO broadband providers, Starlink holds the greatest potential for highly accurate pseudorange-based LEO PNT. Up to 16 Starlink beams may be directed to a single service cell, providing favorable instantaneous geometry. Additionally, Starlink’s wide bandwidth signals ($F_s = 240$ MHz) contain deterministic pilot sequences that enable extremely accurate TOA measurements [26], [29], [30]. Stand-alone opportunistic pseudorange-based PNT with Starlink appears intractable due to onboard clocks exhibiting inconsistent, unmodelable behavior [26]. Therefore, we look to a network-aided or fused approach for this paper’s pseudorange-based PNT formulation, aiming for meter-accurate positioning and nanosecond-accurate timing.

Given that timing errors for Starlink SVs are in the worst case beam-specific [26], a network-aided opportunistic PNT solution would be especially costly: the inter-station spacing for the reference network would need to be smaller than the Starlink beam diameter, or approximately 20 km. Thus, a fused communications-and-PNT solution, as advocated in [12], appears more favorable for the Starlink constellation. Nonetheless, a fused solution may be simulated with a network-aided solution: the same corrections to ephemeris and timing models, together with pilot symbols for greater detection sensitivity

and improved TOA and Doppler measurement precision, can be sent to a user receiver by the reference network as if it had been sent by the Starlink network itself. Thus, a network-aided implementation amounts to a mock fused implementation.

In this paper we propose treating the Starlink constellation as though it were a fully operational fused LEO GNSS. This approach involves two receivers, a reference station, designated as the control segment (CS), and a user receiver, designated as the customer terminal (CT). The CS functions as a ground station, equipped with a high-gain antenna and a high-quality clock. It monitors SVs transmitting to its service cell and gathers key information, including transmission timestamps, SV clock errors, and data symbols. Using these observations along with its known position, the CS derives accurate clock and orbital models, which it then relays to the CT as if sent within the data message of transmitting Starlink SVs. The CT, in turn, uses this information to form pseudorange measurements and compensate for SV clock errors in its PNT solution.

The CT is assumed to be a consumer-grade device that, in the worst case, operates with a low-gain omnidirectional antenna, a low-quality clock, and a narrow-bandwidth front-end. Under these conditions, it may receive signals in a low signal-to-noise ratio (SNR) regime, making it difficult to extract TOA measurements using the known pilot sequences alone. To mitigate this, the CS may designate additional pilot sequences beyond those used for communications. These additional pilots would enhance TOA measurement accuracy, even in low SNR conditions. However, such data resources would not be allocated arbitrarily for two reasons. First, improper selection could introduce errors due to high sidelobes in the autocorrelation function of the pilot-bearing downlink signal. Second, any allocation of symbols as pilots prevents their use as conveyors of data, cutting into downlink data rate capacity. Instead, a select few data resources are chosen optimally to minimize the Ziv-Zakai bound (ZZB) as described in [14].

We implement the proposed approach in a field campaign. CS and CT receivers are used to simultaneously record data within the same Starlink service cell. The recorded data are

processed as described, and a PNT solution is determined using nonlinear batched least squares (NBLs). The resulting solution achieves 10-meter-level accuracy despite suboptimal capture bandwidth and ephemeris errors, providing motivation for two further developments: (1) the implementation of a dedicated fused LEO GNSS service (as discussed in [12]), or (2) the establishment of a third-party service that deploys reference stations in key service cells, offering the necessary information to client receivers for opportunistic navigation with the Starlink signal, as discussed in [10].

The rest of this paper is organized as follows. In Section II, we present the system concept of operations and its relation to our mock implementation. Section III discusses the Starlink Ku-band downlink signal structure and defines relevant terminology. In Section IV, we present the necessary SV and receiver clock model. Section V enumerates three potential data-aiding strategies to improve TOA measurement accuracy. In Section VI, we explain how the CS provides ephemeris and clock corrections. The estimator used to produce a PNT solution is developed in Section VII. Section VIII describes our specific experiment and presents the PNT results. Finally, Section IX, provides some further insights about the work.

II. CONCEPT OF OPERATIONS

This section outlines the key assumptions underlying the fused LEO GNSS concept as regards the CS, the SV constellation, and the CT. It then provides a brief overview of this paper’s experiment and explains how each component aligns with the established assumptions.

A. Fused LEO GNSS

We assume a constellation of SVs primarily dedicated to providing broadband internet to its users, with a secondary mission to offer PNT services at minimal additional cost to the user and minimal opportunity cost (i.e., reduced throughput) to the communications service. This dual-purpose system, referred to as fused LEO GNSS in [12], will be referred to as such throughout this paper. Additionally, GNSS as we know them (i.e., GPS, Galileo, Beidou, etc.) will be referred to as traditional GNSS from this point forward.

A fused LEO GNSS could elect to rely on traditional onboard GNSS receivers to perform orbit determination and clock disciplining. However, this approach subjects itself to all the weaknesses that a fused LEO GNSS seeks to overcome. That is, service would be denied when traditional GNSS would be denied. It is worth mentioning that despite this key flaw, using traditional onboard GNSS receivers could be a near-term solution for existing constellations such as Starlink and OneWeb as there is nearly no deployment cost.

In the long term, in a fused LEO GNSS, a network of ground stations is essential to perform precise orbit determination and clock monitoring of the constellation. This network of ground stations, referred to as the CS, provides ephemeris and clock corrections to the SVs, which are then disseminated to end users.

The constellation of SVs may operate in a transparent or regenerative scheme. They may transmit a waveform of any

type, across any number of frequency channels, at varying levels of directionality, and using any polarization desired by the service provider to ensure sufficient orthogonality.

The constellation is assumed to operate in LEO and to be a so-called mega-constellation. A mega-constellation consists of a large number of SVs, ranging from a few hundred (e.g., Eutelsat’s Oneweb) to several thousand (e.g., SpaceX’s Starlink). Constellations of this size are crucial as they provide broad communications coverage and excellent geometric diversity for PNT.

The CT is assumed to be a consumer device capable of engaging in whatever resource allocation scheme the network has employed. For example, if the network’s constellation of SVs transmits narrow beams across multiple frequency channels, the CT is assumed to be equipped with a phased array antenna capable of receiving the highly-directional beams and rapidly switching between them as scheduled, along with the necessary radio frequency (RF) chains to process signals at any of the frequency channels.

It is not necessary for every CT to decode all the data transmitted by a given SV. Data destined for other users are assumed to be encrypted and thereby inaccessible to a given CT.

If the CT is designated for PNT, it need only be capable of computing a PNT solution using the network resources allocated for PNT services. This opens the possibility for a distinct category of CT: the PNT-only device. Since the signal-to-noise ratio (SNR) requirements for PNT are significantly lower than those for broadband communication, a service provider could feasibly offer a compact, portable PNT-only CT with a low-gain, possibly omnidirectional, antenna. However, even a PNT-only CT would still require sufficient bandwidth to precisely measure TOA using pilot symbols.

B. Mock Implementation

Since no fused LEO GNSS constellation is yet deployed, we develop an experimental demonstration of the concept via a mock implementation, with Starlink as our target constellation. Our mock implementation retains all essential elements of the fused approach.

In our implementation, the CS consists of a single reference station at a known position, mounted on the roof of the parking garage adjoining the Aerospace Engineering Building at the University of Texas. Equipped with a high-gain antenna, the CS is capable of receiving full-bandwidth ($F_s = 240$ MHz) Ku-band downlink signal data from a single Starlink SV at a time. By processing the digitized received signal, augmented with knowledge of its own position, access to precise GPS time (GPST), and atmospheric models, the CS provides differential corrections to the CT.

Our implementation’s CT is assumed to be PNT-only and is situated at a location unknown to the CS. Specifically, the CT comprises a low-gain antenna and an RF front-end capable of capturing only a 25-MHz real-time bandwidth—a small fraction of the total bandwidth of Starlink Ku-band downlink waveform. The CT is also limited to receiving data on a single frequency channel at a time. This configuration

represents a pessimistic scenario, as the CT cannot reliably and accurately measure TOAs based solely on existing Starlink pilot sequences that support communications, such as the PSS and the SSS revealed in [29]. Instead, it must be fed additional pilot data by the CS to overcome weak SNR by additional processing gain.

Moreover, since the CS only tracks and produces corrections for a single SV at a time, the CT is unable to compute a single-epoch pseudorange-based PNT solution, as is possible with traditional GNSS. Instead, it emulates single-epoch PNT by tracking a sequence of Starlink SVs with an RF front-end driven by a highly stable receiver clock so that it need only estimate a single receiver clock offset and a single clock offset rate, as would be the case for a single-epoch solution.

III. SIGNAL STRUCTURE

This section introduces the Starlink frame structure, drawn from [29], and additional concepts and terminology needed to understand Starlink operation in general and Starlink-based PNT in particular.

A. Frame Structure

As shown in Fig. 1, each Starlink downlink orthogonal frequency division multiplexing (OFDM) frame consists of 302 intervals of length $T_{\text{sym}} = 4.4\mu\text{s}$ plus a frame guard interval T_{fg} , for a total frame period of $T_f = 1/F_f$ s, where $F_f = 750$ Hz is the frame rate. Each frame begins with the primary synchronization sequence (PSS), which is natively represented in the time domain, followed by the secondary synchronization sequence (SSS), which is formatted as a standard 4PSK OFDM symbol. The final occupied symbol interval in each frame is the coda symbol, which is followed by the frame guard interval. Subsequent frames may be present or not, depending on user demand.

The time-domain modulation sequence of the PSS, and the frequency-domain symbol sequence of the SSS were revealed in [29]. These sequences allow a Starlink receiver to perform the channel estimation necessary to demodulate each OFDM symbol in the frame. In particular, the PSS allows the receiver to precisely identify the beginning of each frame, while the SSS allows it to perform equalization across all subcarriers. The PSS and SSS have a fixed phase relationship, which allows coherent integration across the combined PSS + SSS interval, permitting more accurate synchronization and carrier frequency offset (CFO) estimation, as shown in [26].

B. Assigned Beams

At any given time, a user within Starlink's $\pm 53^\circ$ latitude primary coverage area may have a direct line of sight to dozens of overhead SVs above 20° elevation [12]. Each Starlink SV is capable of simultaneously directing up to 48 downlink beams to terrestrial service cells, 16 for each of its three downlink phased arrays. A proprietary beam assignment procedure assigns each beam to a ~ 20 -km-diameter service cell [31]. A service cell may be illuminated by up to 16 beams simultaneously, two beams for each of the eight frequency

channels identified in [29]. We call the N_{bi} beams deliberately directed to the i th service cell *assigned beams*.

The number of distinct SVs casting assigned beams onto the i th service cell, $N_{asi} \leq N_{bi} \leq 16$, may be less than N_{bi} because a given SV may project multiple beams onto the same cell, each on a different channel or with a different polarization (right-hand vs. left-hand circular polarization). As Starlink user density has increased over the past few years, N_{asi} and N_{bi} have generally increased. It was shown in [32] that a search procedure prioritizing high-elevation SVs was effective at finding multiple illuminating SVs.

C. Fixed Assignment Interval

Several publications have noted the existence of a 15-s Starlink network reconfiguration interval [32]–[35]. Because all beam assignments remain fixed over the duration of this interval, we call it the *fixed assignment interval* (FAI). FAI boundaries appear to be approximately aligned with GPST [32]. As discussed in [26], noticeable transitions in Starlink frame timing behavior occur at FAI boundaries.

IV. CLOCK MODELS

This section introduces the Starlink clock models described in [26]. These models will be necessary to derive the clock corrections provided by the CS in Section VI.

A. Beam-Specific Clock Cascade Model

Fig. 2 presents a clock cascade model for a single Starlink downlink beam. The model is beam-specific because, as shown in [26], frame timing behavior may differ from beam to beam over the same FAI for the same SV. This remarkable observation implies that Starlink signal timing is much different from that of traditional GNSS, in which a single clock governs the whole transmission cascade across all frequencies. To be sure, in traditional GNSS timing offsets may be present between different spreading codes on the same carrier and across multiple carriers, resulting in so-called differential code biases [36]. But the code rates as transmitted are all a constant multiple of a single base clock's reference frequency. By contrast, the frame sequences on different beams from the same Starlink SV may differ in both time offset and rate, a fact with significant implications for Starlink-based PNT.

The root of each clock cascade is a base oscillator. As shown in [26], this oscillator is shared across all beams from a given SV. An onboard GNSS receiver driven by the base oscillator produces a clock disciplining signal s_d that corrects the base oscillator towards GPST. Together, the base oscillator and the GNSS receiver form a closed feedback loop that produces the GNSS-disciplined timing signal t_d .

B. Frame and Carrier Clocks

The GNSS-disciplined timing signal t_d drives the frame and carrier clocks. In turn, the frame clock's signal t_f governs the timing of frames transmitted by the SV, and the carrier clock's signal t_c drives the carrier onto which the information-bearing




Fig. 1: Frame layout for the Ku-band Starlink downlink along time-frequency dimensions, from [29]. Indices along the horizontal axis enumerate the 303 intervals that constitute a single frame. The quantity T is the useful (non-cyclic) OFDM symbol interval, T_g is the symbol guard interval (cyclic prefix), and F is the subcarrier spacing. Other quantities are defined in the text.




Fig. 2: Clock cascade model for a single Starlink downlink beam, from [26].

frames are modulated. The frame clock is likely a software-based clock whose output t_f depends on when baseband frames are loaded into buffers for mixing to RF and subsequent transmission. The carrier clock is likely transparent, meaning that $t_c = t_d$, but it is represented in a manner identical to the frame clock for full model generality.

Note that in traditional GNSS the clock driving each SV's spreading code (analogous to the frame clock) also drives the underlying carrier that the spreading code modulates [37]. Thus, the code and the carrier—as transmitted—are locked together such that the code chipping rate is a constant rational multiple of the carrier frequency. In the case of Starlink downlink signals, the frame and carrier clocks operate somewhat independently, which is why they are represented separately in Fig. 2 [26].

We represent frame and carrier clock offsets from t_d as Δt_f and Δt_c . These are related to t_d , t_f , and t_c by

$$t_d(t) = t_f(t) - \Delta t_f(t) \quad (1)$$

$$t_d(t) = t_c(t) - \Delta t_c(t) \quad (2)$$

where t represents true time, or time according to an ideal clock, such as is closely realized by GPST [38]. In this paper, true time and GPST are taken to be synonymous. The frame




Fig. 3: Frame sequence timing diagram showing the transition from the $(l-1)$ th FAI to the l th FAI.

and carrier clocks are related to t by

$$t = t_f(t) - \delta t_f(t) \quad (3)$$

$$t = t_c(t) - \delta t_c(t) \quad (4)$$

where $\delta t_f(t)$ and $\delta t_c(t)$ are the frame and carrier clock offsets.

The time derivatives of $\delta t_f(t)$ and of $\delta t_c(t)$ are called the frame and carrier clock drift. They are equivalent to the instantaneous fractional frequency deviation, written generically as $y(t)$, on which clock stability analysis is based [39].

C. Discrete-Time Frame Clock

The frame sequence timing diagram in Fig. 3 offers further details about the frame clock. Each frame as transmitted has a duration of exactly T_f according to the frame clock. Within each FAI, the frame slot index increments from 0 to $N_a - 1$, with $N_a = 11250$ being the number of frame slots in one FAI. Each FAI starts at the beginning of frame slot 0 and lasts $N_a T_f = 15$ s. The interval of unoccupied frame slots at the beginning of the l th FAI, called the FAI guard interval $T_{ag}(l) = N_{ag}(l)T_f$, spans a variable number of frame slots $N_{ag}(l)$. Note that, for any FAI index l , frame slot $N_{ag}(l)$ is occupied by definition, but other frame slots may not be occupied.

Let $t_f(i, l, m)$ be the frame clock time for the i th Starlink SV at the instant when the frame in the m th frame slot of the l th FAI begins to pass through the phase center of the SV's downlink antenna, where l and m are zero-based indices. By definition, we take this to be

$$t_f(i, l, m) \triangleq 15l + mT_f \quad (5)$$

Let $t^*(i, l, m)$ and $\delta t_f(i, l, m)$ be the corresponding GPST and frame clock offset. Then

$$t^*(i, l, m) = t_f(i, l, m) - \delta t_f(i, l, m) \quad (6)$$

D. Receiver Clock

Now consider the clock of a receiver (e.g., the CT) tracking signals from the Starlink downlink. The receiver clock time t_r is related to true time t by

$$t = t_r(t) - \delta t_r(t_r) \quad (7)$$

The receiver clock offset $\delta t_r(t_r)$ is represented as a function of t_r because it is natively ordered in receiver time in the course of solving for a position and time solution. The time derivative of δt_r with respect to t , denoted $\dot{\delta t}_r(t_r(t))$, is called the receiver clock drift.

Let $t_r(i, l, m)$ be the time of reception, according to the receiver clock, of the frame transmitted by the i th SV at true time $t^*(i, l, m)$. Let $\delta t_r(i, l, m)$ be the corresponding receiver clock offset and $t_*(i, l, m)$ be the corresponding true time of reception. More precisely, $t_r(i, l, m)$ is the receiver clock time at which the frame transmitted at true time $t^*(l, m)$ from the i th SV's downlink antenna's phase center first reached the receiver antenna's phase center. The receipt time $t_r(i, l, m)$ can be related to $t_*(i, l, m)$, $t^*(i, l, m)$, and $t_f(i, l, m)$ by

$$t_*(i, l, m) = t_r(i, l, m) - \delta t_r(i, l, m) \quad (8)$$

$$t^*(i, l, m) = t_r(i, l, m) - \delta t_r(i, l, m) - \delta t_{\text{tof}}(i, l, m) \quad (9)$$

$$t_f(i, l, m) = t_r(i, l, m) - \delta t_r(i, l, m) - \delta t_{\text{tof}}(i, l, m) + \delta t_f(i, l, m) \quad (10)$$

where $\delta t_{\text{tof}}(i, l, m)$ is the frame's time of flight from transmission to reception, expressed as an interval in true time.

V. DATA AIDING

The inclusion of additional pilot sequences in the local replica generated by the CT improves both detection and TOA measurement performance. This section discusses three data-aiding strategies that affect the CT's processing gain: (1) no data aiding: the CT only uses the PSS and SSS (PSS+SSS); (2) using the PSS, SSS, and optimally allocated subcarrier pilots, according to the ZZB (Optimal Pilots); and (3) using the entire data frame within the CT's front-end bandwidth as the local replica, either decoded by the CT or allocated by the CS (Full Frame).

A. PSS + SSS

Given that the CT can detect the Starlink signal using a local replica of the PSS + SSS alone, TOA measurements obtained this way should provide reasonable accuracy. In [26], the authors show that sub-centimeter accuracy is theoretically achievable at moderate SNR with full bandwidth. Since our implemented CT is limited to only 25 MHz of bandwidth, a new calculation is necessary. A lower bound on TOA error variance for a signal with Fourier transform $S(f)$, energy E , and noise spectral density N_0 can be expressed by [40, Eq.11.15]

$$\sigma_\tau^2 \geq \frac{1}{2\gamma^2 \frac{E}{N_0}} \quad (11)$$

where γ^2 is the effective squared bandwidth defined as

$$\gamma^2 \triangleq \frac{\int_{-\infty}^{\infty} (2\pi f)^2 |S(f)|^2 df}{\int_{-\infty}^{\infty} |S(f)|^2 df}. \quad (12)$$

The PSS signal given in [29, Eq. 34] has a Fourier transform

$$S_{\text{PSS}}(f) \triangleq \frac{1}{F_s} \sum_{k=-N_g}^{N-1} \text{rect}\left(\frac{f}{F_s}\right) \exp\left(-j2\pi(k + N_g)\frac{f}{F_s}\right) p_k \quad (13)$$

and, when sampled at 25 MHz, an effective bandwidth of $\gamma_{\text{PSS}} = 3.9977 \times 10^7$ Hz. Similarly, the SSS signal given in [29, Eq. 1.37] has a Fourier transform

$$S_{\text{SSS}}(f) \triangleq \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} X_{m1k} \exp(-j2\pi F T_g k) G_s(f - Fk) \quad (14)$$

$$G_s(f) \triangleq \mathcal{F}[g_s(t)] = T_{\text{sym}} \text{sinc}(T_{\text{sym}} f) \exp(-j\pi T_{\text{sym}} f) \quad (15)$$

and, when sampled at 25 MHz, an effective bandwidth of $\gamma_{\text{SSS}} = 4.6063 \times 10^7$ Hz. The PSS and SSS signal combined in a single frame have a Fourier transform of $S_1(f) = S_{\text{PSS}}(f) + \exp(-j2\pi T_{\text{sym}} f) S_{\text{SSS}}(f)$ and when sampled at 25 MHz, an effective bandwidth of $\gamma_{\text{PSS+SSS}} = 4.3283 \times 10^7$ Hz. Importantly, if the PSS + SSS can be used to form a TOA measurement, the network would incur no loss of data capacity.

B. Optimal Pilots

Due to the low received SNR at the CT, the CS may transmit demodulated data from the Starlink frame to the CT. The CT can then correlate these data resources in addition to the already-known synchronization sequences, allowing for a more accurate TOA estimate than would be achievable with the PSS and SSS alone. This technique can be considered as a form of data aiding in our experiment. In practice, however, the Starlink network would designate these data resources as known pilots, and the CT would have *a priori* knowledge of them.

In this scheme, it is critical that the overhead required to transmit this data to the CT is minimized, ensuring reliable and low-latency data transmission. To achieve this, the CT




Fig. 4: The optimized selections of 12 out of 106 available subcarriers for data aiding as a function of pre-integration SNR.

receives only a small selection of demodulated data resources corresponding to a sparse subset of subcarriers that span the CT's available bandwidth. However, the specific selection of subcarriers will affect the TOA autocorrelation function. If the CT is provided with a suboptimal allocation, its autocorrelation function may exhibit ambiguities or high sidelobe levels, leading to misdetections and increased TOA estimation errors. Conversely, if the CT is provided an optimal allocation, the autocorrelation function will exhibit a sharp mainlobe and adequately suppressed sidelobes, improving the TOA accuracy for the given SNR.

Using the noncoherent integer-constrained ZZB-minimization techniques from [14], an optimized set of subcarrier selections was generated for pre-integration SNRs ranging from -20 dB to 20 dB. In this minimization, the ZZB assumed an *a priori* TOA window with a duration equal to that of the cyclic prefix. The returned allocations are near-optimal, achieving a ZZB TOA variance within 5% of the true optimal selection. Each selection consists of 12 out of 106 available subcarriers within the CT's 25 MHz bandwidth. These subcarrier selections are shown in Fig. 4. At low SNR, the selection is more evenly distributed across the available bandwidth, while at high SNR, the selection favors the extremities of the band, thereby maximizing the mean-squared bandwidth.

The CS can then select subcarriers based on the CT's SNR, demodulate all data resources in the Starlink frame corresponding to these subcarriers, and transmit the decoded data to the CT as though they were pilots sent by the network. If the Starlink network were to implement this, the transmitting SV would need to sacrifice 12 subcarriers per OFDM symbol in each frame. Given that each Starlink OFDM symbol consists of 1024 subcarriers, the absolute data rate loss for the entire frame would be 1.17%.

To compute a lower bound on TOA measurement errors, the Fourier transform of the entire frame at a given post-integration SNR can be described by $S_0 = S_1(f) +$

$\sum_{k=2}^{300} \exp(-j2\pi kT_{\text{sym}}f) S_S(f)$, where $S_S(f)$ represents the optimal subcarrier allocation for a given symbol at a specific SNR. Assuming that the SNR remains constant over the course of a frame, we can compute S_0 for each SNR, determine the effective bandwidth, and calculate the bound on TOA accuracy.

It should be noted that these allocations are optimal on a per-symbol basis. In other words, a TOA measurement computed using a single optimal subcarrier allocation across the entire frame may not necessarily be optimal. For the purposes of this paper, we assume that the difference will be marginal. However, we acknowledge that the potential for an optimal allocation throughout the entire frame warrants further investigation in future work.

C. Full Frame

Finally, the CT may use the entire frame it has access to in order to form a TOA measurement. This could happen in two ways. In the first, the network provider may choose to dedicate the entire bandwidth of the CT to pilot subcarriers. This would mean that the CT would be designated as PNT-only, as there would no longer be data to decode. In this case, the SV would allocate 106 subcarriers per symbol per frame, assuming a CT with 25 MHz bandwidth. The absolute data rate loss incurred by this allocation would be 10.35%.

Alternatively, if the CT were not PNT-only, but rather had sufficient gain to permit decoding data, it could use the decoded data as a local replica to form a TOA measurement as described in [15]. This *decision-directed* scheme would incur no cost to the network, but would require a more complex CT. Additionally, any decoding errors made by the CT would affect TOA measurement accuracy.

To compute a lower bound on TOA measurement errors, we assume that the entire frame within the CT's front-end bandwidth is transmitted as pilots. The Fourier transform of the entire frame can be described by $S_F(f) = S_1(f) + \frac{1}{F_s} \sum_{k=2}^{300} \exp(-j2\pi kT_{\text{sym}}f) \text{rect}\left(\frac{f}{F_s}\right)$, assuming that the Fourier transform of the remaining symbols in the frame is rectangular. The effective bandwidth of the entire frame under this assumption, given a limited bandwidth of 25 MHz is $\gamma_F = 4.6434 \times 10^7$ Hz.

Fig. 5 plots root mean squared error (RMSE) lower bounds for the three paradigms discussed in this section over the range of -20 dB to 20 dB. As expected, the PSS + SSS provides the worst TOA accuracy but still achieves sub-meter accuracy at relatively low SNR. Using the optimal pilot allocations results in a significant performance improvement, theoretically reaching sub-centimeter accuracy at high SNR. The use of the full frame offers the best accuracy, theoretically reaching sub-centimeter accuracy approximately 10 dB earlier than the optimal pilot configuration.

VI. SATELLITE CLOCK AND EPHEMERIS CORRECTIONS

This section introduces the pseudorange measurement model and identifies where errors due to the ephemeris and transmitter clock may arise. Additionally, we explore how these errors are mitigated in practice and in our specific experiment.




Fig. 5: TOA RMSE bounds for three levels of data aiding to a CT having a front-end bandwidth of 25 MHz. Using the PSS+SSS alone incurs no loss of data capacity, while employing 12 optimal pilot subcarriers incurs a 1.17 % loss of data capacity, and the use of the entire 25 MHz band would incur a 10.35 % loss of data capacity.

A. Pseudorange Measurement Model

In a fused LEO GNSS based on Starlink, a CT would apply (5) to obtain the time of transmission (TOT), according to the i th SV's frame clock, of the m th frame in the l th FAI. For its part, the network would send the frame clock offset $\delta t_f(i, l, m)$ so that the CT could obtain the corresponding TOT in GPST, as in (6). Likewise, for our experiment, the frame clock offset $\delta t_f(i, l, m)$ is estimated by the CS for provision to the CT, as described below.

Rearranging (10) to isolate $\delta t_f(i, l, m)$ yields

$$\delta t_f(i, l, m) = t_f(i, l, m) - t_r(i, l, m) + \delta t_r(i, l, m) + \delta t_{\text{tof}}(i, l, m) \quad (16)$$

If the CS can calculate, measure, or model all the quantities on the right-hand side of this equation, it will obtain an estimate of $\delta t_f(i, l, m)$ that it can pass to the CT. To obtain $t_f(i, l, m)$, it calculates nominal frame timing according to (5). Note that there is no ambiguity in this calculation because, as explained in [26], the CS can mark the end of a FAI by observation and count frame slots within the subsequent FAI.

The CS obtains $t_r(i, l, m)$ by measurement. Let $\mathcal{M}_{il}$ be the set of occupied-frame indices for the l th FAI of the i th SV. A sequence of highly precise frame TOA measurements corresponding to $\mathcal{M}_{il}$ may be extracted by the CS using the maximum-likelihood (ML) methods described in [41]. For $m \in \mathcal{M}_{il}$, these are modeled as

$$\tilde{t}_r(i, l, m) = t_r(i, l, m) + w_r(i, l, m) \quad (17)$$

where $w_r(i, l, m)$ is zero-mean Gaussian measurement error with variance $\sigma_r^2(i, l, m)$. As explained in [41], if the CS has access to the full 240-MHz channel bandwidth and an antenna with sufficient gain to provide a pre-correlation SNR of 14 dB or more, then ML TOA estimation over a full frame is remarkably precise, with range-equivalent RMSE $c\sigma_r(i, l, m)$

less than one centimeter, where c is the speed of light in vacuum.

To obtain $\delta t_r(i, l, m)$, the CS employs simultaneous Starlink and GNSS signal capture, accounting for all cable delays, as described in [26]. Finally, the CS models $\delta t_{\text{tof}}(i, l, m)$ using meter-accurate SV ephemeris models and atmospheric models, as also described in [26].

For this paper's demonstration, meter-accurate ephemerides based on observables from a GNSS receiver onboard the Starlink SVs were not available as was the case in [26]. Instead, we were obliged to use publicly accessible Starlink ephemerides, either from space-track.org, whose ephemerides are accurate (at least periodically) to within about 10 m [42], or from NORAD-provided TLEs, which have much larger errors. Such errors caused significant errors in our modeled $\delta t_{\text{tof}}(i, l, m)$. Thus, the estimated $\delta t_f(i, l, m)$ sent from the CS to the CT, denoted $\bar{\delta}t_f(i, l, m)$, is modeled as

$$\bar{\delta}t_f(i, l, m) = \delta t_f(i, l, m) + \eta(i, l, m)/c + w_r(i, l, m) \quad (18)$$

where $\eta(i, l, m)$, given in meters and discussed further in the following subsection, is the error in $\bar{\delta}t_f(i, l, m)$ due to ephemeris errors.

We now turn our attention to the CT, which calculates TOA just as the CS does but with a possibly reduced bandwidth and, in the case of a PNT-only CT, with ML estimation based on additional CT-provided pilot data, which provides the ML estimator a boost in processing gain to compensate for the CT's low-gain antenna. Let $\tilde{t}_u(i, l, m)$ denote the CT's measured frame TOA for the corresponding i, l , and $m \in \mathcal{M}_{il}$, which is modeled as

$$\tilde{t}_u(i, l, m) = t_u(i, l, m) + w_u(i, l, m) \quad (19)$$

where $t_u(i, l, m)$ is the true TOA according to the CT's clock, and $w_u(i, l, m)$ is zero-mean Gaussian measurement error with variance $\sigma_u^2(i, l, m) \geq \sigma_r^2(i, l, m)$. The CT's pseudorange measurement, defined as

$$\rho(i, l, m) = c [\tilde{t}_u(i, l, m) - t_f(i, l, m)] \quad (20)$$

may then be modeled as

$$\rho(i, l, m) = r_u(i, l, m) + c\delta t_u(i, l, m) + \eta(i, l, m) + c\delta t_{\text{atm}}(i, l, m) + w(i, l, m) \quad (21)$$

where $r_u(i, l, m)$ is the true range between the i th Starlink SV and the CT, $\delta t_u(i, l, m)$ is the CT clock error, $\delta t_{\text{atm}}(i, l, m)$ denotes atmospheric delays along the line-of-sight path, and $w(i, l, m)$ is zero-mean Gaussian measurement error with variance $c^2 [\sigma_r^2(i, l, m) + \sigma_u^2(i, l, m)]$.

B. Ephemeris Error Modeling

If Starlink were providing PNT services as a fused LEO GNSS, the ephemeris errors would likely be low as the satellite ephemeris could be packaged in a navigation message. However, during our experiment, only TLEs were available. As a result, it is important to discuss the differential errors caused by the large ephemeris errors inherent in using TLEs.

At any given epoch, the position of a Starlink SV can be retrieved from publicly available TLEs. The accuracy of the

SV position derived from a TLE is typically accurate to the kilometer-level. However, after 24 hours since the TLE was last updated, these spatial errors can grow to be 1 to 10 km [43]. This subsection aims to characterize how these spatial ephemeris errors affect the differential pseudorange error when measured at the CS and CT. For simplicity, we consider a single satellite, a single FAI, and a single frame, and thus omit the corresponding notation.

Let $\mathbf{p}_s$ denote the true position of the SV, $\mathbf{p}_r$ the known position of the CS, and $\mathbf{p}_u$ the position of the CT. The true ranges between $\mathbf{p}_r$ and $\mathbf{p}_s$, and $\mathbf{p}_u$ and $\mathbf{p}_s$, are denoted r_r and r_u , respectively. These ranges are expressed as

$$r_r = \sqrt{(\mathbf{p}_r - \mathbf{p}_s)^\top (\mathbf{p}_r - \mathbf{p}_s)} \quad (22)$$

$$r_u = \sqrt{(\mathbf{p}_u - \mathbf{p}_s)^\top (\mathbf{p}_u - \mathbf{p}_s)} \quad (23)$$

Now, suppose that the assumed SV position contained a 3D spatial error ϵ . The resulting ranges, when assuming the erroneous SV position at the CS, $\tilde{r}_r$, and at the CT, $\tilde{r}_u$, are expressed as

$$\tilde{r}_r = H_r \epsilon + r_r \quad (24)$$

$$\tilde{r}_u = H_u \epsilon + r_u \quad (25)$$

where H_r and H_u are the 1×3 Jacobians with respect to the SV's position, expressed as

$$H_r = \frac{\partial r_r}{\partial \mathbf{p}_s} = -\hat{\mathbf{p}}_r = -\frac{\mathbf{p}_r - \mathbf{p}_s}{r_r} \quad (26)$$

$$H_u = \frac{\partial r_u}{\partial \mathbf{p}_s} = -\hat{\mathbf{p}}_u = -\frac{\mathbf{p}_u - \mathbf{p}_s}{r_u} \quad (27)$$

The pseudorange measurement error at both the CS and the CT due to the spatial errors in the SV's position will manifest differently due to differences in geometry. However, if the positions of the CS and CT are close to each other, the resulting errors will be relatively small. The relative differential pseudorange error between the CS and the CT, caused by spatial ephemeris errors, is denoted as η , and is expressed as

$$\begin{aligned} \eta &= (\tilde{r}_r - r_r) - (\tilde{r}_u - r_u) \\ &= (H_r - H_u) \epsilon \end{aligned} \quad (28)$$

For example, consider a specific $\mathbf{p}_r$, $\mathbf{p}_u$, $\mathbf{p}_s$, and 3D satellite spatial error $\epsilon \sim \mathcal{N}(\mathbf{0}, P_\epsilon)$, then

$$\eta \sim \mathcal{N}(0, \sigma_\eta^2) \quad (29)$$

with

$$\sigma_\eta^2 = (H_r - H_u) P_\epsilon (H_r - H_u)^\top \quad (30)$$

The distribution of η in (29) requires the knowledge of $\mathbf{p}_u$ to compute the Jacobian which is required for σ_η^2 . In practice, $\mathbf{p}_u$ is unknown, but the CT will have *a priori* information about $\mathbf{p}_u$ relative to $\mathbf{p}_r$. For instance, the CT might know that it is located within the same beam footprint as the CS. In this case, $\mathbf{p}_u$ is modeled as

$$\mathbf{p}_u = \mathbf{p}_r + \boldsymbol{\xi} \quad (31)$$

where $\boldsymbol{\xi}$ is a 3D random offset with respect to $\mathbf{p}_r$, following some distribution with covariance P_ξ . Since $\mathbf{p}_u$ is now treated as a random variable, $H_u(\boldsymbol{\xi})$ and $\sigma_\eta^2(\boldsymbol{\xi})$ also become random variables. $H_u(\boldsymbol{\xi})$ is expressed as

$$H_u(\boldsymbol{\xi}) = H_r + \boldsymbol{\xi}^\top A \quad (32)$$

with

$$A^\top = \frac{\partial H_r^\top}{\partial \mathbf{p}_r} = r_r^{-1} (\hat{\mathbf{p}}_r \hat{\mathbf{p}}_r^\top - \mathbb{I}) \quad (33)$$

where A is the Jacobian of H_r with respect to $\mathbf{p}_r$. The derivation of this Jacobian can be found in [44]. Consequently, (30) is reduced to

$$\sigma_\eta^2(\boldsymbol{\xi}) = \boldsymbol{\xi}^\top A P_\epsilon A^\top \boldsymbol{\xi} \quad (34)$$

Now, $\sigma_\eta^2(\boldsymbol{\xi})$ is a function of the unknown random vector $\boldsymbol{\xi}$, so solving for $\bar{\sigma}_\eta^2$ is required for modeling purposes. Note, (34) is a quadratic form, thus

$$\bar{\sigma}_\eta^2 = \mathbb{E} [\sigma_\eta^2(\boldsymbol{\xi})] = \text{trace} (A P_\epsilon A^\top P_\xi) \quad (35)$$

Finally, the distribution of η when $\mathbf{p}_u$ is unknown can be modeled as

$$\eta \sim \mathcal{N}(0, \bar{\sigma}_\eta^2). \quad (36)$$

Consider the scenario where the ephemeris error is zero-mean Gaussian in the ENU frame, with covariance matrix $P_\epsilon = \sigma_{\text{eph}}^2 \mathbb{I}$. Also, assume that $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, P_\xi)$, where $P_\xi = \text{diag}(\sigma_e^2, \sigma_n^2, \sigma_u^2)$, centered at the CS. The choices of σ_{eph}^2 , σ_e^2 , σ_n^2 , and σ_u^2 are treated as tuning parameters. The values of σ_e^2 , σ_n^2 , and σ_u^2 depend on the assumed distance between the CS and CT.

In the case of Starlink, it is required that the CS and CT occupy the same service cell. Following this, the CT's position is modeled as Gaussian with $P_\xi = \text{diag}(3300^2, 3300^2, 50^2)$ centered at the CS, which puts the CT inside of a service cell 95% of the time. In this paper, σ_{eph}^2 is scaled with the age of TLE. For a TLE that is 12 hours and 24 hours old, choices of $\sigma_{\text{eph}} = 2.5$ km, and $\sigma_{\text{eph}} = 5$ km are conservatively used, respectively. With a typical Starlink satellite to service cell geometry, $\bar{\sigma}_\eta = 17.8$ m for a 12-hour-old TLE, and $\bar{\sigma}_\eta = 35.6$ m for a 24-hour-old TLE. If a CT had access to high-quality ephemeris data with $\sigma_{\text{eph}} = 10$ m, then $\bar{\sigma}_\eta = 0.07$ m, which would be negligible. Since the CS and CT are within the same service cell, a small ephemeris error would affect the pseudoranges at the CS and CT in a similar manner. The distribution of η will be further exploited in Section VII.

VII. ESTIMATION

This section uses the measurement model derived in VI to develop a NBLS estimator for obtaining the PNT solution of the CT. The NBLS estimator introduced here is similar to the traditional GNSS NBLS estimator, with additional constraints tailored to the specific requirements of our experiment.

A. Estimator Formulation

In practice, the CT would be expected to compute an instantaneous PNT solution using pseudoranges from multiple SVs all received at nearly the same time, with little ephemeris error. This estimator could be similar to what is used in traditional GNSS. The estimator developed in this section can use measurements from multiple SVs at a time, but these measurements will not be received simultaneously and, in some cases, may be spaced minutes apart. Additionally, they will also suffer from large ephemeris errors since only TLEs were available at the time of the experiment. Appropriate design changes are made to compensate for these differences. We begin by making the simplification that, for our experiment, only one FAI is used for each satellite. Therefore, we drop the FAI index from here onward. This also implies that the set $\mathcal{M}_{il} = \mathcal{M}_i$. Then, we develop the NBLS estimator by rewriting the pseudorange measurement model in (21) as

$$\rho(i, m) = r_u(i, m) + c[\delta t_u(0, 0) + \Delta t(i, m)\delta \dot{t}_u] + \eta(i) + w(i, m) \quad (37)$$

where $\Delta t(i, m)$ is the time separation between the first SV's first frame and the i th SV's m th frame in seconds, $\delta \dot{t}_u$ is the CT's clock drift, which is assumed to be constant over the course of the experiment, and $\eta(i)$ is as described in Section VI. Here, we simplify by assuming that $\eta(i)$ remains constant over all FAIs and frames for a given satellite. Then, we can define our measurement vector $\mathbf{z}$ for n SVs as

$$\mathbf{z} = \begin{bmatrix} \rho(0) \\ \rho(1) \\ \vdots \\ \rho(n-1) \end{bmatrix}, \quad \boldsymbol{\rho}(i) = \begin{bmatrix} \rho(i, 0) \\ \rho(i, 1) \\ \vdots \end{bmatrix} \quad (38)$$

where the length of $\boldsymbol{\rho}(i)$ is the cardinality of the set of occupied frames used to determine pseudorange measurements during the i th SV's FAI, denoted by $\mathcal{M}_i$.

Next, we define the state vector. Given the nature of the current experiment and the gravity of the ephemeris errors present in $\mathbf{z}$, it is necessary to estimate $\eta(i)$ for each SV i . However, to limit the number of total states to be estimated, we combine the CT's clock offset with $\eta(0)$ as $\delta t_\eta = c\delta t_u(0, 0) + \eta(0)$. We also define $\delta \eta(i) = \eta(i) - \eta(0)$. The measurement model may then be reformulated as

$$\rho(i, m) = r_u(i, m) + \delta t_\eta + \delta \eta(i) + c\Delta t(i, m)\delta \dot{t}_u + w(i, m) \quad (39)$$

Where, by inspection, $\delta \eta(0) = 0$. Then, the state vector $\mathbf{x}$ is given by

$$\mathbf{x} = [\mathbf{p}_u^\top \quad \delta t_\eta \quad \delta \boldsymbol{\eta}^\top \quad c\delta \dot{t}_u]^\top \quad (40)$$

where $\mathbf{p}_u \in \mathbb{R}^3$ are the position states of the CT in the Earth-centered, Earth-fixed coordinate system and $\delta \boldsymbol{\eta} \in \mathbb{R}^{n-1}$ is the vector of all $\delta \eta(i), i \neq 0$. Finally, we can define the design matrix H for a single row as

$$H(i, m) = \begin{bmatrix} \frac{\partial \rho(i, m)}{\partial \mathbf{p}_u} & \frac{\partial \rho(i, m)}{\partial \delta t_\eta} & \frac{\partial \rho(i, m)}{\partial \delta \boldsymbol{\eta}} & \frac{\partial \rho(i, m)}{\partial c\delta \dot{t}_u} \end{bmatrix} \quad (41)$$

$$= [\hat{\mathbf{p}}(i, m) \quad 1 \quad \mathbf{1}_i \quad \Delta t(i, m)]$$

Where $\mathbf{1}_i \in \mathbb{R}^{n-1}$ is a vector with 1 in the i th position and zeros elsewhere. We also make the important distinction that $\mathbf{1}_0 = \mathbf{0} \in \mathbb{R}^{n-1}$, or equivalently, $\mathbf{1}_0$ is the zero vector.

B. Measurement Covariance

If all measurements were independent and identically distributed (IID), we could proceed directly to form the estimate $\hat{\mathbf{x}}$ with ordinary least squares. However, there are two primary factors that affect the distribution of the measurements, preventing them from being IID. First, the received SNR of the i th SV's m th frame may differ from SV to SV and frame to frame at both the CS and CT, which influences the variance of $w(i, m)$ over time. Second, the shared CT clock parameters between measurements introduce correlation.

To begin, we first assume that the SNR for the i th satellite remains relatively constant for the entire FAI, as observed at both the CS and the CT. Therefore, we can define $w(i, m) \approx w(i) \sim \mathcal{N}(0, \sigma^2(i)), m \in \mathcal{M}_i$. The measurement covariance matrix of the i th SV $R_a(i) = \sigma_\rho^2(i)\mathbb{I}_i$, where $\mathbb{I}_i$ is the identity matrix with size equal to the cardinality of $\mathcal{M}_i$. Consequently, the measurement covariance matrix for all SVs, due to TOA estimation error is then the block diagonal matrix

$$R_a = \begin{bmatrix} R_a(0) & 0 & \dots & 0 \\ 0 & R_a(1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & R_a(n-1) \end{bmatrix} \quad (42)$$

where each 0 is an appropriately sized matrix of all zeros.

To address the clock errors, we first point to the original measurement model in (21). Dropping the FAI index, the CT clock offset term $\delta t_u(i, m)$ is a random variable, which is modeled as a combination of a Wiener process and an integrated Wiener process [45]. It's instantaneous variance is given by $\sigma_{\delta t_u}^2(i, m) = S_f \Delta t(i, m) + \frac{S_g \Delta t^3(i, m)}{3}$ in [46], where $S_f = \frac{h_0}{2}$ and $S_g = 2\pi^2 h_{-2}$ are the power spectral densities of the phase and frequency random processes of the receiver clock, respectively. Typically, the coefficients h_0 and h_{-2} are determined through laboratory experiments.

The measurement covariance matrix introduced by the clock offset random process is determined as follows. For the receiver clock offset at the time of the i th and j th SV's m th and k th frames, respectively the covariance is

$$\mathbb{E}[\delta t_u(i, m)\delta t_u(j, k)] = S_f \Delta t_m + \frac{S_g \Delta t_m^3}{3} \quad (43)$$

where $\Delta t_m = \min(\Delta t(i, m), \Delta t(j, k))$

Therefore, we may define each element of the measurement covariance due to CT clock errors R_{clk} as in (43). As a simple example, consider the 3×3 case where three TOA measurements are extracted from three frames from two SVs. Then, R_{clk} would be

$$R_{\text{clk}} = \begin{bmatrix} \sigma_{\delta t_u}^2(0, 0) & \sigma_{\delta t_u}^2(0, 0) & \sigma_{\delta t_u}^2(0, 0) \\ \sigma_{\delta t_u}^2(0, 0) & \sigma_{\delta t_u}^2(1, 0) & \sigma_{\delta t_u}^2(1, 0) \\ \sigma_{\delta t_u}^2(0, 0) & \sigma_{\delta t_u}^2(1, 0) & \sigma_{\delta t_u}^2(1, 1) \end{bmatrix} \quad (44)$$

From there, it is straightforward to extrapolate R_{clk} for multiple TOA measurements from various satellites and frames.

Then, the final measurement covariance matrix R is formed as

$$R = R_a + c^2 R_{\text{clk}} \quad (45)$$

C. Maximum a Posteriori Least Squares Update

At this point, the weighted NBLS estimator would allow the states $\delta\eta$ to explore any value. However, since we have *a priori* information about the distributions of $\delta\eta$, we can incorporate this information within the NBLS update step as follows. First, consider the traditional weighted NBLS update step at the k th iteration with residuals $\delta\mathbf{z}$

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_{k-1} + (H^\top R^{-1} H)^{-1} H^\top R^{-1} \delta\mathbf{z} \quad (46)$$

We may apply the *a priori* distribution of $\hat{\mathbf{x}}_0$ to the update, assuming that $\hat{\mathbf{x}}_0 \sim \mathcal{N}(\mathbf{x}_0, P)$. The update step then becomes

$$\hat{\mathbf{x}}_k = \hat{\mathbf{x}}_{k-1} + (H^\top R^{-1} H + P^{-1})^{-1} (H^\top R^{-1} \delta\mathbf{z} + P^{-1} \delta\mathbf{x}) \quad (47)$$

where $\delta\mathbf{x} = \hat{\mathbf{x}}_{k-1} - \mathbf{x}_0$, and $\mathbf{x}_0$ is the initial guess of the iterative NBLS process. Assuming no prior knowledge of the position, clock offset, or clock drift states, the *a priori* covariance matrix $P \in \mathbb{R}^{5+n-1}$ will be diffuse for the first four states and the final state. Said differently, the first four diagonal elements of P and the last diagonal element of P will be set to ∞ to reflect a complete lack of prior knowledge. The elements corresponding to the $\delta\eta$ states will be set to P_η and the elements corresponding to the covariance between $\delta\eta$ and the diffuse states will be set to 0.

The matrix P_η is defined using the distributions derived in Section VI. To define P_η in this way, we first make the simplification that $\eta(i) \sim \mathcal{N}(0, \bar{\sigma}_\eta^2(i))$, $i \in \mathcal{I}$, where $\mathcal{I}$ is the set of all SVs used in the NBLS estimator. After making this simplification, we can recognize that

$$\delta\eta \sim \mathcal{N}(\mathbf{0}, P_\eta) \quad (48)$$

$$P_\eta = \begin{bmatrix} \bar{\sigma}_\eta^2(0) + \bar{\sigma}_\eta^2(1) & \bar{\sigma}_\eta^2(0) & \cdots & \bar{\sigma}_\eta^2(0) \\ \bar{\sigma}_\eta^2(0) & \bar{\sigma}_\eta^2(0) + \bar{\sigma}_\eta^2(2) & \cdots & \bar{\sigma}_\eta^2(0) \\ \vdots & \vdots & \ddots & \vdots \\ \bar{\sigma}_\eta^2(0) & \bar{\sigma}_\eta^2(0) & \cdots & \bar{\sigma}_\eta^2(0) + \bar{\sigma}_\eta^2(n-1) \end{bmatrix} \quad (49)$$

Then, the initial guess is set to $\mathbf{x}_0 = [\mathbf{p}_0 \ \mathbf{0}]^\top$, where $\mathbf{p}_0$ is some diffuse initial position guess, and $\mathbf{0}$ is the zero vector of length $n+1$. From there, given a set of measurements $\mathbf{z}$ the CT may form a solution using the weighted NBLS estimator by iteratively applying (47) over a predefined number of iterations or until the second quantity in (47) becomes sufficiently small.

VIII. EXPERIMENTAL RESULTS

This section discusses the experiment we devised to test our mock implementation of fused LEO GNSS, including the signal capture setup, the specific scenario, and the results obtained.

A. Signal Capture

To capture signal data, a CS receiver and a CT receiver were set up. The CS was located on the roof of the Aerospace Engineering Building on the University of Texas campus. The CS used a steerable 90-cm offset parabolic dish with an approximately 3 degree beamwidth to receive Starlink signals. The dish was steered to desired SVs using TLEs, with the beamwidth limiting captures to a single SV at a time. The parabolic dish was equipped with a feedhorn connected to a low-noise block (LNB), which has a conversion gain of 60 dB and a noise figure of 0.8 dB. The LNB downconverts 10.7–11.7 GHz signals to 950–1950 MHz, or 11.7–12.75 GHz to 1100–2150 MHz. The antenna's nominal gain is 40 dBi at 12.5 GHz, with at least 4-5 dB of losses due to the absence of a circular-to-linear polarizer and feedhorn misalignment.

The CS also featured a signal capture system capable of continuous sampling at 250 Msps. The usable bandwidth of the capture is around 200 MHz, which was ample to obtain high-quality TOAs from Starlink's $F_s = 240$ MHz Ku-band downlink signal. The hardware then performed downmixing to baseband, bandpass filtering, and 16-bit complex sampling. The capture system can simultaneously capture on two channels, with the limitation that the sampling rate must be identical for both channels to start sampling simultaneously. This setup is referred to as a dual capture. To extract TOTs as described in [26], a dual capture of both Starlink and GNSS data was performed for all captures treated in this paper. The signal capture system and the LNB local oscillators were both driven by an external GPS disciplined oven-controlled crystal oscillator (GPSDO).

The CT was located in the parking lot of the Whitaker tennis courts in Austin, TX, approximately 3.009 km from the CS. The CT used the same feedhorn and LNB as the CS, but did not have a parabolic dish. Instead, the CT faced the feedhorn directly towards zenith to collect data. The CT also featured a signal capture system. It was only capable of continuous sampling at 25 Msps. The hardware performed the same steps as in the CS, including downmixing to baseband, bandpass filtering, and 16-bit complex sampling. Like the CS, the CT's signal capture system can also perform a dual capture. For ground truth, a dual capture of both Starlink and GNSS data was performed at the CT as well. Since the CS and CT were only able to observe one SV at a time, there were often long gaps between observations. To make long gaps appear short, the CT signal capture system's local oscillator was also driven by an external GPSDO, and the LNB's local oscillator was driven by the same GPSDO.

B. Scenario Description

Starlink signal data was recorded over approximately a 13-minute period in January 2025. During this time, the CT was continuously capturing to ensure local clock coherence between SVs. The CS sporadically captured as SVs became observable, capturing data in 60-second periods. Both the CS and the CT captured data at a center frequency of 12.075 GHz, corresponding to channel 6 from [29]. Over this period, 5 SVs were visible. Fig. 6 shows the ground tracks of the recorded




Fig. 6: The ground tracks of the recorded Starlink SVs. In total, 5 SVs were recorded, with identifiers, in order of capture, being 5135, 31868, 3160, 3569, and 31848. The black diamond marks the location of the CT and the black square marks the location of the NBLS initialization.

Starlink SVs. For all 5 SVs, measurements were generated over a 10-second period during an FAI. The gaps between SVs ranged from as few as 100 seconds and as many as 280 seconds.

TOA measurements were made at the CT by using full frames as local replicas. Since Starlink frames change frequently and the modulation of each symbol in each frame also changes, sometimes within a symbol (as shown in [41]), decoding enough frames to generate a significant number of measurements proved too cumbersome. Instead, we opted to directly correlate the CS in-phase and quadrature samples with those of the CT. This process is similar to the *decision-directed* approach described in Section V. We acknowledge that this is a suboptimal method to extract these measurements, and improvements could be made in the future. In total, $N_z = 6384$ TOA measurements were extracted from 5 SVs over 50 seconds of data. The measurements from a single SV are spaced at arbitrary multiples of the frame rate $1/750$ seconds.

For these data, only TLEs were available. Most of the TLEs were 24 hours old, with one being just 12 hours old. We determined P_η from Section VII-C by using the method described in Section VI-B. To do this, we assumed that the position of the CT was Gaussian distributed with $\sigma_e = \sigma_n = 3.3$ km and $\sigma_u = 50$ m around the CS. Since the TLEs were around 24 hours old, $\sigma_{\text{eph}} = 5$ km was used. Given that the CT's local oscillator was driven by a GPSDO, we set $h_0 = 5 \times 10^{-23}$ and $h_{-2} = 3 \times 10^{-26}$ to model an especially good OCXO.

C. Results

The NBLS estimator was set up as in Section VII and initialized in Monterrey, MX, as shown in Fig. 6, approximately 572 km away from the CT. Fig. 7 plots the results of the NBLS using the measurements obtained during our field experiment. At the top, the residuals are shown. The residuals are nearly zero-mean and show the long gaps between SV observations. The final solution is plotted below the residuals. The solution has an absolute 3D error of 6.47 m, with a 95% confidence ellipse that has a semi-major axis of 11.96 m and a semi-minor axis of 6.31 m. The reported altitude has a 95% confidence interval of ± 18.6 m around its solution. Finally, the clock offset state δt_η is resolved within a 95% confidence interval of ± 17.29 m around its solution. Considering that this estimate is the combination of the true clock offset $\delta t_u(0,0)$ and the TLE error $\eta(0)$, we can incorporate the distribution of $\eta(0)$ into the confidence interval for resolving the true clock offset. Assuming a 24 hour old TLE, the confidence interval expands to ± 73.27 m around the true clock offset. Converting to seconds this allows us to resolve clock offset to ± 244.39 ns with 95% confidence.

IX. CONCLUSION

In this paper we presented a method for producing a pseudorange-based PNT solution using Starlink's Ku-band downlink signal. We tested this method in a field campaign and achieved 10-meter-level position and timing accuracy, despite challenges presented by large ephemeris errors, a low-gain receiver antenna, and a narrow-bandwidth front-end. SpaceX could readily implement a full-fledged fused LEO GNSS service using the methods described here. The accuracy of such a service would exceed what we have demonstrated, provided that SpaceX has access to meter-accurate ephemeris data for their Starlink SVs and could actively monitor their SV clocks to within ~ 10 ns. Alternatively, a third-party service could be developed, deploying reference stations at service-cell density to provide timing and ephemeris corrections to users. In either case, we believe that the present work establishes the viability of pseudorange-based LEO PNT as a backup to traditional GNSS in the near term. Looking further ahead, we expect LEO PNT to eventually provide PNT services superior to traditional GNSS.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] O. Kutkov, "Connecting external GPS antenna to the starlink terminal," <https://olegkutkov.me/2023/11/07/connecting-external-gps-antenna-to-the-starlink-terminal/>, 2023, accessed: 2023-03-01.
- [2] T. G. Reid, A. M. Neish, T. Walter, and P. K. Enge, "Broadband LEO constellations for navigation," *NAVIGATION*, vol. 65, no. 2, pp. 205–220, 2018.




Fig. 7: The post-fit residuals from the NBLs estimator (top). The residuals are fairly zero-mean. Also plotted is the NBLs position solution, marked with a white circle, alongside the true location of the CT, marked with a black diamond. The solution has an absolute 3D error of 6.47 m, with a 95% confidence ellipse that has a semi-major axis of 11.96 m and a semi-minor axis of 6.31 m.

[3] F. S. Prol, R. M. Ferre, Z. Saleem, P. Valisuo, C. Pinell, E. S. Lohan, M. Elsanhoury, M. Elmusrati, S. Islam, K. Celikbilek, K. Selvan, J. Yliaho, K. Rutledge, A. Ojala, L. Ferranti, J. Praks, M. Z. H. Bhuiyan, S. Kaasalainen, and H. Kuusniemi, "Position, navigation, and timing (PNT) through low earth orbit (LEO) satellites: A survey on current status, challenges, and opportunities," *IEEE access*, vol. 10, pp. 83 971–84 002, 2022.

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

[5] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, M. L. Psiaki, and T. E. Humphreys, "First results from three years of GNSS interference monitoring from low Earth orbit," *NAVIGATION*, vol. 68, no. 4, pp. 673–685, 2021.

[6] Z. Clements, P. Ellis, and T. E. Humphreys, "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit," in *Proceedings of the IEEE/ION PLANS Meeting*, Monterey, CA, 2023, pp. 458–469.

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

[8] Z. L. Clements, P. B. Ellis, M. J. Murrian, M. L. Psiaki, and T. E. Humphreys, "Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit," *NAVIGATION*, 2025, submitted for review.

[9] W. Van Uytsel, T. Janssen, M. Weyn, and R. Berkvens, "A technical overview of current "new space" LEO-PNT initiatives and their applica-

tion potential," in *2024 IEEE 35th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2024, pp. 1–6.

[10] E. Rubinov, "FrontierSI State of the Market Report LEO PNT 2024 Edition," Jan. 2025, <https://frontiersi.com.au/wp-content/uploads/2025/01/FrontierSI-State-of-Market-Report-LEO-PNT-2024-Edition-v1.1.pdf>.

[11] N. S. Miller, J. T. Koza, S. C. Morgan, S. M. Martin, A. Neish, R. Grayson, and T. Reid, "SNAP: A Xona Space Systems and GPS software-defined receiver," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, 2023, pp. 897–904.

[12] P. A. Iannucci and T. E. Humphreys, "Fused low-Earth-orbit GNSS," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2022.

[13] —, "Economical fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, 2020.

[14] A. M. Graff and T. E. Humphreys, "Ziv-Zakai-optimal OFDM resource allocation for time-of-arrival estimation," *IEEE Transactions on Wireless Communications*, 2025.

[15] —, "OFDM-based positioning with unknown data payloads: Bounds and applications to LEO PNT," *IEEE Transactions on Wireless Communications*, 2024, submitted for review.

[16] Z. M. Kassas, N. Khairallah, and S. Kozhaya, "Ad astra: Simultaneous tracking and navigation with megaconstellation LEO satellites," *IEEE Aerospace and Electronic Systems Magazine*, 2024.

[17] S. Shahcheraghi, J. Saroufim, and Z. M. Kassas, "Acquisition, Doppler tracking, and differential LEO-aided IMU navigation with uncooperative satellites," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.

[18] J. Saroufim and Z. M. Kassas, "Ephemeris and timing error disambiguation enabling precise LEO PNT," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.

[19] S. Zhou, R. Yang, Y. Li, X. Zhan, and H. Qin, "Iridium TOA estimation and positioning based on carrier tracking and beam decoding," *IEEE Transactions on Instrumentation and Measurement*, vol. 74, pp. 1–23, 2025.

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

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

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

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

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

[25] N. Jardak, R. Adam, and Q. Jault, "Leveraging multi-LEO satellite signals for opportunistic positioning," *IEEE Access*, 2024.

[26] W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, "Timing properties of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, 2025, submitted for review.

[27] H. Sallouha, S. Saleh, S. De Bast, Z. Cui, S. Pollin, and H. Wymeersch, "On the ground and in the sky: A tutorial on radio localization in ground-air-space networks," *IEEE Communications Surveys and Tutorials*, vol. 27, no. 1, pp. 218–258, 2025.

[28] A. Baron, P. Gurfil, and H. Rotstein, "Implementation and accuracy of Doppler navigation with LEO satellites," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 2, 2024.

[29] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–16, 2023.

[30] Z. M. Komodromos, W. Qin, and T. E. Humphreys, "Weak signal acquisition and tracking of the Starlink Ku-Band downlink to enable global PNT," in *Proceedings of the ION Joint Navigation Conference (JNC)*, 2024.

[31] S. Pekhterev, "The bandwidth of the Starlink constellation and the assessment of its potential subscriber base in the USA," *SatMagazine*, vol. 11, pp. 54–57, 2021.

[32] 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.

- [33] G. Huston, "A transport protocol's view of Starlink," May 2024, accessed: 2024-11-22. [Online]. Available: <https://blog.apnic.net/2024/05/17/a-transport-protocols-view-of-starlink/>
- [34] N. Mohan, A. E. Ferguson, H. Cech, R. Bose, P. R. Renatin, M. K. Marina, and J. Ott, "A multifaceted look at Starlink performance," in *Proceedings of the ACM on Web Conference 2024*, 2024, pp. 2723–2734.
- [35] E. Grayver, R. Nelson, E. McDonald, E. Sorensen, and S. Romano, "Position and navigation using Starlink," in *2024 IEEE Aerospace Conference*. IEEE, 2024, pp. 1–12.
- [36] O. Montenbruck, A. Hauschild, and P. Steigenberger, "Differential code bias estimation using multi-GNSS observations and global ionosphere maps," *NAVIGATION*, vol. 61, no. 3, pp. 191–201, 2014.
- [37] M. Meurer and F. Antreich, "Signals and modulation," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 91–119.
- [38] C. Jekeli and O. Montenbruck, "Time and reference systems," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 25–58.
- [39] R. Beard and K. Senior, "Clocks," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 121–164.
- [40] M. Skolnik, *Introduction to radar systems 2nd Edition*. McGraw-Hill, 1980.
- [41] W. Qin, Z. M. Komodromos, S. C. Morgan, and T. E. Humphreys, "Maximum likelihood time of arrival and Doppler estimation for precise Starlink-based PNT," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.
- [42] A. Liu, X. Xu, Y. Xiong, and S. Yu, "Maneuver strategies of Starlink satellite based on SpaceX-released ephemeris," *Advances in Space Research*, 2024.
- [43] J. R. Vetter, "Fifty years of orbit determination," *Johns Hopkins APL technical digest*, vol. 27, no. 3, p. 239, 2007.
- [44] Z. Clements, P. Ellis, M. L. Psiaki, and T. E. Humphreys, "Geolocation of terrestrial GNSS spoofing signals from low Earth orbit," in *Proceedings of the ION GNSS+ Meeting*, Denver, CO, 2022, pp. 3418–3431.
- [45] L. Galleani, "A tutorial on the two-state model of the atomic clock noise," *Metrologia*, vol. 45, no. 6, p. S175, 2008.
- [46] R. G. Brown and P. Y. Hwang, *Introduction to Random Signals and Applied Kalman Filtering*. Wiley, 2012.

Network-Aided Pseudorange-Based LEO PNT from OneWeb

Zacharias M. Komodromos*, Samuel C. Morgan[†], Zachary L. Clements[†], Wenkai Qin[†], W. Jeremy Morrison[†],
Todd E. Humphreys[†]

**Department of Electrical and Computer Engineering, The University of Texas at Austin*

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

Abstract—This paper presents a framework and first experimental results for pseudorange-based positioning, navigation, and timing (PNT) exploiting OneWeb’s Ku-band downlink signal and a reference network. As the demand for accurate and resilient positioning and timing solutions grows, the proliferation of low Earth orbit (LEO) satellites, including approximately 650 in OneWeb’s constellation, offers promising opportunities for powerful new PNT services. Newly discovered synchronization sequences embedded in the OneWeb signal open the way for third-party provision of highly accurate clock and orbital models for each OneWeb satellite. Model parameters would be continually estimated by a sparse network of reference stations at known locations having access to an accurate universal time scale such as UTC. Subscribers to the third party’s data feed could treat OneWeb much like a traditional GNSS, extracting pseudorange, carrier phase, and Doppler observables directly from the OneWeb signals, either by exploiting only the newly discovered synchronization sequences or augmenting these with additional sequences decoded continuously by the reference network. The observables would then be processed together with the third-party-provided models to produce highly accurate position, velocity, and timing solutions. Because each OneWeb satellite illuminates a wide geographic area, this technique could be implemented with an attractively low density of reference stations. In this paper we develop the OneWeb satellite clock model, establish its applicability across multiple downlink beams, analyze the effect of ephemeris errors on the resulting solution, and prove the technique in a field experiment with live signals, achieving sub-meter-level positioning and timing.

Index Terms—OneWeb, signal processing, positioning, low Earth orbit, CDGNSS.

I. INTRODUCTION

Dual purposing broadband communication signals from low Earth orbit (LEO) mega-constellations for positioning, navigation, and timing (PNT) is the subject of a rapidly growing body of research [1]–[6]. LEO signals potentially offer more robust, secure, and accurate PNT compared to traditional Global Navigation Satellite System (GNSS) signals due to their significantly larger bandwidth, higher received power, superior geometric diversity, and larger constellation size [1]. Thus, they could be the basis of a PNT service that acts as a complement to or a backup for traditional GNSS amid rising threats of GNSS jamming and spoofing [7]–[11]. OneWeb’s Ku-band downlink signal exemplifies the capabilities of LEO signals used for broadband communication, providing 230 MHz-wide channels from around 650 satellites that are 15 times closer to Earth than traditional GNSS satellites [12].

OneWeb’s polar orbits offer global coverage, and their space-to-Earth link operates in the reasonably accessible 10.7–12.7 GHz band.

The great majority of opportunistic LEO PNT methods presented in the literature are Doppler-based [13]. But compared to pseudorange-based PNT techniques, Doppler-based techniques have worse timing accuracy by many orders of magnitude (milliseconds vs. nanoseconds), even under optimistic measurement noise and satellite clock offset rate assumptions [4], [14], [15]. Recognizing that many PNT applications of practical interest require accurate timing, we focus on pseudorange-based PNT, or PNT based on time-of-arrival (TOA) measurements.

The extraordinary bandwidth of the OneWeb signal is a boon for TOA measurements, providing far more resistance to multipath than the widest traditional GNSS signal, the 51.15-MHz Galileo E5 AltBOC signal. Moreover, the OneWeb signal does not suffer from the side-peak ambiguity problem inherent in Galileo AltBOC processing [16].

The approach proposed in this paper leverages pseudorange measurements for a navigation solution and uses signals from OneWeb’s constellation. Any method relying on pseudorange measurements is limited by the clocks at the receiver and the satellite. As recent work indicates, clocks onboard LEO satellites meant for broadband communication may not meet the stringent requirements of PNT applications [17], [18]. To counteract any timing errors, two receivers can be used to cooperatively form the PNT solution. In this framework, a reference station (RS) can act as a reference receiver with a known position and reliable timing. While OneWeb’s satellites move overhead, the RS can estimate the transmission time of a frame, the channel it was transmitted on, and estimate clock errors. The second receiver, the user equipment (UE), can use this information to form TOA measurements with reduced effects from the satellite’s clock errors. In low signal-to-noise (SNR) ratio regimes, the RS could provide more payload symbols to extend the integration period used to collect TOA measurements. The UE could be consumer-grade hardware with an omnidirectional antenna and a lower-quality clock. Especially with the plethora of reasonably priced feedhorn and low noise blocks (LNBs) for Ku-band due to their mass production, the UE need not be expensive.

This TOA-based approach would require the signal structure to be known, which is a challenge for LEO signals like that of OneWeb since their signals are not standardized, nor designed

for navigation [19]. Progress has been made with Starlink’s signal structure with the discovery of synchronization sequences that unlock the potential for pseudorange-based PNT [5]. The soon-to-be published work on which this paper relies uncovers details about OneWeb’s signal [20]. Perhaps the most significant discovery is a short synchronization sequence that repeats every 1 ms, which unlocks pseudorange-based PNT with OneWeb.

Our proposed framework assumes the RS can communicate timing corrections and payload information to the UE within a reasonable time frame. The information an RS communicates becomes stale as different satellites pass overhead, or as a single satellite’s clock drifts. The minimum refresh interval for the clock model is the duration a single beam illuminates the same area, but could possibly be longer if the beams depend on the same clock. Since OneWeb’s satellites have a fixed beam pattern consisting of 16 beams, a study must be conducted to determine if clock corrections for one beam can be used for the entire satellite footprint. A contribution of this paper is evidence that beam-to-beam clock consistency of a single satellite is consistent enough that the clock corrections from the RS could be used for PNT within at least one beam footprint, which is conservatively around $67 \times 1080 \text{ km}^2$.

A. Related Work

Four strategies for realizing LEO PNT are summarized in Table I. The first is a dedicated approach, where satellites from companies like Xona Systems and TrustPoint have a dedicated PNT mission with high deployment cost. The second is a fused approach, in which broadband LEO communications SVs also offer a communications service, which comes at the cost of some lost communication capacity [1], [6], [23], [24], [29]. Both of these approaches offer high accuracy, but at a cost assumed by the constellation owner.

This paper explores the third strategy, in which a network of reference stations provides corrections to pseudorange, Doppler, or carrier-phase measurements of existing LEO satellite signals. In the simplest embodiment, the network is a single reference station. Previous studies rely on Doppler measurements [26], or have only considered pseudorange measurements from OneWeb in simulation [27]. Authors in [28] conducted a study similar to this paper’s using Iridium’s signal and achieved impressive results, considering the comparatively smaller bandwidth.

Note that Table I also classifies the opportunistic approach presented in [25] as network-aided because the receiver being tracked is aided by GNSS signals during an initialization phase, which allows estimation of LEO satellite ephemeris, clock, and clock rate errors just as occurs with network aiding. In effect, the receiver acts as its own reference station during the initialization phase, holding over key estimated values during the GNSS-denied phase.

The fourth strategy in Table I is the stand-alone opportunistic approach, where unmodified LEO constellations are exploited for PNT using only publicly available data and no network aiding.

The first two strategies rely on cooperation with the constellation owner and are not the focus of this paper. The

stand-alone opportunistic approach has not been demonstrated to offer accuracy competitive with traditional GNSS. Recent results in [18] indicate that stand-alone opportunistic use of Starlink for pseudorange-based PNT would fail to offer solutions remotely competitive with traditional GNSS due to satellite timing peculiarities that appear impossible to model. This is also likely true for OneWeb, though a study like [18] has yet to be done for OneWeb. Thus, a network-aided opportunistic approach remains a promising strategy for realizing LEO PNT, since a third-party’s network would provide measurement corrections, obviating the need for a model of mercurial satellite clock variations. As more becomes known about OneWeb’s signal, and additional synchronization sequences are discovered, the network-aided opportunistic approach could become even more attractive.

B. Contributions

This paper outlines the framework in which an RS could provide information to UE located within the same satellite beam footprint, enabling PNT. We offer 3 contributions. First, we develop the necessary clock model that an RS would need to provide to a UE. Second, we study the effects of ephemeris errors on the final positioning solution, and study the effect on the TOA RMSE of sharing more payload symbols with the UE. Third, we conduct a field experiment with live signals to demonstrate the feasibility of the framework, achieving a sub-meter final solution error. The success of this framework could motivate a third party to provide PNT services using OneWeb’s Ku-band downlink signal, or could provide an example for fused GNSS for the constellation’s owners.

II. CONCEPT OF OPERATIONS

Suppose a third party wants to provide PNT services using OneWeb’s Ku-band downlink signal. They can commission a stationary RS at a known location, with a good quality clock, prior knowledge of the satellite ephemerides, and a communication link for any UE subscribed to their service, as depicted in Fig. 1. The RS clock could be disciplined to GNSS or alternatively, for a GNSS-independent approach, the third party might utilize atomic clocks or time signals from other terrestrial broadcast systems for RS clock disciplining. Additionally, the RS boasts a high gain Ku-band antenna to capture the downlink signal and extract the payload.

OneWeb’s roughly 650 satellites follow polar orbits at an altitude of 1200 km. Their downlink signal has approval for the 10.7–12.7 GHz space-to-Earth frequency band. Each channel supports a $F_{\text{sym}} = 230.4\text{-MHz}$ wide single-carrier Time Division Multiple Access (SC-TDMA) signal, with frames being transmitted every $T_f = 10 \text{ ms}$ [20]. Assuming OneWeb is transmitting a QPSK signal, the RS could easily demodulate frames and re-broadcast the data as an uncompressed constant stream of 57.6 MB/s if every symbol of every frame is required - a rate which would be cost ineffective. Furthermore, using the known position of the RS and the satellite, the RS could provide the time of transmission (TOT) for each frame, with some correction for the satellite’s clock error.

TABLE I: LEO PNT Strategies

	Dedicated [19], [21], [22]	Fused [1], [6], [23], [24]	Network-Aided Opportunistic [25]–[29]	Stand-Alone Opportunistic [4], [14], [30]–[32]
Description	LEO constellation or hosted payloads dedicated solely to PNT.	Fuses a secondary PNT service with a primary communications service.	Like stand-alone opportunistic, but a network of reference stations provides corrections.	Exploits unmodified signals from communications SVs. Public ephemerides. No network aiding.
Marginal deployment cost	high: constellation of SVs	low: uses communications hardware and signals	high: network of reference stations	very low
Potential availability	mid-term	near-term	near-term for local coverage	immediate
Potential accuracy [m]	< 1	< 1	< 1	< 100
Time to fix [s]	< 10	< 10	< 10	< 1000
Dependency on traditional GNSS	@SVs in near-term	@SVs in near-term	@reference stations in near-term	@SVs in some cases

The RS would serve consumer-grade UE with a lower quality clock and an omnidirectional Ku-band antenna. Using the RS's broadcasted data, the UE could form TOA measurements using local replicas from the RS decoded payload. These TOAs with the clock corrected TOTs form pseudorange measurements that the UE can use to estimate its position and clock error.

Before such a framework is implemented, certain challenges must be addressed. The first challenge is that of the SV clock error and stability, where these effects could act much like the clock dithering implemented to intentionally degrade GPS accuracy under the Selective Availability program discontinued in May 2000 [33]. The second challenge is determining the refresh rate of the clock models, and the payload information, since the maximum 57.6 MB/s is a costly burden to meet in real-time. The third challenge is the effect of ephemeris errors on the final positioning solution, since a third-party may not have access to the more precise satellite ephemerides required for collision avoidance.

All these questions are intrinsically linked to the signal structure of OneWeb's signal, their fixed beam pattern, and polar orbits. Given their constellation, only a beam from a single satellite illuminates a receiver at a given time. From the satellite's 16 beams, only 8 are active per channel with roughly 20 seconds between the apex in power of each beam [20].

In this paper we will address the above challenges and provide an experimental demonstration of the framework using at most 4 of the 8 active beams per satellite. Due to our experimental setup, the TOA measurements could not be obtained in quick succession, but rather in bursts with small gaps between the beams, and larger gaps when switching between satellites. This increases the burden on our UE to estimate its clock bias and bias rate during the gaps, as its own clock drifts.

III. SIGNAL MODEL

OneWeb employs a Multi-Frequency Time Division Multiple Access (MF-TDMA) Single Carrier (SC) scheme, on each




Fig. 1: Diagram of the proposed framework. A network of reference stations (RSs) provides timing and ephemeris corrections along with data payload symbols to user equipment (UE) through a side channel. The UE then gathers TOA measurements using the RS-broadcasted payload data and forms pseudorange measurements with the provided TOTs.

of 8 separate frequency channels [20]. Only three channels, centered at 11.075, 11.325, and 11.575 GHz are presently active over our location in Austin, Texas. Each channel is modulated with a QPSK signal, although documentation for their user terminal (UT) suggests support for QPSK, 8PSK, and 16APSK [34].

A. QPSK Signal Model

Fundamentally, a SC signal is a train of pulses, each modulated by a phase and amplitude shift corresponding to the data symbol being transmitted on that pulse. The resulting continuous stream of symbols can be characterized by a simple signal model. The baseband signal model for a SC QPSK signal is given by

$$s(t) = \sum_m \exp(j\pi a_m/2) p\left(\frac{t - mT_{\text{sym}}}{T_{\text{sym}}}\right) \quad (1)$$




Fig. 2: Theorized frame structure of OneWeb's downlink signal.

where statistically independent symbol phases $a_m \in \{0, 1, 2, 3\}$, $m \in \mathbb{Z}$ can be encoded to represent two bits per symbol, and T_{sym} is the symbol period. Following the analysis from [20], $T_{\text{sym}} = 1/230.4\text{MHz}$, and $p(t)$ is the square root raised cosine (SRRC) pulse with a rolloff factor $\beta_r = 0.1$.

B. OneWeb Signal Structure

OneWeb's signal structure remains largely unknown. Over our location in Austin Texas, we have observed that satellite beam transmits a frame every $T_f = 10$ ms. Our current hypothesis is that a frame consists of 10 slots, each with a preamble. Within the preamble exists a synchronization sequence which is common across all preambles. The payload data for each slot should contain unique data for each frame, but remains the same in the case the slot has not been assigned to a user. Fig. 2 is not to scale, as the preamble only lasts about $25 \mu\text{s}$. The synchronization sequence, which lasts only 400 symbols [20], is a poor choice for a ranging sequence even though it is identical across all beams and satellites.

Our current observation is that the payload data used when the slot is not assigned to a user remains the same across frames from the same beam. Thus, as things stand, we need only decode a single 10-ms frame from each beam to accommodate TOA measurements for the entire duration the beam illuminates a receiver.

A stationary receiver can easily distinguish specific intervals corresponding to captures from different beams. Let such an interval be called a beam interval (BI). Given OneWeb only uses half of the 16 beams at a time, the BI can be as long as 20 seconds. A continuous capture of seven BIs from OneWeb-0114 on March 2024 is shown in Fig. 3, where a receiver can easily distinguish the different beams simply by using the received power. For the purposes of this paper, we chose to focus on 12-second BIs centered at the apex of transmitted power to ensure no interference in the data collection and maximize the number of symbols we can reliably decode. These manifest as the colored portions in Fig. 3. The use of 12-second BIs introduces a small 8-second gap between the end of one BI and the beginning of the next.

IV. TIMING

Clocks onboard LEO satellites are of unknown quality, and may exhibit peculiarities as was recently revealed for Starlink's satellites [17], [18]. The clock model we develop in this

section is necessary for the RS to provide timing corrections to the UE.

A. Transmitter Clock Model

We speculate that each satellite must have a base oscillator, probably disciplined to GNSS, that feeds a frame clock and a carrier clock for each beam. It may be that the beams use the same frame and carrier clocks, but this is still under investigation. Let t_d be the GNSS-disciplined timing signal that drives the frame clock timing signal t_f and the carrier clock timing signal t_c . The carrier clock is likely transparent, with $t_c = t_d$, while the frame clock depends on how baseband frames are loaded into buffers for mixing to the carrier frequency before transmission. Each of the two clocks may have offsets from t_d denoted as Δt_f and Δt_c respectively such that

$$t_d(t) = t_f(t) - \Delta t_f(t) \quad (2)$$

$$t_d(t) = t_c(t) - \Delta t_c(t) \quad (3)$$

where t represents true time. Each of these is thus related to true time by

$$t = t_f(t) - \delta t_f(t) \quad (4)$$

$$t = t_c(t) - \delta t_c(t) \quad (5)$$

where $\delta t_f(t)$ and $\delta t_c(t)$ are the frame and carrier clock offsets. This distinction is necessary to inform our estimation process later on, since pseudorange measurements would be affected by the frame clock and Doppler measurements by the carrier clock.

B. Discrete-Time Frame clock

Within a BI, frames arrive at a cadence of 10 ms. The first frame of a BI is somewhat arbitrary, especially in the case of a 12-second BI, since there is no clear feature indicating a first unique frame in a BI. Starting from the first observed frame in a BI, we can index the frames from 0 to $N_a - 1$ where $N_a = 1200$ for a 12-second BI.

Let $t_f(l, m)$ be the frame clock time at the time of transmission of the m th frame in the l th BI, where m and l are zero-indexed. This quantity is defined as

$$t_f(l, m) \triangleq t(l) + mT_f \quad (6)$$

where $t(l)$ is the time of the first frame in the l th BI. Let $t^*(l, m)$ and $\delta t_f(l, m)$ be the corresponding GPS time (GPST) and frame clock offset. Then

$$t^*(l, m) = t_f(l, m) - \delta t_f(l, m) \quad (7)$$

The receiver clock offset $\delta t_r(t_r)$ is represented as a function of t_r because it is natively ordered in receiver time in the course of solving for a position and time solution. The time derivative of δt_r with respect to t , denoted $\delta \dot{t}_r(t_r(t))$, is called the receiver clock drift.




Fig. 3: Normalized power of the signal (left), with the associated ground track (right) of each BI and the receiver (marked by the triangle). Data are for ONEWEB-0114 from a capture centered at F_{c2} taken in March 2024.

Let $t_r(l, m)$ be the time of reception, according to the receiver clock, of the frame transmitted at true time $t^*(l, m)$. Let $\delta t_r(l, m)$ be the corresponding receiver clock offset and $t_*(l, m)$ be the corresponding true time of reception. More precisely, $t_r(l, m)$ is the receiver clock time at which the frame transmitted at true time $t^*(l, m)$ from the satellite's downlink antenna's phase center first reached the receiver antenna's phase center. The receipt time $t_r(l, m)$ can be related to $t_*(l, m)$, $t^*(l, m)$, and $t_f(l, m)$ by

$$t_*(l, m) = t_r(l, m) - \delta t_r(l, m) \quad (8)$$

$$t^*(l, m) = t_r(l, m) - \delta t_r(l, m) - \delta t_{\text{tof}}(l, m) \quad (9)$$

$$t_f(l, m) = t_r(l, m) - \delta t_r(l, m) - \delta t_{\text{tof}}(l, m) + \delta t_f(l, m) \quad (10)$$

where $\delta t_{\text{tof}}(l, m)$ is the frame's time of flight from transmission to reception, expressed as an interval in true time.

C. Timing corrections

TOT calculation on the RS begins with the frame TOA measurement in receiver time. We subtract from this an estimate of the offset $\delta t_r(l, m)$ obtained from the simultaneously captured GNSS signals, accounting for the difference in length of the cables from our capture equipment to the GNSS and Ku-Band antennas. This process allows us to determine $t_*(l, m)$ to within a few ns, which, in turn, is related to $t^*(l, m)$ by

$$t^*(l, m) = t_*(l, m) - \delta t_{\text{tof}}(l, m) \quad (11)$$

What remains is to calculate the frame's time of flight, $\delta t_{\text{tof}}(l, m)$, modeled as

$$\delta t_{\text{tof}}(l, m) = \frac{1}{c} \cdot \|\mathbf{r}_r - \mathbf{r}_t(t^*(l, m))\| + \delta t_{\text{atm}} \quad (12)$$

where $\mathbf{r}_r$ is the receiver's location, $\mathbf{r}_t(t^*(l, m))$ is the transmitter's location at the TOT, and δt_{atm} is the atmospheric (neutral and ionospheric) delay experienced by the signal over its path from transmitter to receiver. Substituting (11) into (12) yields the implicit relationship

$$\delta t_{\text{tof}}(l, m) = \frac{1}{c} \cdot \|\mathbf{r}_r - \mathbf{r}_t(t_*(l, m) - \delta t_{\text{tof}}(l, m))\| + \delta t_{\text{atm}} \quad (13)$$

from which $\delta t_{\text{tof}}(l, m)$ can readily be calculated numerically, provided δt_{atm} and a smooth transmitter location function $\mathbf{r}_t(t)$.

To support determination of $t^*(l, m)$ to the best of our ability, errors in $\mathbf{r}_t(t)$, $\mathbf{r}_r$, and δt_{atm} must be small relative to this amount. For δt_{atm} , a Saastamoinen [35] neutral atmospheric model with average surface parameters was applied with Niell wet and dry mapping functions [36]; ionospheric delays, which are minimal at Ku-band, were ignored. The process is identical to that outlined in [18, Sec. VII-A].

Unlike the work in [18], we did not have access to meter-level accurate ephemerides for the satellites. Instead, we used the Two-Line Element (TLE) data provided by the North American Aerospace Defense Command (NORAD) to estimate the satellite's position at the time of transmission. This introduces an error in the $\delta t_{\text{tof}}(l, m)$ calculation, which we will address in Sec. VI.

V. DATA AIDING

Given the SS for OneWeb is short, a high-gain antenna is required to extract TOA measurements accurately. Since the UE is assumed to be consumer-grade hardware, we must rely on the RS to provide the UE with additional symbols to extend the coherent integration period. This raises the questions of how much data the RS should provide, and how often.

A. Data burst duration

The duration of the data bursts directly impacts the quality of the TOA measurements the UE can obtain. This has been studied extensively in the literature [37], [38], and has been applied to Starlink's signal to provide RMSE bounds on the TOA measurements obtained from their signal [18]. The analysis for OneWeb is no different.

In short, the TOA error variance for a signal with Fourier transform $S(f)$ and post integration SNR E/N_0 can be expressed in the form [39, Eq. 11.15]

$$\sigma_\tau^2 \geq \frac{1}{2\gamma^2 \frac{E}{N_0}} \quad (14)$$

where γ^2 is the effective squared bandwidth defined as

$$\gamma^2 \triangleq \frac{\int_{-\infty}^{\infty} (2\pi f)^2 |S(f)|^2 df}{\int_{-\infty}^{\infty} |S(f)|^2 df}. \quad (15)$$

In Sec. VIII-B, we provide information on the UE setup which is limited to only 25 MHz of the 230.4 MHz wide




Fig. 4: TOA RMSE bounds for SNRs of -5 , 0 , and 5 dB over the coherent processing interval.

signal. The captured bandwidth of the signal on the UE will be approximately flat. Flat spectra have an effective squared bandwidth of $\gamma^2 = W^2/12$, where in the case of the UE the bandwidth is $W = 2\pi \cdot 25 \cdot 10^6 \approx 157.07 \cdot 10^6$ rad/sec. Varying the coherent integration period, and thus the post integration SNR, will lead to different bounds on the TOA error. This is depicted in Fig. 4 for three pre-correlation SNRs of the signal. Larger data bursts improve accuracy but yield diminishing returns.

B. Data cadence

The cadence at which the RS provides the UE with additional symbols determines the number of measurements the UE has to estimate its position and clock bias. Since the system is dynamic, with the satellites moving, more measurements may smooth out short-term variations in errors, and could help improve the Geometric Dilution of Precision (GDOP) of the solution. The closer measurements are taken, the more correlated they are, and the less unique information they provide. A truly informed decision on the cadence of the data bursts would require a realization of satellite tracks.

The error bounds on the final solution will depend on $P = (H^T R^{-1} H)^{-1}$, where R is the measurement covariance matrix and H is the measurement matrix. The measurement matrix H can only be formed with some realization of the satellite tracks. Once an error covariance is computed, we can transform the ECEF errors to ENU, and calculate the semi-major axis of the error ellipse. This is depicted in Fig. 5 for a single satellite moving directly overhead, and a satellite reaching at most 70 degrees elevation. The bounds on the error will decrease as more satellites are used in the solution, but the relationship between the data burst cadence and error remains the same.

VI. EPHEMERIS ERROR MODELING

At any given epoch, the position of a OneWeb satellite can be retrieved from publicly available TLEs. The accuracy of the satellite position when retrieved from a TLE is typically km-level accurate, with spatial errors reaching 1 to 10 km 24 hours after a TLE file is updated [40]. Spatial ephemeris errors become a source of error in pseudorange measurements.




Fig. 5: Semi-major of the error ellipse of solution over the cadence of data bursts for a satellite moving directly overhead (blue), and a satellite reaching at most 70 degrees elevation (orange).

This subsection aims to characterize how the spatial ephemeris error affects the differential pseudorange error when measured at the RS and UE.

Let $\mathbf{r}_s$ denote the true position of the satellite, $\mathbf{r}_r$ denote the known position of the RS, and $\mathbf{r}_u$ be the true position of the UE. The true pseudoranges between $\mathbf{r}_r$ and $\mathbf{r}_s$, and $\mathbf{r}_u$ and $\mathbf{r}_s$, are denoted ρ_r and ρ_u , respectively, are expressed as

$$\rho_r = \sqrt{(\mathbf{r}_r - \mathbf{r}_s)^T (\mathbf{r}_r - \mathbf{r}_s)} \quad (16)$$

$$\rho_u = \sqrt{(\mathbf{r}_u - \mathbf{r}_s)^T (\mathbf{r}_u - \mathbf{r}_s)} \quad (17)$$

Now, suppose that the assumed satellite position contained a 3D spatial error ϵ . The resulting pseudorange when assuming the erroneous satellite position at the RS, $\tilde{\rho}_r$, and UE, $\tilde{\rho}_u$, are expressed as

$$\tilde{\rho}_r = H_r \epsilon + \rho_r \quad (18)$$

$$\tilde{\rho}_u = H_u \epsilon + \rho_u \quad (19)$$

where H_r and H_u are the 1×3 Jacobians with respect to the satellites position, expressed as

$$H_r = \frac{\partial \rho_r}{\partial \mathbf{r}_s} = -\hat{\mathbf{r}}_r = -\frac{\mathbf{r}_r - \mathbf{r}_s}{\rho_r} \quad (20)$$

$$H_u = \frac{\partial \rho_u}{\partial \mathbf{r}_s} = -\hat{\mathbf{r}}_u = -\frac{\mathbf{r}_u - \mathbf{r}_s}{\rho_u} \quad (21)$$

The pseudorange measurement error at the RS and the UE due to the spatial error of the satellite will manifest differently due to the different geometry. However, if the positions of the RS and UE are close to each other, the errors will be small. The relative differential pseudorange error between the RS and the UE due to spatial ephemeris error is denoted as η , and is expressed as

$$\begin{aligned} \eta &= (\tilde{\rho}_r - \rho_r) - (\tilde{\rho}_u - \rho_u) \\ &= (H_r - H_u) \epsilon \end{aligned} \quad (22)$$

For example, consider a specific $\mathbf{r}_r$, $\mathbf{r}_u$, $\mathbf{r}_s$, and 3D satellite spatial error $\epsilon \sim \mathcal{N}(\mathbf{0}, P_\epsilon)$, then

$$\eta \sim \mathcal{N}(0, \sigma_\eta^2) \quad (23)$$

with

$$\sigma_\eta^2 = (H_r - H_u) P_\epsilon (H_r - H_u)^T \quad (24)$$

The distribution of η in (23) requires the knowledge of $\mathbf{r}_u$ to compute the Jacobian which is required for σ_η^2 . In practice, $\mathbf{r}_u$ is unknown, but the CT will have *a priori* information about $\mathbf{r}_u$ relative to $\mathbf{r}_r$. For instance, the CT might know that it is located within the same beam footprint as the CS. In this case, $\mathbf{r}_u$ is modeled as

$$\mathbf{r}_u = \mathbf{r}_r + \boldsymbol{\xi} \quad (25)$$

where $\boldsymbol{\xi}$ is a 3D random offset with respect to $\mathbf{r}_r$, following some distribution with covariance P_ξ . Since $\mathbf{r}_u$ is now treated as a random variable, $H_u(\boldsymbol{\xi})$ and $\sigma_\eta^2(\boldsymbol{\xi})$ also become random variables. $H_u(\boldsymbol{\xi})$ is expressed as

$$H_u(\boldsymbol{\xi}) = H_r + \boldsymbol{\xi}^\top A \quad (26)$$

with

$$A^\top = \frac{\partial H_r^\top}{\partial \mathbf{r}_r} = \mathbf{r}_r^{-1} (\hat{\mathbf{r}}_r \hat{\mathbf{r}}_r^\top - \mathbb{I}) \quad (27)$$

where A is the Jacobian of H_r with respect to $\mathbf{r}_r$. The derivation of this Jacobian can be found in [41]. Consequently, (24) is reduced to

$$\sigma_\eta^2(\boldsymbol{\xi}) = \boldsymbol{\xi}^\top A P_\epsilon A^\top \boldsymbol{\xi} \quad (28)$$

Now, $\sigma_\eta^2(\boldsymbol{\xi})$ is a function of the unknown random vector $\boldsymbol{\xi}$, so solving for $\bar{\sigma}_\eta^2$ is required for modeling purposes. Note, (28) is a quadratic form, thus

$$\bar{\sigma}_\eta^2 = \mathbb{E}[\sigma_\eta^2(\boldsymbol{\xi})] = \text{trace}(A P_\epsilon A^\top P_\xi) \quad (29)$$

Finally, the distribution of η when $\mathbf{r}_u$ is unknown can be modeled as

$$\eta \sim \mathcal{N}(0, \bar{\sigma}_\eta^2). \quad (30)$$

Consider the scenario where the ephemeris error is zero-mean Gaussian in the ENU frame, with covariance matrix $P_\epsilon = \sigma_{\text{eph}}^2 \mathbb{I}$. Also, assume that $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, P_\xi)$, where $P_\xi = \text{diag}(\sigma_e^2, \sigma_n^2, \sigma_u^2)$, centered at the RS. The choices of σ_{eph}^2 , σ_e^2 , σ_n^2 , and σ_u^2 are treated as tuning parameters. The values of σ_e^2 , σ_n^2 , and σ_u^2 depend on the assumed distance between the RS and UE.

In the case of Oneweb, it is required that the RS and UE occupy the same beam footprint. Consider two scenarios: (1) a dense RS network where the UE's position is Gaussian with $P_\xi = \text{diag}(3300^2, 3300^2, 50^2)$ and (2) a sparse RS network where UE's position is Gaussian with $P_\xi = \text{diag}(25000^2, 25000^2, 1000^2)$. Additionally, let $\sigma_{\text{eph}} = 2.5$ km and assume a typical Oneweb to beam footprint geometry. If a dense RS network is available, the spatial relative pseudorange error between RS and UE can be expected to be small, as $\bar{\sigma}_\eta = 9.0$ m. On the other hand, if a sparse RS network is available, the spatial relative pseudorange error becomes much larger, as $\bar{\sigma}_\eta = 67.8$ m. The distribution of η will be exploited in Section VII.

VII. ESTIMATION

The UE has access to its own TOAs $t_r(l, m)$, the TOTs $t^*(l, m)$ that have already been compensated from the RS, its own Doppler measurements $f_D(l, m)$, and the position $\mathbf{r}_s(l, m)$ and velocity $\mathbf{v}_s(l, m)$ of the satellites at the TOT. Given the results in Sec. VI, the UE will not be able to estimate its clock bias in a single variable because each satellite will have a different spatial ephemeris error which will manifest as a different clock bias over a short time duration. As such, for a stationary UE, we must estimate its position $\mathbf{r}_u$, a frame clock offset consisting of the receiver bias and the ephemeris-error pseudorange bias for each satellite $b(s_{l,m}) = \delta t_r + \eta(s_{l,m})$, a frame clock drift $\delta \dot{t}_r$, and a carrier clock drift for each satellite $\delta \dot{t}_c(s_{l,m})$ which is the relative carrier clock drift between the UE and the satellite.

The value $s_{l,m} \in \{0, 1, \dots, N_s - 1\}$ is a mapping from beam and frame index to a satellite index for N_s satellites. For example, if two measurements are from the same satellite and the first is from beam i with frame index k , while the second is from beam j with frame index p , they map to the same satellite $s_{i,k} = s_{j,p}$ and have a common clock biases $b(s_{i,k}) = b(s_{j,p})$. For a measurement identified by (l, m) , our observables are the pseudorange and the range-rate:

$$\begin{bmatrix} \rho(l, m) \\ \dot{\rho}(l, m) \end{bmatrix} = \begin{bmatrix} c[t_r(l, m) - t^*(l, m)] \\ -\lambda f_D(l, m) \end{bmatrix} + \begin{bmatrix} w_\rho(l, m) \\ w_{\dot{\rho}}(l, m) \end{bmatrix} \quad (31)$$

where w_ρ and $w_{\dot{\rho}}$ are both zero-mean Gaussian noise with variances σ_ρ^2 and $\sigma_{\dot{\rho}}^2$.

It is important to note here that $\dot{\rho}(m, l)$ is the carrier clock drift, and not the rate of the pseudoranges obtained from the frame clock $\rho(m, l)$. A measurement vector $\mathbf{z}$ is constructed of all the pseudorange measurements followed by all the carrier range-rate measurements.

A. Least squares estimation

Suppose the UE has N_f measurements per beam, from N_b beams corresponding to N_s satellites. There are $N_z = N_f N_b$ total pairs of measurements. The state the UE must estimate is

$$\mathbf{x} = \begin{bmatrix} \mathbf{r}_u^\top & \mathbf{b}^\top & \delta \dot{t}_r & \delta \dot{t}_c^\top \end{bmatrix}^\top \quad (32)$$

where $\mathbf{r}_u \in \mathbb{R}^3$, $\mathbf{b}, \delta \dot{t}_c \in \mathbb{R}^{N_s}$, and $\delta \dot{t}_r \in \mathbb{R}$.

The nonlinear measurement model is

$$\begin{bmatrix} h_\rho(\mathbf{x}; l, m) \\ h_{\dot{\rho}}(\mathbf{x}; l, m) \end{bmatrix} = \begin{bmatrix} \|\mathbf{r}(l, m)\| + b(s_{l,m}) + t_r(l, m) \delta \dot{t}_r \\ \hat{\mathbf{r}}^\top(l, m) \mathbf{v}_s(l, m) + \delta \dot{t}_c(s_{l,m}) \end{bmatrix} \quad (33)$$

where $\mathbf{r}(l, m)$ is the vector pointing from satellite $s_{l,m}$ to the rover, and $\hat{\mathbf{r}}(l, m)$ is its unit vector. All measurements are in meters and meters per second.

Let H be the linearized measurement model

$$\begin{aligned} H_\rho &= \begin{bmatrix} \frac{\partial h_\rho(\mathbf{x}; 0, 0)}{\partial \mathbf{x}} & \dots & \frac{\partial h_\rho(\mathbf{x}; N_b - 1, N_f - 1)}{\partial \mathbf{x}} \end{bmatrix}^\top \\ H_{\dot{\rho}} &= \begin{bmatrix} \frac{\partial h_{\dot{\rho}}(\mathbf{x}; l, m)}{\partial \mathbf{x}} & \dots & \frac{\partial h_{\dot{\rho}}(\mathbf{x}; N_b - 1, N_f - 1)}{\partial \mathbf{x}} \end{bmatrix}^\top \\ H &= \begin{bmatrix} H_\rho \\ H_{\dot{\rho}} \end{bmatrix} \end{aligned} \quad (34)$$

With the linearized measurement model, an increment from the current estimate $\Delta \mathbf{x}$ can be calculated and used to update the estimate $\hat{\mathbf{x}} + \Delta \mathbf{x}$, repeating for as many iterations as desired.

B. Altitude Constraint

The estimator includes an altitude constraint by employing a Gaussian distributed pseudo-measurement parameterized by a mean μ_{alt} and variance σ_{alt}^2 of the altitude with respect to the WGS-84 ellipsoid. This pseudo-measurement constrains the altitude of the UE and allows us to keep position in the ECEF coordinate system instead of converting at each step to a geodetic coordinate system. To apply the constraint, we can add a term to the measurements and extend the measurement model similar to [11]:

$$z_{2N_z+1} = \mu_{\text{alt}} \quad (35)$$

$$h_{\text{alt}}(\mathbf{x}) = \text{alt}(\mathbf{r}_u)$$

$$H_{2N_z+1} = [\cos(\phi_{\text{lat}})\cos(\lambda_{\text{lon}}), \cos(\phi_{\text{lat}})\sin(\lambda_{\text{lon}}), \sin(\phi_{\text{lat}}), 0] \quad (36)$$

where $\text{alt}(\mathbf{r}_u)$ returns the altitude of $\mathbf{r}_u$ with respect to the WGS-84 ellipsoid, ϕ_{lat} and λ_{lon} are the latitude and longitude of $\mathbf{r}_u$, respectively. The measurement covariance matrix R is also augmented with an additional element on the diagonal, σ_{alt}^2 that is independent from all other measurements.

C. Frame Clock bias constraint

As shown in (33), the pseudorange measurements from satellite $s_{l,m}$ contain a constant measurement bias $b(s_{l,m}) = \delta t_r + \eta(s_{l,m})$, where $\eta(s_{l,m})$ is an additional bias term due to the spatial ephemeris error of satellite $s_{l,m}$. As the UE gets pseudorange measurements from a new satellite $s_{i,k}$, there will be an instantaneous jump in pseudorange bias due to the new $\eta(s_{i,k})$. Rather than estimating each $b(s_{l,m})$ independently, the UE can exploit the distribution of $\eta(s_{l,m})$ to constrain the estimate of $b(s_{l,m})$ for $s_{l,m} = 1, 2, \dots, N_s - 1$ with respect to $b(0)$. This can be accomplished via $N_s - 1$ pseudo-measurements, where each pseudo-measurement is the difference between $b(s_{l,m})$ and $b(0)$ for $s_{l,m} = 1, 2, \dots, N_s - 1$. The Jacobian H_b and its corresponding measurement covariance matrix R_b becomes

$$H_b = [\mathbf{0}_{N_s-1 \times 3}, \mathbf{1}_{N_s-1 \times 1}, -\mathbb{I}_{N_s-1 \times N_s-1}, \mathbf{0}_{N-1 \times N+1}] \quad (37)$$

$$R_b = \text{diag}(\bar{\sigma}_{\eta}^2(1), \bar{\sigma}_{\eta}^2(2), \dots, \bar{\sigma}_{\eta}^2(N_s - 1)) + \bar{\sigma}_{\eta}^2(0) \mathbf{1}_{N_s-1 \times N_s-1} \quad (38)$$

This set of $N_s - 1$ pseudo-measurements are independent of the other measurements, but have a common covariance due to $\eta(0)$. The values of $\bar{\sigma}_{\eta}(s_{i,k})$ are calculated for each satellite, given the satellite's position and the *a priori* information of $\mathbf{r}_u$ with respect to $\mathbf{r}_r$. As the UE gets farther away from the RS, the constraint on the bias becomes weaker. If the UE does not have *a priori* information of $\mathbf{r}_u$ with respect to $\mathbf{r}_r$, this will result in estimating each $b(s_{l,m})$ independently. However, this is an unlikely scenario because the UE will need to know which RS to retrieve timing corrections from. This constraint

on $\mathbf{b}$ significantly improves observability, resulting in a smaller error ellipse.

VIII. RESULTS

Our field experiment consisted of a single RS located on the campus of The University of Texas at Austin, and a single UE located roughly 3 km Northeast of the RS. Both the RS and UE were centered to capture at 11.575 GHz, or channel 4 from [20]. This section presents the results of the experiment given the advocated framework in this paper.

A. RS hardware and capabilities

The RS is equipped with a steerable 90-cm offset parabolic dish with an approximately 3-degree beamwidth. Using publicly available ephemerides from NORAD in the form of TLEs, we can steer the dish to track OneWeb satellites overhead. Our antenna's beamwidth limits captures to a single satellite at a time, which for OneWeb's fixed beam approach is sufficient.

Fig. 6 outlines the hardware used to capture the raw IQ samples. Our parabolic dish is equipped with a feedhorn connected to a low-noise block (LNB) with a conversion gain of 60 dB and a noise figure of 0.8 dB. The LNB downconverts 10.7–11.7 GHz signals to 950–1950 MHz, or 11.7–12.75 GHz to 1100–2150 MHz. The antenna's nominal gain is 40 dBi at 12.5 GHz, with at least 4-5 dB of losses due to lack of circular-to-linear polarizer and feedhorn misalignment. The antenna is located on the campus of the University of Texas at Austin, with a clear view of the sky.

The downstream hardware performs additional downmixing, bandpass filtering, and 16-bit complex sampling. Both the LNB and downstream downconversion and sampling hardware are locked to the same GPS-disciplined oven-controlled crystal oscillator (OCXO).

The usable bandwidth of the capture is 200 MHz. The capture system is capable of capturing on two channels at once, with the limitation that the sampling rate must be identical for the sampling to begin simultaneously. Alongside the OneWeb captures, we simultaneously capture GPS centered at L1 to obtain a true time of arrival for the OneWeb signal at the RS. While our capturing equipment is capable of continuous capture, we opted for 70-second captures due to storage limitations. The 70 seconds are enough for clear captures of three to four beams per satellite.

For OneWeb, the RS must be able to capture the entire signal's bandwidth to decode it, since it is a SC signal. This offers a good point of comparison with an OFDM-based signal like Starlink, since a RS under the same framework for Starlink could make do with some subset of the bandwidth, and in fact wouldn't gain much from capturing at a wider bandwidth than the UE. A high gain antenna on the RS is key, since the higher the received SNR the less bit error rate (BER) the RS will have after decoding, maximizing the correlation on the UE side.

B. UE Setup

Our UE is equipped with a Ku-band feedhorn pointing towards zenith. The feedhorn is connected to an LNB identical




Fig. 6: Block diagram of the OneWeb signal capture process at the RS.

to the one at the RS. The hardware performs downmixing to baseband, bandpass filtering, and 16-bit complex sampling. Both the downmixing in the LNB and in the sampling operations in the downstream hardware are phase-locked to a common GPS-disciplined OCXO to minimize the effects of receiver clock variations. The capture system is also capable of capturing on two channels at once. We take advantage of this dual capture process to simultaneously capture GPS centered at L1 for truth data. Both channels capture at 25 Msps. On the UE, we capture continuously for around 30 minutes to maintain coherence between the captures of the satellite beams.

We had initially planned on estimating only a single frame clock offset for the entire capture. As mentioned in Sec. VII-C, the different errors from each TLE manifested as different frame clock biases per satellite, rendering the problem infeasible to solve without estimating the biases separately. In a more realistic scenario, the RS may have access to ephemerides of better quality than TLEs. Alternatively, if a network of RS are used with visibility of the same satellites, the RS network could estimate its own corrections to the ephemerides.

Unlike the RS, the UE need not have a large gain antenna, or even a large bandwidth since no decoding occurs on the UE. These factors enable a cheaper and more portable UE. Our realization of the UE relies on GPS-disciplining to keep the clock drift low. Since the RS doesn't continuously capture, and the gaps in the captures while repositioning to a new satellite are large, long idle periods would result in a poor-quality clock to drift without measurements to correct it. In practice, where both RS and UE would be continuously capturing, the UE could forgo the GPS-disciplining. Additionally, the RS could capture on multiple channels, allowing the UE to switch between them to minimize the gaps between beams.

C. Experimental Results

As a satellite moves overhead, and its signal is captured by both the RS and UE, the RS begins decoding the frames as they come in. This process is identical to the one outlined in [20]. For each frame, the RS also obtains the TOT as outlined in Sec. IV-C. For our experiment, we opted to provide the

UE with 2 ms of data every 50 ms which, if done in real time would require a link capable of 2.304 MB/s on average. Compared to the data required to be sent over per frame, the data burden for clock and orbital corrections in the form of TOTs and satellite positions and velocities are negligible.

On the other end the UE uses the received partial payload data to generate a local replica with which to correlate the incoming signal against. For our experiment we used the BLS process outlined in Sec. VII in post-processing. In practice this would be replaced by a sequential estimator that would use new TOAs to update the current state estimate as they come in. The TOAs along with the reference-provided TOTs are used to estimate the position and clock error of the UE.

Our data was collected over roughly a 30-minute period in January 2025. The captures consisted of 12 satellites, the ground-tracks of which are shown in Fig. 7. We ran the BLS process once to obtain the standard deviation of the residuals. Using twice the standard deviation of the first run in our noise covariance matrix, we ran the BLS process again to obtain our solution. For each satellite, the standard deviation of the noise provided for the pseudoranges was between 3.3 and 4.1 m, and for the Doppler measurements between 0.2 and 0.9 m/s. The residuals from the second run are shown in Fig. 8, and the final position estimate is shown in Fig. 9.

The bias constraints from Sec. VII-C were calculated to assume that the UE's position was Gaussian-distributed with $\sigma_e = \sigma_n = 3.3$ km and $\sigma_u = 50$ m around the RS. The TLE's were 12 hours old, so $\sigma_{eph} = 2.5$ km was used. The BLS process was robust enough where even initializing the position estimate 250 km away from the true position, the solution converged to the same point. The final position estimate was within 0.35 m of the true position, with the solution contained within a 95% error ellipse with a semi-major of 36.38 m and a semi-minor of 2.40 m. Importantly, timing was resolved to within about 69.49 ns with 95% confidence. When assuming a sparse RS network, the UE's position was modeled as Gaussian distributed with $\sigma_e = \sigma_n = 25$ km and $\sigma_u = 1$ km around the RS, resulting in the semi-minor growing to 15.17 m. The results show that the framework outlined in this paper is feasible and can provide accurate position estimates,




Fig. 7: The track each of the 12 satellites took, projected on the Earth, only while their beams were active.




Fig. 8: Pseudorange (top) and Doppler (bottom) residuals from the BLS process for 6815 measurements of 12 satellites, with 1-4 beams captured per satellite, spanning about 30 minutes.

along with clock bias and bias rate estimates.

Future work will focus extending this work by reducing some sources of error in the process described so far. A sub-sample super-resolution method to obtain TOA and Doppler measurements would significantly reduce the measurement noise. Additionally, while the pseudorange measurements on the UE benefited from RS corrections, the range-rate measurements were not corrected. Jumps in the range-rate measurements in Fig. 8 are due to carrier clock jumps from the satellites and remained unmodeled in the BLS process.

IX. CONCLUSIONS

With this paper we outline a framework for a pseudorange-based PNT solution using OneWeb's downlink signal and the cooperation of a third-party providing timing corrections and data aiding to extend the coherent integration period. Further, we demonstrate through an experiment the feasibility of the proposed framework. Given the additional challenges our experimental setup faced compared to a real-world scenario, we believe the framework and our results show that using broadband communications signals from LEO for PNT is competitive with traditional GNSS. If implemented correctly, it could provide better timing and positioning services than traditional GNSS.



Fig. 9: Final UE position estimate, with the true position indicated with a star. The error of the final estimate is 0.35m, with the solution contained within a 95% error ellipse, which has a semi-major of 36.38m and a semi-minor of 2.40m.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] P. A. Iannucci and T. E. Humphreys, "Fused low-Earth-orbit GNSS," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2022.
- [2] N. Jardak and Q. Jault, "The potential of LEO satellite-based opportunistic navigation for high dynamic applications," *Sensors*, vol. 22, no. 7, p. 2541, 2022.
- [3] Z. Kassas, in *Navigation from low earth orbit — Part 2: Models, implementation, and performance in Position, Navigation, and Timing Technologies in the 21st Century*, 2021, vol. 2.
- [4] M. L. Psiaki, "Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids," *NAVIGATION*, vol. 68, no. 3, pp. 621–641, 2021.
- [5] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–16, 2023.
- [6] 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.
- [7] T. E. Humphreys, "Interference," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 469–503.
- [8] M. J. Murrian, L. Narula, P. A. Iannucci, S. Budzien, B. W. O'Hanlon, M. L. Psiaki, and T. E. Humphreys, "First results from three years of GNSS interference monitoring from low Earth orbit," *NAVIGATION*, vol. 68, no. 4, pp. 673–685, 2021.
- [9] Z. Clements, P. Ellis, and T. E. Humphreys, "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit," in *Proceedings of the IEEE/ION PLANS Meeting*, Monterey, CA, 2023, pp. 458–469.
- [10] G. S. Workgroup, "GPS spoofing: Final report of the GPS spoofing workgroup," OPSGROUP, Tech. Rep., 2024. [Online]. Available: <https://ops.group/blog/gps-spoofing-final-report>
- [11] Z. L. Clements, P. B. Ellis, M. J. Murrian, M. L. Psiaki, and T. E. Humphreys, "Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit," *NAVIGATION*, 2025, submitted for review.
- [12] W. S. Limited, "Amendment to modification application for U.S. Market Access Grant for the OneWeb Ku- and Ka-Band system," https://licensing.fcc.gov/myibfs/download.do?attachment_key=3495551, Jan. 2021, SAT-APL-20210112-00007.

- [13] H. Sallouha, S. Saleh, S. De Bast, Z. Cui, S. Pollin, and H. Wymeersch, "On the ground and in the sky: A tutorial on radio localization in ground-air-space networks," *IEEE Communications Surveys and Tutorials*, vol. 27, no. 1, pp. 218–258, 2025.
- [14] B. McLemore and M. L. Psiaki, "Navigation using Doppler shift from LEO constellations and INS data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4295–4314, 2022.
- [15] A. Baron, P. Gurfil, and H. Rotstein, "Implementation and accuracy of Doppler navigation with LEO satellites," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 2, 2024.
- [16] T. Fan, T. Zhang, H. Zhang, J. Mo, and X. Niu, "A double sideband combined tracking method for Galileo E5 AltBOC signals," *Satellite Navigation*, vol. 4, no. 1, p. 27, 2023.
- [17] 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.
- [18] W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, "Timing properties of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, 2025, submitted for review.
- [19] E. Rubinov, "FrontierSI State of the Market Report LEO PNT 2024 Edition," Jan. 2025, <https://frontiersi.com.au/wp-content/uploads/2025/01/FrontierSI-State-of-Market-Report-LEO-PNT-2024-Edition-v1.1.1.pdf>.
- [20] Z. M. Komodromos, Z. L. Clements, and T. E. Humphreys, "Signal parameter estimation and demodulation of the OneWeb Ku-Band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, 2025, in preparation.
- [21] W. Van Uytsel, T. Janssen, M. Weyn, and R. Berkvens, "A technical overview of current "new space" LEO-PNT initiatives and their application potential," in *2024 IEEE 35th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*. IEEE, 2024, pp. 1–6.
- [22] N. S. Miller, J. T. Koza, S. C. Morgan, S. M. Martin, A. Neish, R. Grayson, and T. Reid, "SNAP: A Xona Space Systems and GPS software-defined receiver," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, 2023, pp. 897–904.
- [23] P. A. Iannucci and T. E. Humphreys, "Economical fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, 2020.
- [24] A. M. Graff and T. E. Humphreys, "Ziv-Zakai-optimal OFDM resource allocation for time-of-arrival estimation," *IEEE Transactions on Wireless Communications*, 2025.
- [25] Z. M. Kassas, N. Khairallah, and S. Kozhaya, "Ad astra: Simultaneous tracking and navigation with megaconstellation LEO satellites," *IEEE Aerospace and Electronic Systems Magazine*, 2024.
- [26] S. Shahcheraghi, J. Saroufim, and Z. M. Kassas, "Acquisition, Doppler tracking, and differential LEO-aided IMU navigation with uncooperative satellites," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.
- [27] J. Saroufim and Z. M. Kassas, "Ephemeris and timing error disambiguation enabling precise LEO PNT," *IEEE Transactions on Aerospace and Electronic Systems*, 2025.
- [28] S. Zhou, R. Yang, Y. Li, X. Zhan, and H. Qin, "Iridium TOA estimation and positioning based on carrier tracking and beam decoding," *IEEE Transactions on Instrumentation and Measurement*, vol. 74, pp. 1–23, 2025.
- [29] S. C. Morgan, Z. M. Komodromos, W. Qin, Z. L. Clements, A. M. Graff, W. J. Morrison, and T. E. Humphreys, "A mock implementation of fused LEO GNSS," in *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.
- [30] Z. Tan, H. Qin, L. Cong, and C. Zhao, "Positioning using IRIDIUM satellite signals of opportunity in weak signal environment," *Electronics*, vol. 9, no. 1, p. 37, 2019.
- [31] M. Neinavaie, J. Khalife, and Z. M. Kassas, "Acquisition, Doppler tracking, and positioning with Starlink LEO satellites: First results," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 3, pp. 2606–2610, 2022.
- [32] N. Jardak, R. Adam, and Q. Jault, "Leveraging multi-LEO satellite signals for opportunistic positioning," *IEEE Access*, 2024.
- [33] D. Odijk and L. Wanninger, "Differential positioning," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 753–780.
- [34] Radio Communications & EMC, "OneWeb ow70I UT test report FCC id xxz-intow70ldac," <https://fcc.report/FCC-ID/XXZ-INTOW70LDAC/5364479.pdf>, 2021, XXZ-INTOW70LDAC.
- [35] J. Saastamoinen, "Atmospheric correction for the troposphere and stratosphere in radio ranging of satellites," in *Geophysical Monograph 15*, S. W. Henriksen, Ed. Washington, D.C.: American Geophysical Union, 1972, pp. 247–251.
- [36] A. E. Niell, "Global mapping functions for the atmosphere delay at radio wavelengths," *Journal of Geophysical Research*, vol. 101, pp. 3227–3246, 1996.
- [37] S. Stein, "Algorithms for ambiguity function processing," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 29, no. 3, pp. 588–599, June 1981.
- [38] J. A. Nanzer, M. D. Sharp, and D. Richard Brown, "Bandpass signal design for passive time delay estimation," in *2016 50th Asilomar Conference on Signals, Systems and Computers*, Nov. 2016, pp. 1086–1091.
- [39] M. Skolnik, *Introduction to radar systems 2nd Edition*. McGraw-Hill, 1980.
- [40] J. R. Vetter, "Fifty years of orbit determination," *Johns Hopkins APL technical digest*, vol. 27, no. 3, p. 239, 2007.
- [41] Z. Clements, P. Ellis, M. L. Psiaki, and T. E. Humphreys, "Geolocation of terrestrial GNSS spoofing signals from low Earth orbit," in *Proceedings of the ION GNSS+ Meeting*, Denver, CO, 2022, pp. 3418–3431.

Maximum Likelihood Time of Arrival and Doppler Estimation for Precise Starlink-Based PNT

Wenkai Qin*, Zacharias M. Komodromos[†], Samuel C. Morgan*, Todd E. Humphreys*

**Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin*

[†]*Department of Electrical and Computer Engineering, The University of Texas at Austin*

Abstract—We present a maximum likelihood (ML) Doppler and time-of-arrival (TOA) estimation framework for opportunistic tracking of Starlink downlink signals. Extending previous approaches that rely solely on known pilot symbols, we incorporate full-frame ML estimation to harness the data payload, significantly improving Doppler and TOA estimation accuracy. Using live Starlink transmissions, we validate our ML estimator and compare its performance against pilot-based cross-ambiguity function (CAF) and pilot-only ML estimation methods. Results show that the full-frame ML estimator achieves a 10^3 factor improvement in Doppler accuracy over the pilot-only CAF method and 10^2 factor improvement over the pilot-only ML method, reducing post-fit residual RMSE from 1469.20 Hz and 752.43 to 6.34 Hz, respectively. TOA estimation sees a smaller improvement. The findings highlight the value of leveraging the entire OFDM frame for estimation. Additionally, we newly identify two OFDM symbol modulation schemes in use by Starlink.

I. INTRODUCTION

The past decade has seen an explosion of satellite operations. In 2013, only 1,187 active satellites orbited the Earth; as of 2023, this number has ballooned to 9,115 [1]. This rapid growth can be in part attributed to a historic interest in low Earth orbit (LEO) satellites for communications—companies such as SpaceX, OneWeb, and Amazon are pursuing the promise of high-bandwidth, worldwide internet as thousands more vehicles are set to be launched into orbit. Meanwhile, traditional global navigation satellite systems (GNSSs) are facing challenges that threaten its reliability. Due to their low power and wide-area broadcast nature, GNSS signals are inherently susceptible to jamming and interference [2], [3]. Further, GNSSs' constellations have a low vehicle count and are relied upon by many military applications, making them attractive targets in the age of anti-satellite warfare [4], [5]. Recent events have further emphasized the need to protect, augment, and toughen traditional GNSS, as evidenced by the rise of GNSS jamming and spoofing attacks in Europe Mediterranean [6], [7].

LEO communications satellites offer a unique opportunity for GNSS augmentation. Due to their stronger signals, wider bandwidth, and large vehicle count, LEO communications downlinks are inherently resilient to jamming, spoofing, and anti-satellite warfare. As such, researchers have already begun to investigate the feasibility and performance of LEO position, navigation, and timing (PNT) systems [8], [9]. Other researchers exploring the possibility of free-to-use LEO PNT have demonstrated meter-level positioning by opportunistically

tracking the Doppler of various communications constellations' transmissions [10]–[15]. Others have noted that compared to pseudorange-based PNT techniques, Doppler-based techniques have worse timing accuracy by many orders of magnitude (milliseconds vs. nanoseconds), even under optimistic measurement noise and satellite clock offset rate assumptions [16]–[18]. Recognizing that many PNT applications of practical interest require accurate timing, this paper provides a method for improved observable extraction for both Doppler- and pseudorange-based approaches to PNT.

The opportunistic Starlink PNT receivers currently available in literature extract PNT observables via (1) frequency tracking of leakage tones present between channels, (2) matched filtering against known pilot symbols unveiled by [19] or (3) frequently observed portions of the Starlink frame [20]. The authors of [21] present theoretical root-mean-square error (RMSE) bounds based only on the pilot symbols mentioned by item (2). However, these papers fail to fully achieve the estimate precision available within the Starlink signal, instead exploiting only two of the 302 OFDM symbols transmitted per frame. The authors of [22] demonstrate the importance of harnessing the full OFDM frame in Doppler estimation. When correlating against the template frame presented in [23], the Doppler ambiguity function's first nulls occur at ± 750 Hz, a drastic improvement when compared to the ± 113 kHz achievable when correlating against only two pilot symbols. Nonetheless, the reported performance is only available when a Starlink SV transmits the template frame rather than user data, which necessarily interferes with its communications mission. Quite possibly, the template frame is transmitted less frequently in areas and during times of high Starlink user data traffic. The theoretical framework presented in [24] presents possible maximum-likelihood (ML) and decision-directed estimation framework but falls short of its implementation or empirical validation.

To address these deficiencies in existing opportunistic LEO PNT systems, we apply the theory presented by [24] to implement, demonstrate, and evaluate a Starlink-based, ML Doppler and TOA estimation algorithm. The paper is organized as follows: Section II introduces the signal model. Section III defines the full-frame ML estimation framework. Section IV describes the signal capture system used to generate the results in this paper. Section V provides a detailed explanation of the estimation framework's practical implementation. Section VI presents the key results. Finally, Section VII closes the paper.




Fig. 1: Frame layout for the Ku-band Starlink downlink along time-frequency dimensions, from [19]. Indices along the horizontal axis enumerate the 303 intervals that constitute a single frame.

II. SIGNAL MODEL

This section first reiterates key theoretical concepts and measurement models presented by [24], then adapts them for realistic application to Starlink transmissions.

As shown by [19], the Starlink Ku-band downlink utilizes orthogonal frequency-division multiplexing (OFDM) modulation, widely used for high-speed wireless communication systems due to its robustness against multipath interference, efficient spectrum utilization, and adaptability to varying channel conditions. While the full details are presented in [19], we present a brief overview of the Starlink signal structure for the reader's convenience. Fig. 1, replicated from [19], offers a reference for the Starlink frame layout.

Starlink OFDM signals are transmitted over one of eight channels, which altogether span the 10.7 – 12.7 GHz band allocated for the constellation's downlink transmissions. The i th channel's center carrier frequency F_{ci} can be derived as $F_{ci} = 10.7 + \frac{F}{2} + F_{\delta}(i - 0.5)$ MHz, where F is the OFDM subcarrier spacing and $F_{\delta} = 250$ MHz is the Starlink channel spacing. A single Starlink OFDM symbol transmitted over any single channel is composed of $N = 1024$ mutually orthogonal data subcarriers spread across its $F_s = 240$ MHz channel bandwidth, resulting in a subcarrier spacing of $F = 234,375$ Hz. Let k represent a single subcarrier's assigned index such that $k \in \mathcal{K} = \{0, 1, \dots, N - 1\}$, and d_k represent the subcarrier offset in frequency from the carrier frequency in units of subcarriers. The subcarrier offset is defined as $d_k = k$ for $k = 0, 1, \dots, \frac{K}{2} - 1$ and $d_k = k - N$ for $k = \frac{K}{2}, \frac{K}{2} + 1, \dots, K - 1$. Then, one may calculate i th channel's k th subcarrier center frequency as $F_{csi k} = F_{ci} + d_k F$.

In OFDM, the k th subcarrier is modulated by a complex-valued frequency-domain coefficient $x_k \in \mathbb{C}$, which encodes one or more bits of information depending on the OFDM modulation scheme (e.g. 1 for BPSK, 2 for QPSK, 3 for 8PSK, etc.). Let $\mathbf{x} \triangleq [x_0, x_1, \dots, x_{N-1}]^T$ represent the payload of an entire OFDM symbol, which may contain either pilot resources, data resources, or nothing (i.e. $\mathbf{x} = \mathbf{0}$). The set of OFDM modulation schemes previously known to be used

by Starlink are QPSK, 4QAM, and 16QAM [19]. However, this paper shows as of January 2025, $\pi/4$ PSK and 32QAM are now used in addition to the aforementioned modulation schemes. In this paper, we differentiate between QPSK and 4QAM OFDM modulation schemes as follows: whereas QPSK refers to the scheme using x_k drawn from the axis-aligned symbol constellation $\mathcal{C}^{\text{QPSK}} \triangleq \{1, j, -1, -j\}$, 4QAM refers to as the scheme using coefficients drawn from the 45° -rotated constellation $\mathcal{C}^{\text{4QAM}} \triangleq \{e^{j\pi/4}, e^{3j\pi/4}, e^{5j\pi/4}, e^{7j\pi/4}\}$.

Suppose a single Starlink OFDM symbol is sent over the i th channel with channel coefficients that are constant across frequency and time. The signal experiences LOS time delay τ , phase shift ϕ , and complex additive white Gaussian noise (AWGN) $v_k \sim \mathcal{CN}(0, \sigma^2)$. Assuming negligible intercarrier interference (ICI) due to small Doppler, and intersymbol interference (ISI) due to a sufficiently long cyclic prefix, the baseband received signal y_k is modeled in the frequency domain as

$$y_k = \alpha_k x_k + v_k \quad (1)$$

$$\alpha_k = \sqrt{g} \exp(-j2\pi d_k F \tau + j\phi) \quad (2)$$

where g is the channel gain. The negligible-ICI assumption, inappropriate for the LEO-Earth wireless channel, will be discussed later in this section.

In all wireless OFDM protocols, sequences of OFDM symbols are packaged sequentially in time into groups, commonly referred to as frames. A Starlink downlink frame is composed of $N_{sf} = 302$ OFDM symbols, each of length $T_{sym} = 4.4 \mu s$, plus a frame guard interval T_{fg} for a total frame period of $T_f = 1/750$ s. Let x_{mk} represent the modulation coefficient of the m th OFDM symbol's k th subcarrier, where $m \in \mathcal{M} = \{0, 1, \dots, N_{sf} - 1\}$, and let $\mathbf{x}_m$ be defined similarly to $\mathbf{x}$. All Starlink frames begin with two known OFDM pilot symbols, referred to as the primary synchronization sequence (PSS), which appears at $m = 0$, and the secondary synchronization sequence (SSS), which appears at $m = 1$. Let $\mathcal{M}_p = \{0, 1\}$ represent the subset of $\mathcal{M}$ such that $\mathbf{x}_m$ contains known pilot resources, and $\mathcal{M}_d = \{2, 3, \dots, N_{sf} - 1\}$ the subset

such that x_m contains unknown data resources. Note that $\mathcal{M} = \mathcal{M}_p \cup \mathcal{M}_d$. While the PSS is natively represented in the time domain, the SSS is formatted as a standard QPSK OFDM symbol [19]. The remaining 300 nonzero OFDM symbols $\{x_m \mid 2 \leq m \leq N_{sf} - 1\}$ were previously thought to contain unknown data symbols; however, other work has shown that these symbols may be more predictable than previously believed [19], [22], [23]. As such, one could likely construct an *a priori* probability function for the likelihood that any single x_{mk} takes on value $c_{\text{sym}} \in \mathcal{C}$ as

$$P(x_{mk} = c_{\text{sym}}) \forall c_{\text{sym}} \in \mathcal{C} \mid \sum_{c_{\text{sym}} \in \mathcal{C}} P(x_{mk} = c_{\text{sym}}) = 1 \quad (3)$$

Distillation of (3) for all $m \in \mathcal{M}, k \in \mathcal{K}, c_{\text{sym}} \in \mathcal{C}$ would likely require mass decoding for a consistency study across hundreds, if not thousands, of Starlink frames.

Now, consider an entire Starlink frame sent over the i th channel that experiences Doppler effects with related Doppler shift frequency F_D . This frequency shift is parameterized as the carrier frequency offset (CFO) parameter $\beta \triangleq -F_D/F_{ci}$, but note that Doppler effects arising from the considerable relative motion between a Starlink satellite vehicle (SV) and stationary ground receiver are not limited to a Doppler frequency shift—the signal also experiences a time-domain compression/dilation of the baseband signal, which [19] showed was non-negligible for LEO communications constellations transmitting OFDM frames with frame durations on the order of 10^{-3} s and bandwidth on the order of 10^8 Hz. As such, the combined Doppler effects across a Starlink frame with CFO parameter β must be modeled as

$$\tau_m = \tau_0 - mT_{\text{sym}} \left(\frac{\beta}{1 - \beta} \right) \quad (4)$$

$$\approx \tau_0 - mT_{\text{sym}}\beta \quad (5)$$

$$\phi_m = \phi_0 + mT_{\text{sym}}\beta F_{ci} \quad (6)$$

where τ_m and ϕ_m are the time delay and phase shift experienced by the m th OFDM symbol, respectively. Accordingly, τ_0 and ϕ_0 are the time delay and phase shift experienced by the OFDM symbol with index $m = 0$, i.e. the PSS. As experimental results show that Starlink channel β -values channel reside on the order of 10^{-6} , the approximation $\frac{\beta}{1 - \beta} \approx \beta$ is valid, considering that $|\beta| \ll 1$.

However, such values of β remain too large to assume negligible ICI: assuming a signal transmitted on channel $i = 5$ (with $F_{ci} = 11.825$ GHz) experiences Doppler effects with CFO parameter $\beta = 5 \times 10^{-6}$, the Doppler frequency shift of $F_D = \beta F_{ci} = 59,125$ Hz is a significant 24.7% of the subcarrier spacing $F = 234,275$ Hz. As such, assume that a receiver first conducts a coarse Doppler acquisition granting *a priori* CFO parameter estimate β_0 such that $\beta = \beta_0 + \delta\beta$, where the error ratio $\delta\beta/\beta_0$ can be as large as approximately 10% in magnitude.

Then, the baseband received signal for the m th OFDM symbol's k th subcarrier modulation y_{mk} can be modeled in

the frequency domain as

$$y_{mk} = \alpha_{mk}x_{mk} + v_{mk} \quad (7)$$

$$\alpha_{mk} = \sqrt{g} \exp(-j2\pi d_k F \tau_m + j\phi_m) \quad (8)$$

$$\tau_m = \tau_0 - mT_{\text{sym}}\delta\beta \quad (9)$$

$$\phi_m = \phi_0 + mT_{\text{sym}}\delta\beta F_{ci} \quad (10)$$

where x_{mk} is the m th OFDM symbol's k th subcarrier modulation coefficient, AWGN $v_{mk} \sim \mathcal{CN}(0, \sigma^2)$, and the model for τ_m makes a short sequence of approximations under the fact that $\beta_0, \beta_m \ll 1$.

III. ESTIMATION FRAMEWORK

Following the lead of [24, Sections III, IV], this section derives a generalized ML estimator for the unknown time delay and phase shift experienced by a single, noisy OFDM symbol. It then derives log-likelihood functions specific for respective use in the pilot-only and data-only ML estimators. It then extends the framework described in [24], which does not consider Doppler, by showing how the CFO parameter β can be estimated using a sequence of ML estimates $\hat{\theta}$.

A. Maximum Likelihood Estimation

In general, given a probability distribution function $\Lambda_Y(\theta) \triangleq p(Y|\theta)$ that describes the likelihood of observing the set of measurements $Y \in \mathbb{R}^{N_y}$ as it relates to unknown parameter vector $\theta \in \mathbb{R}^{N_\theta}$, the ML estimate for θ is the value of θ that maximizes $\Lambda_Y(\theta)$. This is typically written as

$$\hat{\theta}_{\text{ML}} = \underset{\theta}{\operatorname{argmax}} \Lambda_Y(\theta) \quad (11)$$

Maximization of the log-likelihood function $\log \Lambda_Y(\theta)$ is often a convenient and equivalent substitute for maximization of the simple likelihood function $\Lambda_Y(\theta)$ in the framework of ML estimation, as (1) $\Lambda_Y(\theta)$ often involves products of probability functions and (2) the natural logarithm is a homogeneous function.

$$\hat{\theta}_{\text{ML}} = \underset{\theta}{\operatorname{argmax}} \log \Lambda_Y(\theta) \quad (12)$$

Suppose a receiver receives and Doppler-precompensates, using an *a priori* value β_0 , a noisy OFDM symbol with symbol index $m \in \mathcal{M}$ and frequency-domain representation $y_m \triangleq [y_{m0}, y_{m1}, \dots, y_{m(N-1)}]^T$. The signal y_m experiences unknown time delay τ and phase shift ϕ , where the measurement model for y_{mk} is defined by (7-10). In this context, the measurement vector can then be defined as $Y \triangleq y_m$, and the unknown parameter vector as $\theta \triangleq [\tau, \phi]^T$. Assuming that v_{mk} are independent $\forall \{k_1, k_2 \in \mathcal{K} \mid k_1 \neq k_2\}$, the likelihood function can be written and subsequently decomposed as

$$\Lambda_{y_m}(\theta) = p(y_m|\theta) \quad (13)$$

$$= \prod_{k \in \mathcal{K}} p(y_{mk}|\theta), \quad (14)$$

giving rise to log-likelihood function

$$\log \Lambda_{y_m}(\theta) = \sum_{k \in \mathcal{K}} \log p(y_{mk}|\theta) \quad (15)$$

and ML estimate

$$\hat{\theta}_{\text{ML}} = \underset{\theta}{\operatorname{argmax}} \sum_{k \in \mathcal{K}} \log p(y_{mk} | \theta) \quad (16)$$

With the single-symbol ML estimator now formulated, the log-likelihood function $\log p(y_{mk} | \theta)$ needs to be derived under two separate conditions: (1) when x_m carries unknown data resources, and (2) when x_m carries known pilot resources.

B. Log-Likelihood Derivation

First, let τ be normalized by the CFO-compensated OFDM sampling period $T_s = \frac{1}{(1-\beta_0)F_s} \approx \frac{1}{F_s}$, creating $z \triangleq \frac{(1-\beta_0)\tau}{T_s} \approx \frac{\tau}{T_s}$ where z is in units of samples.

1) *Unknown Data Resources:* Consider the case where x_m contains unknown data resources. The likelihood of y_{mk} conditioned on the parameter vector θ and knowledge of x_{mk} is

$$p(y_{mk} | x_{mk}, \theta) \quad (17)$$

$$= \frac{1}{\pi\sigma^2} \exp\left(\frac{-1}{\sigma^2} |y_{mk} - \mu_{mk}\nu_k(\theta)|^2\right),$$

where $\nu_k(\theta) \triangleq \exp(-j2\pi z d_k/N + j\phi)$ and $\mu_{mk} \triangleq \sqrt{g}x_{mk}$. Assuming a known prior distribution function for c_{sym} expressed as $p(x_{mk} = c_{\text{sym}})$, which is trivially independent from the unknown parameter θ , $p(y_{mk} | \theta)$ is decomposed as

$$p(y_{mk} | \theta) \quad (18)$$

$$= \sum_{c_{\text{sym}} \in \mathcal{C}} p(y_{mk} | x_{mk} = c_{\text{sym}}, \theta) P(x_{mk} = c_{\text{sym}})$$

As determination of $P(x_{mk} = c_{\text{sym}})$ is not within the scope of this paper, a uniform probability distribution is assumed for the distribution of subcarrier modulations such that $P(c_{\text{sym}i}) = P(c_{\text{sym}j}) = \frac{1}{|\mathcal{C}|}$ for all $c_{\text{sym}i}, c_{\text{sym}j} \in \mathcal{C}$. The likelihood of y_{mk} is then given by

$$p(y_{mk} | \theta) = \frac{1}{|\mathcal{C}|} \sum_{c_{\text{sym}} \in \mathcal{C}} p(y_{mk} | x_{mk} = c_{\text{sym}}, \theta), \quad (19)$$

which can be recognized as a Gaussian mixture distribution function. Following the steps taken in [24, Eq. 19], the log-likelihood then becomes

$$\log p(y_{mk} | \theta) \quad (20)$$

$$= \log \frac{1}{|\mathcal{C}|} + \log \sum_{c_{\text{sym}} \in \mathcal{C}} p(y_{mk} | x_{mk} = c_{\text{sym}}, \theta)$$

$$= \log \frac{1}{\pi\sigma^2|\mathcal{C}|} \quad (21)$$

$$+ \log \sum_{c_{\text{sym}} \in \mathcal{C}} \exp\left(\frac{-1}{\sigma^2} |y_{mk} - \mu_{mk}\nu_k(\theta)|^2\right)$$

$$= \log \frac{1}{\pi\sigma^2|\mathcal{C}|} - \frac{1}{\sigma^2} |y_{mk}|^2 \quad (22)$$

$$+ \log \sum_{c_{\text{sym}} \in \mathcal{C}} \exp\left(\frac{1}{\sigma^2} (2\Re\{y_{mk}^* \mu_{mk}\nu_k(\theta)\} - g|c_{\text{sym}}|^2)\right).$$

2) *Known Pilot Resources:* Now consider the case where x_m contains pilot resources, i.e. $x_{mk} = p_{mk}$ where p_{mk} is known *a priori*. Then, $\log p(y | \theta)$ is evaluated in a straightforward manner:

$$\log p(y_{mk} | \theta) \quad (23)$$

$$= \log \left(\frac{1}{\pi\sigma^2} \exp\left(\frac{-1}{\sigma^2} |y_{mk} - \mu_{mk}\nu_k(\theta)|^2\right) \right) \quad (24)$$

$$= \log \frac{1}{\pi\sigma^2} + \frac{1}{\sigma^2} (2\Re\{y_{mk}^* \mu_{mk}\nu_k(\theta)\} - |y_{mk}|^2 - g|p_{mk}|^2) \quad (25)$$

where ν_k maintains the previous definition but $\mu_{mk} \triangleq \sqrt{g}p_{mk}$.

C. Least Squares CFO Estimation

While [24] proposes a full-frame ML estimator that grants a unified estimate for τ and ϕ per frame, it fails to estimate Doppler shift F_D . A natural extension of the approach taken in [24] would augment measurement models [24, Eqs. 2, 3] with an additional phase adjustment coefficient of $\exp(j2\pi F_D)$ and carry through the same analysis to arrive at an adjusted ν_k formulation. Otherwise, the ML estimation technique remains the same as discussed. However, consider that ML estimators are often implemented using grid searches when input data is noisy, the likelihood surface is irregular, or both. In the context of a grid search, the above approach adds a factor of N_β to the search space complexity $N_\tau \cdot N_\phi$, where N_β , N_τ , and N_ϕ are the number of elements in the β , τ , and ϕ search spaces. Instead, this paper proposes a two-stage process that first makes a series of estimates for ϕ and τ , then stitches estimates together across an entire frame to make a least-squares (LS) estimate of a unitary β for the frame. This process also generates the unknown parameters τ_0 and ϕ_0 .

Suppose an entire Starlink frame is pre-compensated with β_0 , after which the single-symbol estimator is individually applied to all $y_m \forall m \in \mathcal{M}$ to construct a sequence of estimates for time delay $\hat{\tau} \triangleq [\hat{\tau}_0, \hat{\tau}_1, \dots, \hat{\tau}_{N-1}]^T$ and phase shifts $\hat{\phi} \triangleq [\hat{\phi}_0, \hat{\phi}_1, \dots, \hat{\phi}_{N-1}]^T$. Assume that the estimates are subject to AWGN error $v_\tau \sim \mathcal{N}(\mathbf{0}, \sigma_\tau^2 I)$ and $v_\phi \sim \mathcal{N}(\mathbf{0}, \sigma_\phi^2 I)$, where $v_\tau, v_\phi \in \mathbb{R}^{|\mathcal{M}|}$.

Based on the dynamics model (9, 10), one can then form a measurement model for $\hat{\tau}$ and $\hat{\phi}$ based on unknown parameters τ_0 , ϕ_0 , and $\delta\beta$.

$$\tau_m = \tau_0 - mT_{\text{sym}}\delta\beta + v_{\tau m} \quad (26)$$

$$\phi_m = \phi_0 + mT_{\text{sym}}\delta\beta F_{ci} + v_{\phi m}, \quad (27)$$

where the $(\cdot)_m$ subscript indicates the m th element of $(\cdot)$ and $(\hat{\cdot})$ notation has been dropped for clarity.

Measurement models (26) and (26) may then be used to form LS estimators for τ_0 , ϕ_0 and $\delta\beta$. While it is possible for $\delta\beta$ to be jointly estimated via (26) and (26), practical experimentation showed that solely using ϕ for LS $\delta\beta$ -estimation generally provided for better results. As such, the LS estimator is formed in a two-step process: First, ϕ is used to make LS estimates for ϕ_0 and $\delta\beta$. Second, $\delta\beta$ is assumed in estimation of τ_0 .




Fig. 2: Block diagram of the Starlink signal capture system.

The LS estimator for ϕ_0 and $\delta\beta$ is

$$\begin{bmatrix} \hat{\phi}_0 \\ \hat{\delta\beta} \end{bmatrix} = (H^T H)^{-1} H^T \phi \quad (28)$$

$$H = \begin{bmatrix} 1 & 0 \\ 1 & T_{\text{sym}} F_{ci} \\ \vdots & \vdots \\ 1 & (N_{\text{sf}} - 1) T_{\text{sym}} F_{ci} \end{bmatrix}, \quad (29)$$

and the LS estimator for τ_0 , taking $\hat{\delta\beta}$ as *a priori* from (28), is The LS estimator for ϕ_0 and $\delta\beta$ is

$$\hat{\tau}_0 = (H^T H)^{-1} H^T (\tau - \mathbf{b}) \quad (30)$$

$$H = \mathbf{1}_{|\mathcal{M}|} \quad (31)$$

$$\mathbf{b} = \hat{\delta\beta} [0, 1, \dots, N_{\text{sf}} - 1]^T, \quad (32)$$

where $\mathbf{1}_{|\mathcal{M}|}$ is the 1-vector of size $|\mathcal{M}| \times 1$.

IV. SIGNAL CAPTURE

In this paper, live Starlink transmissions are used to implement a tangible proof-of-concept demonstration of the ML estimation framework proposed by [24]. This section outlines the signal capture setup used to receive and record the signals used for software testing and validation.

Fig. 2 shows a block diagram of the signal capture system. The Starlink capture chain begins with an offset parabolic dish antenna mounted on an azimuth-elevation actuated mount that uses publicly accessible Starlink ephemerides in the form of two-line elements to steer the antenna boresight towards Starlink SVs passing overhead our signal capture site in Austin, Texas.

To avoid possible estimation performance issues stemming from low SNRs, the results presented in this paper use only dominant signals from assigned beams possessing signal-to-noise ratios (SNRs) of 16 dB or higher, using the parlance of [21]. However, future work could show the performance of the proposed estimator under varying SNR conditions to verify the theoretical and simulated estimation bounds proposed in [24]. The dish antenna's approximately 3° beamwidth helps ensure that signals captured mostly originate from a single SV. Any inadvertently-captured side beams presented as secondary signals with typical SNRs -10 dB or lower than those of

assigned beams, greatly simplifying the process of isolating a composite signal's dominant signal.

The antenna focuses captured signals onto a feedhorn connected to a low-noise block (LNB) with a conversion gain of 60 dB and a noise figure of 0.8 dB. The LNB is dual-band, downconverting either 10.7–11.7 GHz (the lower band) to 950–1950 MHz, or 11.7–12.75 GHz (the upper band) to 1100–2150 MHz. The antenna's nominal gain is 40 dBi at 12.5 GHz, but there are losses of at least 4-5 dB due to lack of a circular-to-linear polarizer and feedhorn misalignment. A downstream radio frequency signal analyzer (RFSA) conducts further bandpass filtering, 16-bit complex sampling, and digital downconversion with a total bandwidth of $F_{\text{sr}} = 250$ Msps, although only ~ 200 Msps are usable due to the RFSA's internal front-end filtering. Note that the same GPS-disciplined 10 MHz oven-controlled crystal oscillator was used to simultaneously drive the Ku-band to L-band downmixing performed by the LNB and the wideband sampling performed by the RFSA.

V. IMPLEMENTATION

This section opens with a description of two key challenges to be addressed, then details all major portions of the ML estimation framework's software implementation.

A. Key Challenges

Before delving into the specifics of software implementation, several key challenges in applying the theoretical concepts presented in Section III to live signals are highlighted. The approaches taken to address these issues are mentioned before further elaboration in Sections V-B through V-F.

1) *Unknown Modulation Scheme*: The framework presented by Section III generally assumes *a priori* knowledge of any x_m 's corresponding modulation scheme $\mathcal{C}_m$. For opportunistic receiver of Starlink signals, both these assumptions are false. Starlink is known to begin all frames with $\mathcal{C}_m = \mathcal{C}^{\text{QPSK}}$ for $m = \{1, 2, 3, 4, 5\}$ [2], but its behavior $\forall m \geq 6$ is known to be unpredictable. To provide the data-only ML estimator with the required $\mathcal{C}_m$, a modulation detector is implemented, as detailed in Section V-D.

2) *Phase Estimate Ambiguity*: Close inspection of the data-only likelihood function formulation (19) reveals that there could exist phase ambiguities of some magnitude θ_a in $\hat{\phi}$, depending on the shape of $\mathcal{C}_m$. Note that calculate of the log-likelihood function involves a sum of $p(y_{mk}|x_{mk} = c_{\text{sym}}, \theta)$ across all $c_{\text{sym}} \in \mathcal{C}$, necessarily losing information regarding which exact value of c_{sym} was the greatest contributor to $p(y_{mk}|\theta)$. Say, for example, that $\mathcal{C}_m = \mathcal{C}^{\text{BPSK}} = \{-1, 1\}$. Then,

$$p(y_{mk}|\tau, \phi) = \frac{1}{2} [p(y_{mk}|x_{mk} = -1, \tau, \phi) + p(y_{mk}|x_{mk} = 1, \tau, \phi)] \quad (33)$$

$$= \frac{1}{2} [p(y_{mk}|x_{mk} = 1, \tau, \phi + \pi) + p(y_{mk}|x_{mk} = -1, \tau, \phi + \pi)] \quad (34)$$

$$p(y_{mk}|\tau, \phi) = p(y_{mk}|\tau, \phi + \pi), \quad (35)$$

thus granting a $\theta_a = \pi$ phase estimation ambiguity in $\hat{\phi}$. One can think of this as rotating the constellation diagram of $\mathbf{y}_m$ to best fit the constellation shape of $\mathcal{C}_m$, without further care for the particular orientation of $\mathbf{y}_m$ within the confines of $\mathcal{C}_m$'s constellation shape.

For the QPSK, 4QAM, 16QAM, and 32QAM modulation schemes, this phenomenon results in a $\pi/2$ phase estimate ambiguity. For the 8PSK scheme, a $\pi/4$ phase estimate ambiguity arises. The presence of θ_a limits the entire estimator's pull-in range for β estimation, as $\beta F_{ci} T_{\text{sym}} \ll \theta_a$ is needed to prevent inter-symbol cycle slips. Hard decoding decisions may be used to resolve the issue of phase ambiguity.

B. Signal Preprocessing

Starlink signals are captured through the system described in Section IV. Before the estimation framework described in Section III can be applied, some signal preprocessing is required. First, the captured signal is resampled at the Starlink information symbol rate $F_s = 240$ MHz such that all downstream processing may proceed with minimal ambiguity.

As stated in Section II, application of the single-symbol ML estimators outlined in Section III requires Doppler-precompensation of y_{mk} to minimize ICI. As such, a local replica p_k consisting of the coherent time-domain concatenation of the PSS and SSS is used to conduct preliminary acquisition using a standard cross-ambiguity function (CAF) [25], granting coarse initial Doppler and TOA estimates β_{0n} and t_{*0n} for the n th frame originating from the dominant signal of a short, seconds-long capture. This process is outlined in Algorithm 1.

The ML estimation framework operates on a single Starlink frame at once. To isolate a single frame's signal, signal data corresponding to time range $[t_{*0}, t_{*0} + T_f]$ is extracted from a seconds-long signal recording. To pre-compensate, this single-frame signal is then resampled from F_s to $(1 - \beta_0)F_s$, and shifted by $F_D = \beta F_{ci}$, where i is found by matching the recording's center frequency F_{cr} to known Starlink channel bands as published in [19]. This grants a Doppler-compensated time domain signal corresponding to a single Starlink frame. Note that while some CFO error still exists such that $|\delta\beta| > 0$,

this level of Doppler precompensation sufficiently mitigates the effects of ICI.

Finally, the standard OFDM steps of (1) stripping the cyclic prefix and (2) taking a fast Fourier transform of the captured data are performed, granting $\mathbf{y}_m$.

Algorithm 1: Coarse Doppler-TOA Acquisition

Input: Captured signal $\mathbf{y}$, recording sampling rate F_{sr} , recording center frequency F_{cr} , Doppler search space $F_{D\text{search}}$

Output: Estimated frame start index k_0 , estimated CFO parameter β_0

Generate local synchronization sequence p_k ;

Define time vector t_{p_k} based on sampling rate;

Initialize correlation tracking arrays;

foreach Doppler shift F_D in $F_{D\text{search}}$ **do**

 Apply frequency shift: $p_k \leftarrow p_k \cdot \exp(j2\pi F_D t_{p_k})$;

 Compute cross-correlation of $\mathbf{y}$ with p_k ;

 Extract lag values and correlation peaks;

 Store maximum correlation and corresponding lag index;

end

Find Doppler shift F_D that maximizes correlation;

Compute frame start index k_0 by scaling lag index;

Compute F_{ci} from F_{cr} ; Compute $\beta_0 = -F_D/F_{ci}$;

return k_0, β_0 ;

C. Pilot Symbol Estimator

The pilot symbol estimator is used to generate $\hat{\theta}_m = [\hat{z}_m, \hat{\phi}_m]^T$ for the first two symbols of each Starlink frame, $\mathbf{y}_0$ and $\mathbf{y}_1$. First, the channel gain g is estimated as the median per-subcarrier power and the PSS and SSS are scaled to obtain $\mu_{mk} = \sqrt{g}p_{mk}$. Then, the log-likelihood $\log \Lambda_{\mathbf{y}_m}(\theta)$ is computed according to (22) for each member of a two-dimensional z - ϕ search space. Note that for the pilot symbol estimator, the search space defined as between ± 2 samples in the z -axis, and $\pm 2\pi$ rad in the ϕ -axis. An example empirical log-likelihood for the pilot symbol estimator is shown in the top plot of Fig. 3. Finally, $\hat{z}_{\text{ML}}$ and $\hat{\phi}_{\text{ML}}$ are taken as the pair corresponding to the maximum value of $\log \Lambda_{\mathbf{y}_m}(\theta)$ and stored in $\hat{\mathbf{z}}$ and $\hat{\theta}$. The processes described in Sections V-C through are outlined in Algorithm 1.

D. Modulation Scheme Detector

As previously discussed, $\mathcal{C}_m$ must be detected per-symbol before data ML estimation can proceed. First, $\mathbf{y}_m$ is rendered in constellation and adjusted using the pilot-estimated subsample delay $\hat{z}_1$ as

$$\mathbf{y}_m \leftarrow \mathbf{y}_m \exp \left(-j2\pi \frac{\hat{z}_1 d_k}{N} \right), \quad (36)$$

which helps to de-scatter the constellation representation. This effect can be observed in the transition between the left and middle plots of Figs. 4 and 5. This step is necessary for reliable differentiation between high-BPS modulation schemes, particularly 16QAM and 32QAM, for which even a subsample delay




Fig. 3: Empirical log-likelihood functions for a single pilot symbol (top) and single data symbol (bottom). Note that the data symbol's log-likelihood function is evaluated only over $0 \leq \phi < 90^\circ$, whereas the pilot symbol is evaluated over $0 \leq \phi < 360^\circ$. Nonetheless, in both cases, only a single peak is observed.

as small as $z = 0.1$ could render the constellation diagram too noisy for precise modulation determination. For each possible BPS value $n_{\text{bps}} \in \mathcal{N}_{\text{bps}} \triangleq \{2, 3, 4, 5\}$, k-means clustering is conducted with $k = 2^{n_{\text{bps}}}$ with 20 replicates, and the clustering performance is rated using the median silhouette score (MSS), which measures the overall agreement of $\mathbf{y}_m$ with its assigned cluster map [26]. MSS is calculated as

$$s[k] = \frac{b[k] - a[k]}{\max(a[k], b[k])} \quad (37)$$

$$\text{MSS} = \text{med}(s) \quad (38)$$

where:

- $s[k]$ is the silhouette score for point y_{mk} .
- $a[k]$ is the average distance from y_{mk} to all other points in its own cluster.
- $b[k]$ is the minimum average distance from y_{mk} to points in different clusters, minimized over clusters.

The $\mathcal{C}_m$ was chosen based on the n_{bps} with maximum MSS according to the following table:

Bits Per Subcarrier	Modulation Scheme
2	QPSK/4QAM
3	8PSK
4	16QAM
5	32QAM

TABLE I: Designation of modulation Scheme according to optimal BPS.

When $n_{\text{bps}} = 2$, $\mathbf{x}_m$ could adopt either a 4QAM or QPSK modulation scheme since both schemes use the same BPS. To resolve this, the algorithm initially assumes the modulation remains unchanged from the last occurrence of either 4QAM or QPSK, i.e. ML estimation is conducted under the assumption $\mathcal{C}_m = \mathcal{C}_M \in \{\mathcal{C}^{\text{QPSK}}, \mathcal{C}^{\text{4QAM}}\}$, where $\mathcal{C}_M$ is the most recent occurrence of QPSK or 4QAM modulation.

If this assumption is incorrect, the estimated phase of the current symbol $\hat{\phi}_m$ should exhibit a characteristic phase shift of approximately $\pi/4$ rad relative to $\hat{\phi}_{m-1}$. To verify whether a modulation change has occurred, the angle difference $\Delta\theta_{\pi/2}(\hat{\phi}_m, \hat{\phi}_{m-1})$ is measured after data symbol estimation, where $\Delta\theta_{\pi/2}$ is the angle difference function under $\pi/2$ wrapping as defined below.

$$\Delta\theta_{\pi/2}(\theta_1, \theta_2) = \text{mod}\left(\theta_2 - \theta_1 + \frac{\pi}{4}, \frac{\pi}{2}\right) - \frac{\pi}{4} \quad (39)$$

If $\Delta\theta_{\pi/2}(\hat{\phi}_m, \hat{\phi}_{m-1})$ is greater than a predefined tolerance θ_{tol} , $\hat{\phi}_m$ and $\mathcal{C}_m$ are retroactively adjusted to fit the new hypothesis.

E. Data Symbol Estimator

The data symbol estimator takes essentially the same steps as the pilot symbol estimator with two exceptions: (1) An additional software loop sums over all possible $c_{\text{sym}} \in \mathcal{C}_m$ to calculate $\log \Lambda_{\mathbf{y}_m}(\boldsymbol{\theta})$, formulated according to (25). This is practically implemented as a matrix summation for speed. (2) The search space is defined as between $\pm\pi/2$ rad in the ϕ -axis, because $p(y_{mk}|\tau, \phi) = p(y_{mk}|\tau, \phi + n\pi/2)$ for 4QAM, QPSK, 16QAM, 32QAM as previously discussed. An example empirical log-likelihood for the data symbol estimator is shown in the bottom plot of Fig. 3. Together, the modulation scheme detector and data symbol estimator are run per-symbol for all data symbols, granting $\hat{\boldsymbol{\theta}}_m = [\hat{z}_m, \hat{\phi}_m]^\top \forall m \in \{2, 3, \dots, N_{\text{sf}} - 1\}$.

F. Doppler-TOA Extraction

Remembering that $\hat{\phi}$ exhibits an $n\pi/2$ ambiguity, $\hat{\boldsymbol{\theta}} = [\hat{z}, \hat{\phi}]^\top$ are first preprocessed to eliminate outliers and unwrap the phase estimate before use for Doppler-TOA extraction. Outliers are identified via a RANSAC one-dimensional polynomial fit on z . Use of the phase-wrapped $\hat{\phi}$ (blue circles in Fig. 6) for outlier identification is unfavorable, as the data follows a sawtooth waveform with sharp jumps at $\hat{\phi} = \pi/2$. Outliers often arose due to insufficient SNR, leading to misidentification of the modulation scheme and/or improper realignment of $\mathbf{y}_m$'s constellation with that of $\mathcal{C}_m$. Examples




Fig. 4: Tripanel constellation representation of single OFDM symbol $\mathbf{y}_m$ with $\mathcal{C}_m = \mathcal{C}^{\text{QPSK}}$ after standard OFDM processing (left), pre-compensation with $\hat{z}_1$ from SSS ML estimate (center), and after correction with $\hat{z}_m$ and $\hat{\phi}_m$ (right). Colors indicate k-means clusters. In the middle panel, the black squares show the centroid of each k-means cluster. In the right panel, the black circles show the position of $\mathcal{C}_m$.




Fig. 5: As Fig. 4 but for an OFDM symbol where $\mathcal{C}_m = \mathcal{C}^{16\text{QAM}}$.

of excluded data can be seen in Fig. 6 as red x-marks. The remaining data $\hat{\phi}$ is then unwrapped with a $\pi/2$ wrapping point.

Using $\hat{\theta}_m = [\hat{z}_m, \hat{\phi}_m]^T$, the theory presented by Section III-C is equivalently implemented as LS one-dimensional polynomial fitting. The fit trends can be seen in Fig. 6 as yellow lines. These steps are outlined in Algorithm 3. This process generates the time delay correction $\hat{\tau}_0$ and CFO parameter correction $\delta\hat{\beta}$, which may be used to refine the initial estimates given by coarse acquisition.

$$\hat{t}_* = t_{*0} - \hat{\tau}_0 \quad (40)$$

$$\hat{\beta} = \beta_0 + \delta\hat{\beta} \quad (41)$$

The CFO parameter estimate $\hat{\beta}$ from the n th frame is used to conduct signal preprocessing for the $(n+1)$ st frame, preventing uncontrolled drift of β away from β_0 .

VI. RESULTS

A. New Starlink OFDM Modulation Schemes

Two OFDM modulation schemes were observed that have not yet been published, to the author's knowledge. 32QAM was regularly observed in captures ranging from September 2024 to January 2025, as shown in Fig. 10.

Further, what initially appeared to be 8PSK modulation in a capture taken on January 13, 2025 turned out to be subcarrier-sliced $\pi/4$, as can be appreciated from Fig. 11.

B. Doppler-TOA Tracking Performance

Using 1.7 seconds of STARLINK-1274's downlink signal data as captured on January 13, 2025, the Doppler-TOA tracking performance of the implemented ML estimator is compared against two other candidates: (1) the standard CAF-based approach using PSS+SSS, and (2) an ML estimation approach using PSS+SSS. The latter was calculated by excluding $\hat{z}_m$ and $\hat{\phi}_m \forall m \in \mathcal{M}_d$, resulting in $\hat{\mathbf{z}}, \hat{\phi} \in \mathbb{R}^{2 \times 1}$.

Algorithm 2: ML Estimation of Delay and Phase

Input: Captured signal $\mathbf{y}$, valid indices k_{valid}
Output: Delay estimate $\hat{z}$, Phase estimate $\hat{\theta}$
Perform OFDM processing on $\mathbf{y}$;
Shorten subcarriers to valid indices k_{valid} ;
Pilot symbol ML estimation:
Estimate channel gain g ;
Set search spaces Z and Φ ;
foreach symbol $\mathbf{y}_m \mid m \in \{0, 1\}$ **do**
 foreach $z \in Z$ **do**
 foreach $\phi \in \Phi$ **do**
 Compute $\log \Lambda_{\text{pilot}}(\mathbf{y}_m, z, \phi)$;
 end
 end
end
Find and save maximum likelihood estimates $\hat{z}$ and $\hat{\phi}$
to $\hat{z}[m]$ and $\hat{\theta}[m]$;
Data symbol ML estimation:
Set search spaces Z and Φ ;
foreach symbol $\mathbf{y}_m \mid m \in \{2, 3, \dots, N_{sf} - 1\}$ **do**
 Determine modulation scheme $\mathcal{C}_m$; Estimate
 channel gain g ;
 foreach $z \in Z$ **do**
 foreach $\phi \in \Phi$ **do**
 foreach $c_{\text{sym}} \in \mathcal{C}_m$ **do**
 Increment $\log \Lambda_{\text{data}}(\mathbf{y}_m, z, \phi)$;
 end
 end
 end
Find and save maximum likelihood estimates $\hat{z}$ and
 $\hat{\phi}$ to $\hat{z}[m]$ and $\hat{\theta}[m]$;
Handle 90-degree phase ambiguity adjustments;
end

Then, to estimate β for each frame, the LS approach is not applied, as a one-dimensional line fit is only minimally identifiable with just two measurements. These tracking results are shown in Figs. 8 and 9. To evaluate the estimate variation, a second-order polynomial was fit to the trace of each estimate parameterized by the frame slot number, which identifies a frame's position within the whole FAI and is labeled as N_{ag} in Fig. 7. The post-fit residual RMSEs were evaluated and are presented Table II.

While the improvement in F_D estimation is substantial, the variance in the t_* estimate shows little reduction. Notably, the t_* estimate measured manifests timing issues present upon transmission by the Starlink SV, which are known to manifest abrupt adjustments and significant jitter, even under nominal behavior in the parlance of [21]. Further analysis using a two-receiver capture system, both with known locations, and a time difference of arrival approach could possibly better reveal the precision improvements afforded by the ML technique presented.

While these RMSE values are not yet close to the theoretical Cramer-Rao bounds reported in [21], there is still practical improvement over the current standard CAF-based methods.




Fig. 6: Example full-frame results for $\hat{z}$ (top), $\hat{\phi}$ (bottom). Note the excluded estimates in red, which all originate from 32QAM modulation schemes, as well as the sawtooth shape of the wrapped $\hat{\phi}$. The unwrapped $\hat{\phi}$ is seeded with value $\hat{\phi}_1$, the phase shift estimated using the PSS-based pilot symbol estimator, which does not exhibit the $\pi/2$ phase ambiguity characteristic of the data symbol estimator.

Algorithm 3: Extract Doppler-Delay Adjustments

Input: Delay estimate $\hat{z}$, Phase estimate $\hat{\theta}$
Output: Initial sample offset z_0 , initial phase offset ϕ_0 , CFO parameter adjustment $\delta\beta$
Remove outliers:
Perform RANSAC on $\hat{z}$ to get m_{inlier} ;
Invert m_{inlier} to get m_{excluded} ;
Remove excluded estimates from ϕ, z ;
LS estimation:
Perform polyfit($\hat{z}$) to get fit coefficients p_z ;
 $\phi = \text{unwrap90}(\phi)$;
Perform polyfit($\hat{\phi}$) to get fit coefficients p_ϕ ;
Compute F_{ci} from F_{cr} ; Compute $\delta\beta = -F_D/F_{ci}$;
Set $z_0 = p_z(2)$;
Set $\phi_0 = p_\phi(2)$;

The marked improvement of the full-frame ML estimator over the pilot-only ML estimator shows the value of harnessing the full frame to estimate Doppler and TOA.

VII. CONCLUSION

This paper presents a maximum likelihood Doppler and time-of-arrival estimation framework for opportunistic tracking of Starlink signals, extending previous work by incorporating both pilot and data symbols into the estimation process. Through theoretical derivation, practical implementation, and empirical validation using live Starlink signals, we demonstrate a significant improvement in Doppler estimation accuracy compared to standard pilot-only approaches. Our results show that full-frame ML estimation substantially reduces




Fig. 7: Frame sequence timing diagram showing the transition from the $(l-1)$ th FAI to the l th FAI.

Method	F_D RMSE (Hz)	t_* RMSE (ns)
Pilot-Based CAF	1469.20	2.090
Pilot-Only ML	752.43	1.629
Full-Frame ML	6.34	1.626

TABLE II: Post-fit residual RMSEs by estimation method.

Doppler estimation RMSE, highlighting the advantage of leveraging data payloads in addition to known pilot symbols. However, while TOA estimation benefits from the proposed approach, its improvement remains modest, suggesting the need for further refinements such as super-resolution decision-directed techniques. Future work will focus on refining the estimation process under varying SNR conditions and investigating advanced data decoding strategies to further enhance positioning accuracy.


ACKNOWLEDGMENTS

Research was supported by the Department of Defense through the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program, the U.S. Department of Transportation under Grant 69A3552348327 for the CAR-MEN+ 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.


REFERENCES

- [1] "Number of active satellites from 1957 to 2022," <https://www.statista.com/statistics/897719/number-of-active-satellites-by-year/>, 2023, accessed: 2023-10-20.
- [2] T. E. Humphreys, "Interference," in *Springer Handbook of Global Navigation Satellite Systems*. Springer International Publishing, 2017, pp. 469–503.
- [3] R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. E. Humphreys, "Know your enemy: Signal characteristics of civil GPS jammers," *GPS World*, Jan. 2012.
- [4] E. C. Dolman, "New frontiers, old realities," *Strategic Studies Quarterly*, vol. 6, no. 1, pp. 78–96, 2012.
- [5] W. J. Broad and D. E. Sanger, "China tests anti-satellite weapon, unnerving U.S.," <https://www.nytimes.com/2007/01/18/world/asia/18cnd-china.html>, January 2007, accessed: 2023-11-1.
- [6] S. Gebrekidan, K. R. Lai, P. Robles, and J. White, "Why GPS is under attack," <https://www.nytimes.com/interactive/2024/07/02/world/gps-threats.html>, July 2024, accessed: 2024-10-23.
- [7] R. Huntley, "GNSS jamming and spoofing are a daily occurrence," Oct 2024. [Online]. Available: <https://www.eetimes.eu/gnss-jamming-and-spoofing-are-a-daily-occurrence/>
- [8] P. A. Iannucci and T. E. Humphreys, "Fused low-Earth-orbit GNSS," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–1, 2022.


- [9] L. Ries, M. C. Limon, F.-C. Grec, M. Anghileri, R. Prieto-Cerdeira, F. Abel, J. Miguez, J. V. Perello-Gisbert, S. D'Addio, R. Ioannidis, A. Ostilio, M. Rapisarda, R. Sarnadas, and P. Testani, "LEO-PNT for augmenting europe's space-based PNT capabilities," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, 2023, pp. 329–337.
- [10] Z. M. Kassas, N. Khairallah, and S. Kozhaya, "Ad astra: Simultaneous tracking and navigation with megaconstellation LEO satellites," *IEEE Aerospace and Electronic Systems Magazine*, 2024.
- [11] M. Neinavaie and Z. M. Kassas, "Cognitive sensing and navigation with unknown OFDM signals with application to terrestrial 5G and Starlink LEO satellites," *IEEE Journal on Selected Areas in Communications*, 2023.
- [12] N. Jardak and R. Adam, "Practical use of Starlink downlink tones for positioning," *Sensors*, vol. 23, no. 6, p. 3234, 2023.
- [13] C. Yang and A. Soloviev, "Starlink Doppler and Doppler rate estimation via coherent combining of multiple tones for opportunistic positioning," in *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, 2023, pp. 1143–1153.
- [14] S. Kozhaya and Z. M. Kassas, "On the fundamental tracking performance and design considerations of radio navigation," *IEEE Journal on Selected Areas in Communications*, 2024.
- [15] H. Sallouha, S. Saleh, S. De Bast, Z. Cui, S. Pollin, and H. Wymeersch, "On the ground and in the sky: A tutorial on radio localization in ground-air-space networks," *IEEE Communications Surveys and Tutorials*, vol. 27, no. 1, pp. 218–258, 2025.
- [16] M. L. Psiaki, "Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids," *NAVIGATION*, vol. 68, no. 3, pp. 621–641, 2021.
- [17] B. McLemore and M. L. Psiaki, "Navigation using Doppler shift from LEO constellations and INS data," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 5, pp. 4295–4314, 2022.
- [18] A. Baron, P. Gurfil, and H. Rotstein, "Implementation and accuracy of Doppler navigation with LEO satellites," *NAVIGATION: Journal of the Institute of Navigation*, vol. 71, no. 2, 2024.
- [19] T. E. Humphreys, P. A. Iannucci, Z. M. Komodromos, and A. M. Graff, "Signal structure of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1–16, 2023.
- [20] S. Kozhaya, S. Joe, and Z. M. Kassas, "Unveiling Starlink for PNT," *NAVIGATION*, vol. 72, no. 1, 2025.
- [21] W. Qin, A. M. Graff, Z. L. Clements, Z. M. Komodromos, and T. E. Humphreys, "Timing properties of the Starlink Ku-band downlink," *IEEE Transactions on Aerospace and Electronic Systems*, 2025, submitted for review.
- [22] S. Kozhaya, J. Saroufim, and Z. M. Kassas, "Starlink for PNT: A trick or a treat?" in *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, 2024.
- [23] Z. M. Komodromos, W. Qin, and T. E. Humphreys, "Weak signal acquisition and tracking of the Starlink Ku-Band downlink to enable global PNT," in *Proceedings of the ION Joint Navigation Conference (JNC)*, 2024.
- [24] 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.
- [25] A. Rihaczek, *Principles of high-resolution radar*. McGraw-Hill, 1969.
- [26] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of computational and applied mathematics*, vol. 20, pp. 53–65, 1987.



(a) Pilot-based CAF acquisition. Equivalent to coarse acquisition using the coherent concatenated PSS + SSS local replica.




(b) Pilot-only ML estimation. In this case, $\hat{z}, \hat{\phi} \in \mathbb{R}^{2 \times 1}$, the minimum number of measurements to estimate β .




(c) Full-frame ML estimation. The full technique as proposed in Section III


Fig. 8: β tracking performance using various estimation techniques. All axes are matched across plots.



(a) See caption of Fig. 8a.



(b) See caption of Fig. 8b.



(c) See caption of Fig. 8c.

Fig. 9: As Fig. 9 but for t_* tracking performance.




Fig. 10: Example of 32QAM modulation, observed in a decoded signal originally captured in January 2025.




Fig. 11: Side and top views of subcarrier-sliced $\pi/4$ PSK. The top view (top) is the typical constellation representation in the complex plane, where the symbol seems to take on 8PSK modulation. However, closer inspection of the side view (bottom) reveals that the modulation clearly switches between 4QAM and QPSK modulations twice, at two discrete subcarrier indices.

HOOC-EM: Fast Beam Sweeping for LEO Mega-Constellation Customer Terminals

Samuel C. Morgan and Todd E. Humphreys

Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin

BIOGRAPHIES

Samuel C. Morgan (B.S.E., Murray State University; M.S., Auburn University) is a Ph.D. student in the department of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin, and a member of the UT Radionavigation Laboratory. His research interests include fused LEO GNSS and signal processing.

Todd E. Humphreys (B.S., M.S., Utah State University; Ph.D., Cornell University) holds the Ashley H. Priddy Centennial Professorship in Engineering in the Department of Aerospace Engineering and Engineering Mechanics at The University of Texas at Austin. He is Director of the Wireless Networking and Communications Group and of the UT Radionavigation Laboratory, where he specializes in the application of optimal detection and estimation techniques to positioning, navigation, and timing. His awards include the UT Regents' Outstanding Teaching Award (2012), the NSF CAREER Award (2015), the ION Thurlow Award (2015), the PECASE (NSF, 2019), and the ION Kepler Award (2023). He is Fellow of the Institute of Navigation and of the Royal Institute of Navigation.

ABSTRACT

The popularity of low Earth orbit (LEO) satellite broadband networks has grown significantly in the past decade. These LEO networks transmit at high carrier frequencies due to spectral availability, resulting in narrow beamforming at the satellite and customer terminal (CT). However, existing methods for beam sweeping have only been developed for terrestrial mmWave networks (e.g., IEEE 802.11ad, IEEE 802.11ay, IEEE 802.15.3c, and the 3GPP's 5G NR). These standards call for exhaustive beam sweeping procedures, selecting the best candidate beam pair after each azimuth and elevation is searched. Such methods are unsuitable for LEO network CTs due to satellite dynamics and frequent satellite handoff. This paper proposes a maximum likelihood (ML) technique for CT beam sweeping in LEO broadband networks. The proposed method, called hexagonal odds-based orbital comb and entrainment method (HOOC-EM), first divides the celestial sphere into hexagonal and pentagonal sectors, according to CT beamwidth. A probability density function (PDF) is applied to the sectors and each sector is searched in a ML fashion (i.e., the most likely sector is searched first and so on) stopping once a signal is detected. An empirical *a priori* PDF and a Markov model are developed using orbital simulations of the constellation. The *a priori* PDF is dynamically updated according to the Markov model to account for satellite dynamics. This paper compares HOOC-EM to other ML methods informed by two other *a priori* PDFs, one uniform, and one linear with respect to elevation. These latter methods are not updated dynamically according to the Markov model. HOOC-EM is also compared to a traditional exhaustive raster scan beam sweep. The performance of the proposed method is analyzed by comparing the average time to detection as a function of receiver beamwidth. Numerical results indicate that HOOC-EM is an effective method for LEO broadband beam sweeping and superior to other techniques in cases where the Markov model is recent.

I. INTRODUCTION

High data throughput has rapidly become a necessity, enabling high-quality video streaming, efficient remote work, and realistic, low-latency extended and/or augmented reality experiences, among other applications. To meet the growing demand, broadband network providers have sought to increase data throughput. Achieving high data throughput requires changes to the physical (PHY) layer, particularly in terms of the bandwidth of the transmitted radio frequency (RF) signal, which is directly related to the maximum data rate. Therefore, a wide bandwidth is required to enable high data throughput. To avoid in-band interference, sufficient spectral availability at the transmitted center frequency is necessary.

Lower center frequencies are often desirable because they are less susceptible to attenuation and experience less free space path loss. However, there is limited spectral availability at lower frequencies, leading to the recent popularity of transmitting at a higher center frequencies (e.g., mmWave), where there is ample available spectrum [Rappaport et al., 2014]. The downside to higher frequencies is the much lower received signal power due to attenuation and path loss. To mitigate these downsides, beamforming at both transceivers (TX/RX) is used. Beamforming allows for high frequency signals to be received at adequate power levels by focusing all the TX/RX gain in a narrow beam in a specific direction. However, this does add a spatial component to network access, complicating sustained base station (BS) and customer terminal (CT) communication and CT

network entry [Andrews et al., 2014].

Wireless communications networks prescribe various protocols for the CT to enter the network. Network entry is defined as the CT being able to send and receive data as a fully integrated node. This implies that the network, including the CT, knows enough about the location of the CT, and time and frequency offsets, to engage in the resource allocation scheme the network uses to coordinate traffic (e.g., time and/or frequency division duplexing). The first step in resolving these unknowns and ultimately entering the network in all existing mmWave standards is beam sweeping. Beam sweeping aims to align the beams of the transmitter and receiver. In terrestrial networks this involves the receiver and/or transmitter sweeping all feasible directions using a predefined codebook [Xue et al., 2024].

In three of the IEEE’s mmWave standards (i.e., IEEE 802.11ad, IEEE 802.11ay, and IEEE 802.15.3c) a two-stage beam sweeping procedure is employed, beginning with what is called a sector-level sweep (SLS) [Nitsche et al., 2014, Ghasempour et al., 2017, Baykas et al., 2011]. In the SLS large, low-resolution beams are used to search all directions in the codebook of the CT. Once all sectors are searched, the highest power sector is then searched using a narrow, high-resolution beam, and the optimal beam pair is selected. It is worth noting that both the transmitter and receiver must perform beam sweeping and take on the roles of initiator and responder to establish a beam pair. In the 3GPP’s 5G NR standard, an exhaustive search is performed by the gNodeB and/or CT to determine the best beam pair by checking the reference signal received power and/or the signal-to-interference-plus-noise ratio (SINR) [Heng et al., 2021].

The standards discussed thus far were designed primarily for terrestrial networks, which are inherently limited by existing infrastructure. As a result, terrestrial networks often fail to provide adequate coverage in remote or rural areas. Non-terrestrial networks (NTNs) were introduced to address this issue by providing coverage anywhere on Earth by using satellites in low Earth orbit (LEO), medium Earth orbit (MEO), and geostationary orbit (GEO), together with high-altitude platforms (HAPS). In general, NTNs take two forms: direct to cell (D2C) or very short aperture terminals (VSATs). This paper focuses on NTNs that use VSATs (i.e., a CT is a VSAT). Among these, LEO satellite NTNs using VSATs are the most widespread, with SpaceX’s Starlink being the most notable example.

The beam sweeping protocols discussed in the IEEE standards present obvious challenges if applied to a LEO broadband system. Due to satellite or CT dynamics, the best beam pair at the start of a search may no longer be the best at the end. Additionally, the distance between the satellite and CT may make it difficult for any given CT to participate in the two-way initiator-responder pairing technique. Furthermore, technical issues arise when devices equipped to use the current standards experience Doppler shifts and time delays that would ordinarily be impossible for terrestrial networks, even with high CT dynamics.

The 3GPP began studying NTNs for use with 5G NR as early as Release 15 [3GPP, 2017]. They begin discussing beam management in Release 16, with a primary focus on the satellite side [3GPP, 2023]. To date, no final decisions have been made on how to handle beam management in NTN in NR and 6G, so the basic terrestrial framework developed in Release 15 is used as a baseline. Some researchers have made suggestions to improve beam management in NR NTNs [Rinaldi et al., 2020, Kim et al., 2020, Wei et al., 2023], while some test the current standard. For instance, [Artiga and Vázquez, 2023] evaluate the beam management framework of New Radio for LEO NTNs, finding that satellite tracking can be performed easily and that initial beam sweeping may only take a few seconds. However, their simulations assume a rather optimistic satellite pass (i.e., it always passes over the zenith of the CT) and rather high minimum elevation angles (68 and 51 degrees).

SpaceX’s Starlink system also requires that its CTs must perform a beam sweeping procedure to enter the network. According to [Kutkov, 2023], the Starlink CT may enter the network in two modes: GNSS-assisted or via Starlink positioning. The author states that network entry is relatively fast when GNSS is available, but can be significantly slower when relying solely on Starlink positioning. Additionally, the report indicates that Starlink positioning is not yet functional for CTs on moving platforms.

The present paper makes the following contributions:

- It introduces HOOC-EM, an efficient beam-sweeping strategy for LEO mega-constellation CTs, which pixelates the celestial sphere using hexagonal and pentagonal sectors based on beamwidth. HOOC-EM incorporates satellite dynamics to develop an empirical Markov transition matrix based on constellation simulations. This matrix is used to determine the initial probability density function (PDF) of the sectors. Then the Markov transition matrix is used to update the sector probabilities during beam sweeping to account for satellite dynamics during the search.
- It provides numerical simulations to demonstrate the effectiveness of HOOC-EM, showing that the time for beam sweeping can range from less than a second to several tens of seconds, depending on the beamwidth and simulated scenario. HOOC-EM is shown to perform better than a traditional raster scan as well as two other methods that use the same hexagonal and pentagonal sectors when the Markov model is recent.

The remainder of this paper is organized as follows. Section II contains assumptions about the broadband LEO system that will be used throughout the paper. Section III describes the HOOC-EM algorithm in detail. Section IV provides a description of the simulations conducted and shows associated results. Section V gives conclusions about the study.

Notation

Column vectors are denoted with lowercase bold, e.g., $\mathbf{x}$. Scalars are denoted without bold, e.g., x . The n th entry of a vector $\mathbf{x}$ is denoted $x(n)$. Matrices and tensors are also denoted with bold, e.g., $\mathbf{X}$. For a tensor, the m th i th k th entry is $X_{mi}(k)$, the m th i th vector is $\mathbf{X}_{mi}$, and the m th matrix is $\mathbf{X}_m$. Zero-based indexing is used throughout the paper.

II. BROADBAND LEO SYSTEM DESCRIPTION

In this paper, a conceptual LEO broadband constellation will be assumed. The constellation may be a dedicated communications network such as SpaceX’s Starlink system, or a fused LEO GNSS system as described by [Iannucci and Humphreys, 2022]. In either case, the CT must be granted access to the network at initial network entry (INE). To enter the network, the CT must have some *a priori* knowledge about the network as a whole. This section discusses the necessary assumptions this paper makes about the LEO network to enable HOOC-EM and, ultimately, full network entry.

1. Constellation Design Assumptions

In general, the CT does not need to make any assumptions about the design of the constellation. HOOC-EM is agnostic to the type of constellation, inclination angle, and even altitude. If necessary, this method could be performed on MEO and GEO constellations as well. However, since LEO is the forerunner in space-based broadband communications, this paper will focus on a LEO constellation.

Although a specific type of constellation is not required, some *a priori* information about the constellation could speed up network entry. For example, a given constellation over a given period of time may have certain azimuths and elevations that are more frequently occupied by satellites at certain latitudes. This sort of knowledge about the constellation could be used to construct a PDF that would enable a maximum-likelihood, satellite search that could take the place of a brute-force, blind satellite search. This concept is the foundation of HOOC-EM and will be the focus of this paper, with detailed discussion provided in Section 3.

2. Network Architecture Assumptions

This paper makes the following assumptions about the network. First, it is assumed that the transmitting satellites send downlink signals via directional beams, as is done by Starlink, OneWeb, and Iridium [Humphreys et al., 2023], [Blázquez-García et al., 2023], [Maine et al., 1995]. Consequently, some of the terminology described in [Qin et al., 2023] will be used in this paper as well. In particular, the terms “service cell,” “assigned beam,” and “fixed assignment interval” are of importance and will be redefined here. A service cell refers to the coverage area of a satellite beam, with the center of the service cell determined by user demand. An assigned beam is a beam directed by a satellite toward a designated service cell. A fixed assignment interval (FAI) is the interval of time over which beam assignments remain fixed across the constellation. The broadband network will be assumed to have an FAI of T_{FAI} seconds, during which a satellite transmitting an assigned beam will service the same service cell. Note that multiple satellites may be transmitting assigned beams to a given service cell. The transmitted signals within these beams are assumed to be orthogonal in at least two dimensions out of a possible four: frequency, time, polarization, and direction (azimuth and elevation). For example, each assigned beam is assumed to always be orthogonal in direction, but may share one or more of the other dimensions—frequency, time, or polarization.

The satellites may operate in either a bent-pipe (transparent) or regenerative configuration. The data and synchronization sequences sent in an assigned beam are assumed to be diverse in both frequency and time. A specific combination of frequency and time will denote a channel, and the data sent on that channel can be modulated in any way the network chooses. For example, Starlink uses orthogonal frequency division multiplexing (OFDM) and Iridium uses differentially encoded quadrature phase shift keying (DEQPSK) [Humphreys et al., 2023], [Maine et al., 1995]. Regardless of the modulation scheme, it is assumed that the transmitted signal contains a pilot or synchronization sequence of some kind. While these sequences reduce overall data capacity, they are necessary for channel equalization and symbol timing recovery. Additionally, they can be used in a matched filter for signal detection.

Finally, it is assumed that the network dedicates a time-slot specifically for network entry. In this slot the satellite transmits what is known as the synchronization signal. This signal is known to the CT and can be used within a matched filter to aid in signal detection during the beam sweeping procedure. The synchronization signal will be transmitted only every T_{SS} seconds. Note that the duration of the synchronization signal will be much shorter than T_{SS} . The remainder of the frame is assumed to be occupied by data. The synchronization periodicity T_{SS} should be large compared to the time offset caused by the distance between the satellite and the CT to enable the CT ample time for detection.

Furthermore, this paper will assume that the CT is in an area with a high number of users. That is, the CT will be in a service cell that is always illuminated by at least one transmitting satellite at some frequency band and polarization.

3. Customer Terminal Assumptions

Whenever the CT is powered on, INE will happen on a dedicated frequency band and polarization pair known to the CT. There will only be one satellite transmitting on that frequency band and polarization, specifically for INE. However, the location of that satellite is unknown to the CT when it is powered on, necessitating a satellite search in azimuth and elevation.

The CT is assumed to be equipped with a phased antenna array capable of steering a narrow beam to a desired azimuth and elevation. The array is assumed to be capable of beam width modulation, i.e., it can electronically generate beams of arbitrary beamwidths (using the definition of 3 dB beamwidth shown in Fig. 1). The CT is assumed to possess the necessary RF chains and associated hardware (low-noise amplifier, down-converter, analog to digital converter, etc.) to effectively use a phased array antenna. Additionally, the CT will have an inertial measurement unit (IMU) and a magnetometer to determine its orientation relative to the local East-North-Up coordinate frame. While the CT will be assumed to also have an onboard GNSS receiver, the receiver will be assumed inactive or inoperable in this paper. Many existing systems, including SpaceX's Starlink rely on GNSS for fast network entry. This critical reliance would have severe consequences if GNSS were unavailable, onboard receivers fail, or if there is a system-wide GNSS fault. Therefore, it is important to begin developing GNSS-free technology moving forward.




Figure 1: A notional polar plot of receiving antenna performance is provided to illustrate the definition of 3 dB beamwidth. The typical definition of beamwidth is defined at the 3 dB, or half power point [Maral et al., 2020]. In this context, D is the antenna diameter, α is an arbitrary angle defined from the max power point $G_{\max}$, and $G(\alpha)$ is the gain at α .

III. HOOC-EM ALGORITHM

The HOOC-EM algorithm is a multistage process that begins with defining an appropriate search space. Then, using knowledge of the network of interest, probabilities are assigned to the defined search space. The search space and associated probabilities are then applied to the beam sweeping problem. This section deals with each of these steps in detail.

1. Search Space

The HOOC-EM algorithm begins with defining a suitable search space. For space-to-Earth communications, it is often convenient to consider the space above the observer as a sphere with the observer at its origin. This approach leverages the fact that objects in space move in curvilinear motion. This sphere, referred to as the observer's celestial sphere, will be the search space of interest moving forward.

Most current beam sweeping approaches consider a raster scan of azimuths and elevations. A raster scan consists of searching each azimuth for a given elevation and so on. A traditional raster scan of azimuth and elevation does not account for the spherical nature of the search space (e.g., at an elevation of 90 degrees it is unnecessary to search the full 360 degrees of azimuth), resulting in a prohibitively large search. A method that uses fewer search points while still maintaining full coverage should be considered instead.

To ensure full coverage with a minimal number of points, it is useful to consider various methods of partitioning a sphere. Historically, partitioning the celestial sphere has been valuable for remote sensing applications, as it facilitates convenient computer processing [Chan, 1980, Tegmark, 1996, McCollum, 2001]. In these approaches, having numerous points within partitions of equal area was desirable. For the present use case, having fewer points is advantageous, but equal partitions are especially desirable as they better represent constant beamwidths.

Given that a large number of points is not required for this application, the centroid of each partition can be selected, with the

partition area approximated to match the desired beamwidth. This approach will be employed in this paper.

Next, an appropriate partitioning scheme must be selected. The partitioning methods developed in the aforementioned papers all adequately segment the celestial sphere. To select among them, it is important to consider how the CT's beam intersects with the celestial sphere. The cross section is predominantly circular, with some flattening as elevation approaches 0 degrees. Therefore, a partitioning scheme that approximates a circle should be favored. Scalability is also crucial: the partitioning scheme should be capable of being scaled down to represent an arbitrarily small beamwidth.

With these criteria in mind, Table 1 provides a brief overview of each partitioning scheme. It is clear from the table that the hexagonal partitioning scheme demonstrated by [McCollum, 2001] is the most suitable for the present use case. In contrast, the square partitioning scheme developed by [Chan, 1980] is the least favorable, as it lacks scalability. While some more recent work has developed scalable square partitioning schemes [Dimitrijević et al., 2016], they still fall short of the hexagonal scheme for the present use case in terms of area efficiency. The triangular scheme, or icosahedral tiling, presented by [Tegmark, 1996] is also ill-suited for the present use case due to its area characteristics. As a result, this paper will use a hexagonal partitioning scheme.

Table 1: Partition Methods: Shape, Scalability, Coverage

Method	Shape	Scalability	Fraction of Circle
[Chan, 1980]	Square	No	$\frac{2}{\pi}$
[Tegmark, 1996]	Triangle	Yes	$\frac{3\sqrt{3}}{4\pi}$
[McCollum, 2001]	Hexagon*	Yes	$\frac{3\sqrt{3}}{2\pi}$

* Hexagons and exactly 12 pentagons.

It is important to note that the triangular and hexagonal scheme are intrinsically linked. All hexagonal partitions of a sphere begin as an ordinary unit icosahedron (a polyhedron with 20 triangular faces that is circumscribed within the unit sphere). To achieve a hexagonal partitioning, each edge of the ordinary icosahedron must be truncated at its midpoint. This process results in a truncated icosahedron, which consists of 20 hexagonal faces and 12 pentagonal faces. The resulting faces can then be projected onto the unit sphere to form the partitions. Since this paper is only concerned with the centroids of the partitions, it is sufficient to project the centroid of each face to the unit sphere. This will create partitions with approximately equal area. To map the area of these partitions to beamwidth, the approximate surface area in steradians of each partition may be calculated by

$$S_{\text{partition}} \approx \frac{4\pi}{N_{\text{faces}}} \quad (1)$$

where N_{faces} is the number of faces of the polyhedron. Then, the beamwidth in degrees is approximated as

$$\theta_{3 \text{ dB}} \approx \sqrt{S_{\text{partition}}} \times \frac{180}{\pi} \quad (2)$$

For example, the truncated icosahedron would have partitions representing a beamwidth of $\theta_{3 \text{ dB}} \approx 35.90$ degrees. This beamwidth is quite large. To further reduce it, we begin with the truncated icosahedron. The truncated icosahedron can be transformed back to a polyhedron with triangular faces by taking its dual. The dual operation consists of identifying the centroid of each face of a polyhedron and adding edges connecting each existing vertex to the centroid. The centroid is then projected to have unit distance from the origin, becoming a new vertex. Once the polyhedron has triangular faces again, it may be truncated as usual to produce a polyhedron with hexagonal and pentagonal faces once more. This may be repeated until the desired beamwidth is achieved. Algorithm 1 summarizes this process as an object-oriented program. The user inputs a number of iterations N_{iter} based on a desired beamwidth and the algorithm outputs a polyhedron object. The polyhedron object is initialized as a unit icosahedron and is truncated once before the loop begins. The polyhedron object has vertices, faces, and centers as members. It also has member functions `truncate()` and `dual()` that perform the respective operations on the vertices, faces, and centers of the object. The polyhedron object is dualized and truncated in a loop until the number of iterations equals the input.

Table 2 provides an overview of the characteristics of polyhedra resulting from various values of N_{iter} . The method allows for a wide range of beamwidth approximations, with finer resolution achievable at smaller beamwidths. Notably, the number of faces to search is significantly reduced compared to the points in a traditional raster scan. Interestingly, there are always 12 pentagonal faces, a consequence of Euler's characteristic, which precludes a sphere from being entirely tiled with hexagons [Euler, 1758]. Figure 2 shows an example of a polyhedron with triangular faces before and after truncation with $N_{\text{iter}} = 5$, showing both hexagonal and pentagonal faces.

Algorithm 1: generatePolyhedron

Input: N_{iter} **Output:** polyhedron

```
1 polyhedron  $\leftarrow$  icosahedron
2 polyhedron.truncate()
3 for  $i = 2:N_{\text{iter}}$  do
4   polyhedron.dual()
5   polyhedron.truncate()
6 end
```

Table 2: Polyhedron Partitioning: Face and Beamwidth Analysis

N_{iter}	Faces	Hexagons	Pentagons	Beamwidth (degrees)	Raster Points
1	32	20	12	35.90	66
2	122	110	12	18.39	200
3	272	260	12	12.32	450
4	482	470	12	9.25	780
5	752	740	12	7.41	1225
6	1082	1070	12	6.17	1770
7	1472	1460	12	5.29	2415
9	2432	2420	12	4.11	3872
12	4322	4310	12	3.09	6903
14	5882	5870	12	2.64	9453

2. Assigning Probabilities

Once a suitable search space is defined, it is necessary to apply probabilities to each partition in order to perform a maximum likelihood search. The first step involves setting the probability of impossible search cells to zero. These cells can typically be identified by elevation: hexagonal and pentagonal sectors that fall below a minimum elevation angle should have their probabilities set to zero. The user may set this mask angle based on preference, though most networks will have a prescribed minimum elevation angle below which satellites will not transmit. For example, SpaceX’s Starlink has a minimum elevation angle of 25 degrees for Ku-band operations, while Amazon Kuiper will require a minimum elevation angle of 35 degrees for their customer terminals [Space Exploration Holdings, 2020, Systems, 2019]. Adopting a network-based minimum elevation angle, when available, is advisable, as it can eliminate a significant portion of the search space. While a user-selected minimum elevation angle may also be effective, it risks creating a search space that is either too large or too small.

After reducing the total search space, an appropriate PDF is applied to the remaining search cells. One possibility would be to apply a uniform distribution to the search space, resulting in a completely random search. This would be a suboptimal strategy, as the probabilities would be poorly assigned. A significant improvement could be made by considering a common aspect of satellite communications: satellite networks prioritize transmitting from higher-elevation satellites as these signals have the highest received power. With this in mind, a more effective strategy might be to search higher elevation cells first, as they are more likely to contain a satellite. Since the CT might lack precise knowledge of the exact distribution relative to elevation, an assumed distribution (e.g., linear, exponential) would be necessary. Fig. 3 shows the truncated polyhedron from Fig. 2 with a minimum elevation angle of 35 degrees and a heat map applied according to uniform and linear elevation-based PDFs. The PDFs are mapped to a maximum-likelihood criterion, where 1 indicates the most likely cell and 0 the least likely cell. Clearly, the elevation-based PDF better approximates how communications satellites may transmit in practice. Thus, if the CT had no knowledge of its own network the best approach may be to apply a user-selected minimum elevation angle and use an elevation-based PDF. However, there are better strategies if recent knowledge of some network parameters is available.

When recent information about the state of the constellation is known to the CT, such as if the CT were recently powered off after connection to the network, a more advanced way of determining the PDF can be employed. This approach is crucial for ensuring the speed of the HOOC-EM algorithm. The technique goes as follows. Suppose the network consists of a constellation of satellites, each of which can be described by its position and velocity vectors $\mathbf{r}$ and $\mathbf{v}$, which may be propagated using a numerical integration technique. Consider also that the CT is placed at the origin of the spherical search space described previously and that each hexagonal partition is a cell to be searched in the initial beam sweeping procedure. By propagating the positions and velocities of the satellites within the constellation over a chosen duration, one can record the cell location of each satellite at each time epoch. Using a frequentist approach, the probability density of each cell of the search space can then be inferred. Let N be the total number of hexagonal and pentagonal cells in the search space, $p(i)$ be the probability that the i th




Figure 2: A polyhedron with triangular faces after $N_{\text{iter}} = 5$ iterations (left) and the same polyhedron after the `truncate()` operation is performed (right), showing both hexagonal and pentagonal faces.

cell is occupied by at least one satellite, and $n(i)$ be the number of epochs in which the i th cell has been occupied by at least one satellite. The approximation of $p(i)$ is then given by

$$p(i) \approx \frac{n(i)}{\sum_{j=0}^{N-1} n(j)}, \quad i \in \{0, 1, \dots, N-1\} \quad (3)$$

The complete set of N probabilities forms the PDF of the search space and will be denoted by $\mathbf{p} = [p(0), p(1), \dots, p(N-1)]^T$. Since the satellites in view and their respective azimuths and elevations vary with the CT's latitude, this estimation must be repeated for each latitude where satellites are visible, given that the location of the CT is unknown when it is initially powered on. The resulting PDF for a given latitude λ will be denoted by $\mathbf{p}_\lambda = [p_\lambda(0), p_\lambda(1), \dots, p_\lambda(N-1)]^T$.

If the position of the CT remains unknown after powering on, latitude should be marginalized out of the PDF. This can be done by assigning a PDF to the possible latitudes of the CT, potentially using population density data. Since the CT lacks knowledge of its location, it can most reasonably estimate its latitude by where it is most likely to be powered on: where there are more people. The PDF may be determined by dividing the population density at each latitude by the total population accounted for in the data. For example, Fig. 4 shows the resulting PDF L_λ when this procedure is applied to population density data from [Kummu and Varis, 2011]. Finally, latitude can be marginalized out by

$$p(i) = \sum_{\lambda=\lambda_{\min}}^{\lambda_{\max}} \sum_{i=0}^{N-1} L_\lambda p_\lambda(i) \quad (4)$$

Now that a PDF has been obtained, assigning probabilities to each cell in the search space according to likely satellite positions and independent of latitude, it is still not appropriate to use this PDF directly in a maximum likelihood search. While this PDF will adequately describe where satellites are likely to be located, it does not account for where they are likely to be transmitting. To address this, consider that satellites are more likely to transmit at higher elevations. Therefore, a more suitable PDF would be a joint PDF that considers both satellite position and elevation. To determine the joint pdf, the previously obtained PDF $\mathbf{p}$ and an elevation-based PDF like what is shown in Fig. 3, which will be denoted by $\mathbf{e} = [e(0), e(1), \dots, e(N-1)]^T$ are combined. Note that $\mathbf{e}$ will have its highest probability at the cell at the zenith angle relative to the CT and the probability will decrease as relative elevation angles of the cells decrease until the elevation angles of the cells drop below some minimum elevation angle $\text{el}_{\min}$. All cells below that elevation angle will have a probability of 0. Since $\mathbf{e}$ is already conditioned on the hexagonal cells of the search space, the i th element of the joint PDF $\mathbf{x}$ can be calculated by an element-wise multiplication and normalization as




Figure 3: A heat map of the truncated polyhedron from Fig. 2 after applying a minimum elevation angle of 35 degrees and a uniform (left) and elevation-based (right) probability density function. The probability density function is mapped to a maximum-likelihood search criterion, applying a search likelihood of 1 to the most probable cells and 0 to the least probable.

$$x(i) = \frac{p(i)e(i)}{\sum_{i=0}^{N-1} p(i)e(i)} \quad (5)$$

Further improvement of this joint PDF could be made if even more specific network information was known to the CT, but that is outside the scope of this paper as this information could vary significantly between networks.

The current iteration of the PDF x is now a very good representation of the true PDF of the search space. However, it still fails to account for satellite dynamics. Indeed, a maximum-likelihood search for the best beam pair will likely span multiple time epochs, yet the current PDF only captures a single moment in time. While (3) marginalizes over time to address the CT's lack of knowledge regarding the true time at power-on, this approach is insufficient for an extended search process.

Over the course of a maximum likelihood search, searched cells are set to 0 probability to prevent them from being searched again. Due to the nature of satellite dynamics, permanently assigning 0 probability to a searched cell for the remainder of the search can be erroneous. Over time, the probability that a previously searched cell will be reoccupied by a transmitting satellite increases as satellites transition from one cell to another.

To better model this dynamic behavior, a Markov model is employed. Since the exact time is unknown and higher-order terms in the satellite dynamics model are not accounted for, the transitions of satellites between cells may be considered stochastic from the perspective of the CT. This approach allows for a more accurate representation of the evolving probabilities throughout the search process, leading to improved search efficiency and accuracy. Fortunately, the transition probabilities can be determined during the same simulations used to compute p_λ , noting that transition probabilities will depend on latitude as well. Let M_λ be the Markov transition matrix at latitude λ and $M_{\lambda,i}(j)$ be the probability that a satellite transitions from the i th cell to the j th cell at latitude λ . The stochastic matrix M_λ would be assembled as

$$M_\lambda = \begin{bmatrix} M_{\lambda,0}(0) & M_{\lambda,0}(1) & \dots & M_{\lambda,0}(N-1) \\ M_{\lambda,1}(0) & M_{\lambda,1}(1) & \dots & M_{\lambda,1}(N-1) \\ \vdots & \vdots & \ddots & \vdots \\ M_{\lambda,N-1}(0) & M_{\lambda,N-1}(1) & \dots & M_{\lambda,N-1}(N-1) \end{bmatrix} \quad (6)$$

The elements of M_λ can be approximated through numerical simulations of the constellation in a similar way to p_λ in (3). Specifically, the matrix A_λ can be assembled with elements $A_{\lambda,i}(j)$, which are the number of times any satellite has transitioned from the i th cell to the j th cell at latitude λ . The matrix A_λ is assembled in the same way as in (6). Although A_λ is not needed for any mathematical operations in matrix form, its matrix representation is convenient for data storage. Then, each element of




Figure 4: PDF of the CT latitudes determined by data from [Kummu and Varis, 2011].

M_λ can be approximated as

$$M_{\lambda,i}(j) \approx \frac{A_{\lambda,i}(j)}{\sum_{k=0}^{N-1} A_{\lambda,i}(k)}, \quad i, j \in \{0, 1, \dots, N-1\} \quad (7)$$

Additionally, the corresponding matrix M_λ can be marginalized for latitude by approximating $M_{i,j}$ using $\mathbf{A}$, the marginalized version of $\mathbf{A}_\lambda$. The procedure is the same as in (4)

$$\mathbf{A} = \sum_{\lambda=\lambda_{\min}}^{\lambda_{\max}} L_\lambda \mathbf{A}_\lambda \quad (8)$$

The transition probabilities $M_i(j)$ can then be approximated as in (7). Interestingly, the extraction of the Markov transition matrix eliminates the need to calculate $\mathbf{p}$ directly. Instead, $\mathbf{p}$ can be derived from $\mathbf{M}$. Recognizing that $\mathbf{M}$ is a positive stochastic matrix, it follows that 1 is an eigenvalue of $\mathbf{M}$, with all other eigenvalues having absolute values less than 1. By the Perron-Frobenius theorem, $\mathbf{M}$ must admit a unique normalized steady state vector, which spans the 1-eigenspace [Robinson, 2002]. That is, the steady state solution of the difference equation

$$\boldsymbol{\pi}^\top = \boldsymbol{\pi}^\top \mathbf{M} \quad (9)$$

exists and serves as a valid PDF. This solution represents the steady state behavior of the Markov model, corresponding to the probabilities that the cells in the search space contain at least one satellite, which is the definition of $\mathbf{p}$.

Algorithm 2 summarizes the process of assigning probabilities to the search space in HOOC-EM, structured as an object oriented program. The program takes as input the polyhedron object generated from Algorithm 1, and a constellation object with `simulate()` and `count()` member functions. The `simulate()` function simulates the constellation for a duration T seconds with an update period of Δt seconds, recording the states of each satellite in the constellation at each time epoch. The `count()` function tracks the number of transitions between the cells of the polyhedron object according to a specified latitude λ .

Additional inputs include the simulation duration T , the update period Δt , and minimum and maximum latitudes $\lambda_{\min}$ and $\lambda_{\max}$, respectively. The algorithm outputs $\mathbf{p}$ and $\mathbf{M}$ for use in the subsequent search algorithm.

3. HOOC-EM Search

Once the search space has been defined and probabilities have been assigned, the search can begin. The search proceeds as a typical maximum likelihood search, where the algorithm will search the cell with the highest probability in the PDF and so on. The slight difference is in the inclusion of the Markov model, which will continually update the PDF as the search continues.

It is important to note that the Markov transition matrix must be relatively recent to enable HOOC-EM. In the event of a complete cold start, the CT would first perform an elevation-based search to initiate network entry. Afterwards, and for as long as the CT

Algorithm 2: assignProbabilities

Input: polyhedron, constellation, T , Δt , $\lambda_{\min}$, $\lambda_{\max}$ **Output:** p , M

```
1 constellation.simulate( $T$ ,  $\Delta t$ )
2  $\mathbf{A} \leftarrow \text{zeros}(N, N)$ 
3 for  $\lambda = \lambda_{\min} : \lambda_{\max}$  do
4    $\mathbf{A}_{\lambda} = \text{constellation.count}(\text{polyhedron}, \lambda)$ 
5    $\mathbf{A} += L_{\lambda} \mathbf{A}_{\lambda} /* \text{From (8)} \quad */$ 
6 end
7  $\mathbf{M} = \text{calcMarkovMatrix}(\mathbf{A}) /* \text{From (7)} \quad */$ 
8  $p = \text{calcPDF}(\mathbf{M}) /* \text{Solution to (9)} \quad */$ 
```

is connected to the network, the network will furnish ‘fresh’ Markov transition matrices with which the CT can use to perform HOOC-EM if it is powered off and then powered on in quick succession. How ‘fresh’ the Markov transition matrix needs to be still needs to be studied but the results presented in Section IV indicate that the Markov matrix is still effective after at least one orbital period.

Algorithm 3 provides a programatic explanation of the maximum likelihood search used in HOOC-EM. The algorithm accepts the following inputs: a `searchSpace` object, which can represent an entire polyhedron or a reduced version after applying a minimum elevation angle; $\mathbf{x}$, the PDF; and $\mathbf{M}$, the Markov transition matrix. Note that $\mathbf{x}$ and $\mathbf{M}$ both apply to the entire polyhedron, as opposed to just the search space. This is because the transitions that would occur at the edges of the search space would be more difficult to model. Besides, $\mathbf{x}$ will only have probabilities greater than 0 above the minimum elevation angle anyway due to the application of the elevation-based PDF earlier.

The algorithm begins by identifying the index of the cell with the highest probability. The probability of that cell is then set to zero. The `searchCell()` function determines the azimuth and elevation of the center of that cell, points the beam of the antenna array at that azimuth and elevation, reads samples from the ADC, and correlates them with a local replica of the synchronization signal and measures the resulting signal power. If the power exceeds a certain threshold, detection is returned as **True**, ending the loop, and azimuth and elevation are returned as ‘az’ and ‘el’, respectively. If the power is below the threshold, detection is returned as **False** and the PDF is updated using the Markov transition matrix. The process then continues until detection is **True**. The azimuth and elevation are returned for satellite tracking purposes later on.

Algorithm 3: maximumLikelihoodSearch

Input: searchSpace, $\mathbf{x}$, $\mathbf{M}$ **Output:** az, el

```
1 detection  $\leftarrow$  False
2 while !detection do
3    $i = \text{argmax}(\mathbf{x})$ 
4    $\mathbf{x}(i) = 0$ 
5   detection, az, el = searchCell(searchSpace,  $i$ )
6    $\mathbf{x} = \mathbf{xM}$ 
7 end
```

Finally, the HOOC-EM algorithm can be developed. Algorithm 4 outlines the program as an object oriented program with Algorithms 1, 2, and 3 featured as functions. Of note there are the functions `calcJointPDF()` and `applyMinimumElevationAngle()`. The function `calcJointPDF()` performs the operation shown in (5) by first generating the PDF e , taking into account the minimum elevation angle $\text{el}_{\min}$. The function `applyMinimumElevationAngle()` removes all cells with centers below the elevation angle $\text{el}_{\min}$ from the search space.

IV. SIMULATIONS AND RESULTS

HOOC-EM can now be tested using an exemplar communications constellation. This paper will use an approximation of the Amazon Kuiper-Ka constellation for testing. The exact parameters used in the constellation are taken from [Systems, 2019]. The constellation is a LEO constellation with 3 orbital shells at different altitudes and inclination angles. Table 3 shows the given parameters for each orbital shell. The exact parameters for simulation of the constellation are not given, but it is claimed that the constellation will have good coverage between 56 degrees north and south of the Equator. One way to achieve this would be to simulate each shell as a Walker Delta constellation. That is the approach taken in this paper.

Algorithm 4: HOOC-EM

Input: $N_{\text{iter}}, \text{constellation}, T, \Delta t, \lambda_{\min}, \lambda_{\max}, \text{el}_{\min}$ **Output:** az, el

```
1 polyhedron = generatePolyhedron( $N_{\text{iter}}$ ) /* Algorithm 1 */
2  $\mathbf{p}, \mathbf{M}$  = assignProbabilities(polyhedron, constellation,  $T, \Delta t, \lambda_{\min}, \lambda_{\max}$ ) /* Algorithm 2 */
3  $\mathbf{x}$  = calcJointPDF( $\mathbf{p}, \text{el}_{\min}$ ) /* From (5) */
4 searchSpace = applyMinimumElevationAngle(polyhedron,  $\text{el}_{\min}$ )
5  $\text{az}, \text{el}$  = maximumLikelihoodSearch(searchSpace,  $\mathbf{x}, \mathbf{M}$ ) /* Algorithm 3 */
```

Table 3: Simulated Constellation: Altitudes and Inclinations [Systems, 2019]

Altitude (km)	Inclination (deg)	Planes	Number of Satellites per Plane	Number of Satellites
590	33	28	28	784
610	42	36	36	1296
630	51.9	34	34	1156

Using the Walker Delta constellation, the COE of each satellite can be determined. Using the COE each satellite can be propagated forward in time using a numerical integration technique. For the purposes of assigning probabilities, the constellation will be propagated using the two-body model with no perturbations. The two-body model has an analytical solution and is periodic with one orbital period. After rotating to the Earth-centered, Earth-fixed (ECEF) coordinate system, the constellation will be exactly periodic after a number of sidereal days that are evenly divisible by the orbital period. Furthermore, emerging communications constellations often consist of multiple shells, like the Kuiper-Ka constellation. As a result, the constellation will only be exactly periodic after a number of sidereal days that are evenly divisible by the greatest common multiple of the three orbital periods. This, of course, would be prohibitively long to simulate. Therefore, this paper will only consider one orbital period for testing. That is, the resulting PDF and Markov transition will only be good to one orbital period of the constellation or approximately 80 minutes in the case of the Kuiper-Ka constellation. Let $\mathbf{r} = [x, y, z]^T$ be the position state vector of a satellite in the Earth-centered, inertial (ECI) frame, μ be the standard gravitational parameter for Earth, $\|\cdot\|_2$ be the standard Euclidean norm, and $f(t, \mathbf{r}, \dot{\mathbf{r}})$ be unmodeled, time-dependent perturbations. Then the two-body equations of motion are

$$\ddot{\mathbf{r}} = -\frac{\mu}{\|\mathbf{r}\|_2^3} \mathbf{r} + f(t, \mathbf{r}, \dot{\mathbf{r}}) \quad (10)$$

Now (10) may be propagated using a numerical integration technique. Fig. 5 shows the constellation described in Table 3 propagated over one orbital period. In addition, the average number of satellites above a 35 degree minimum elevation angle is shown as a function of latitude, showing strong correspondence to the same chart shown in [Systems, 2019].

Using the simulated constellation in Fig. 5, the positions of the satellites are rotated from the ECI frame to the ECEF frame and the probabilities are assigned as in Subsection III.2 to the polyhedra described in Table 2 with a minimum elevation angle of 35 degrees and the joint PDF is determined using an exponential, elevation-based PDF. Let beamwidth in degrees be β , then Fig. 6 shows the resulting probability density functions mapped to each search space delineated by their beamwidth according to a maximum likelihood criterion assigning the highest probability cell a 1 and decreasing as probability of each cell decreases.

1. Simulations

Once the probabilities are assigned, Monte Carlo simulations are carried out to verify the efficacy of HOOC-EM. The Monte Carlo simulations are performed as follows. The constellation is simulated using the two-body model over one orbital period. However, for the purposes of the Monte Carlo simulation the J_2 perturbation will be modeled and included in the simulation of the constellation. This is to serve as a stand-in for unmodeled perturbations when determining the Markov transition matrix $\mathbf{M}$ at the control segment. The most significant perturbing accelerations for a LEO satellite are due to Earth's non-uniform gravity. The most severe of which are the zonal terms denoted as J_2 [Curtis, 2009]. Let R_e be the mean radius of the Earth and $J_2 = 1.08262668355 \times 10^{-3}$ be the dimensionless zonal coefficient. The resulting equations of motion of each axis of a LEO satellite in the ECI frame are [Morales et al., 2019]




Figure 5: The simulated LEO constellation shown above Earth (left) and the average number of satellites in view (above a 35 degree elevation angle) at different latitudes (right) determined at increments of 5 degrees with each increment denoted by a triangle.

$$\begin{aligned}
 \ddot{x} &= \frac{-\mu x}{||\mathbf{r}||_2^3} \left[1 + J_2 \frac{3}{2} \left(\frac{R_e}{||\mathbf{r}||_2} \right)^2 \left(1 - 5 \frac{z^2}{||\mathbf{r}||_2^2} \right) \right] \\
 \ddot{y} &= \frac{-\mu y}{||\mathbf{r}||_2^3} \left[1 + J_2 \frac{3}{2} \left(\frac{R_e}{||\mathbf{r}||_2} \right)^2 \left(1 - 5 \frac{z^2}{||\mathbf{r}||_2^2} \right) \right] \\
 \ddot{z} &= \frac{-\mu z}{||\mathbf{r}||_2^3} \left[1 + J_2 \frac{3}{2} \left(\frac{R_e}{||\mathbf{r}||_2} \right)^2 \left(3 - 5 \frac{z^2}{||\mathbf{r}||_2^2} \right) \right]
 \end{aligned} \tag{11}$$

The equations in (11) are integrated using a numerical integration technique with the Walker Delta initial conditions for each satellite. The positions of each satellite in the constellation are then saved. At the beginning of each Monte Carlo simulation a random time within the orbital period is selected and the constellation is propagated from that time. The position of the CT is randomly selected; the latitude of the CT is randomly selected according to the PDF given in Fig. 4, the longitude is selected from a uniform distribution from -180 to 180 degrees, and the altitude is chosen to be 800 m, which is the average height of land above sea level. Among the visible satellites, one is selected to be transmitting at the beginning of the simulation according to the elevation-based PDF in Fig. 3. The selected satellite will remain transmitting for the fixed assignment interval T_{FAI} seconds. Afterwards a new satellite will be selected using the same criterion and will transmit for T_{FAI} seconds and so on. During this time searches will be going on with the following methods: (1) a raster scan of azimuths and elevations, beginning with zenith (2) a maximum likelihood search with the uniform PDF from Fig. 3, (3) a maximum likelihood search with the elevation-based PDF from Fig. 3, and (4) HOOC-EM using the PDFs from 6 that are propagated using the Markov models derived from earlier simulations of the constellation. Note that each method is performed at the beamwidths shown in Fig. 6 and the uniform and elevation-based PDFs are equivalently mapped to the corresponding polyhedra. This proceeds until a detection or acquisition is made.

The acquisition criterion is that the transmitting satellite is within the cell that the antenna of the CT is pointing at during the period T_{FAI} . While no explicit signal power criterion is used here it can be argued that the network may design their signal to function with a given beamwidth. For example, it was shown in [Komodromos et al., 2023] that an OFDM signal may be detected at SNRs as low as -43 dB when the synchronization sequence encompassed the entire Starlink frame. While this would be unreasonable in practice, as no data would be broadcast, it shows how detection at a certain antenna gain or beamwidth could be considered a signal design parameter. Furthermore, the present approach could be argued to be conservative. If a detector based on pure SNR were used a satellite may be detected in adjacent hexagonal cells, depending on the position of the satellite.




Figure 6: The empirical satellite-and-elevation-based probability density functions, marginalized for latitude, applied to the hexagonal and pentagonal sectors above a 35 degree elevation mask for beamwidths from 2.65 to 18.39 degrees.

2. Results

The simulations were performed with $T_{\text{FAI}} = 5$ and 15 seconds. A T_{FAI} of 15 seconds aligns with observations of the Starlink network in [Qin et al., 2023], while a shorter T_{FAI} of 5 seconds could be employed by a future fused LEO GNSS as shorter assignment intervals would allow for greater geometric diversity and reduce the time to first fix. Fig. 7 shows the results of a 500 run Monte Carlo simulation at $T_{\text{FAI}} = 5$ seconds.

On the left, the average time to detection in seconds is plotted against beamwidth in degrees. It clearly shows that the raster scan is the least effective approach, even underperforming compared to the uniform search with hexagonal tiling, which benefits from fewer search points. For wider beamwidths, the elevation-based and HOOC-EM approaches perform similarly, likely due to fewer search points and the reduced likelihood of satellite transitions in larger hexagonal cells. However, as beamwidths narrow—where satellite transitions are more frequent and the number of search points increases—HOOC-EM demonstrates superior performance, which is crucial in high-interference scenarios requiring narrow beams.

On the right, the cumulative density functions (CDFs) for a beamwidth of 2.64 degrees are displayed, further highlighting the inferiority of the uniform and raster scan approaches. The elevation-based and HOOC-EM searches perform similarly for the first 50% of the searches, but beyond this, the elevation-based approach shows multiple plateaus. At this assignment interval and beamwidth, over 90% of the HOOC-EM searches finish in under 40 seconds, while the elevation-based search has much longer tails, with over 10% of the searches taking over a minute to perform. Fig. 8 shows the results when $T_{\text{FAI}} = 15$ seconds. The results are very similar to those shown in Fig. 7, with a general reduction in average detection time. This improvement occurs because the search is more likely to finish before the transmitting satellite switches beams. Overall the trends remain consistent: the raster scan and uniform searches are the slowest, while the elevation-based and HOOC-EM searches are the fastest, with the HOOC-EM search edging out the elevation-based search at narrow beamwidths.

On the right, the CDFs for a beamwidth of 2.64 degrees are shown. The elevation-based and HOOC-EM searches perform similarly up to just over 60% of the searches. At this assignment interval and beamwidth, over 90% of the HOOC-EM searches take under 30 seconds to complete, compared to over 10% of the elevation-based searches taking over 50 seconds to complete. If the user has some information about the position of the CT, performance can be improved significantly improved. Fig. 9 shows the results at a known position for elevation-based and HOOC-EM searches at $T_{\text{FAI}} = 15$ seconds. The Markov matrix M was determined by a simulation of the constellation over one orbital period with the CT at the known position. With knowledge of





Figure 7: The average time to detection of the various beam sweeping methods (left) and their cumulative density functions (right) at a beamwidth of 2.64 degrees with $T_{FAI} = 5$ seconds.

the position of the CT, the HOOC-EM search is convincingly better than the elevation-based search in terms of average time to detection. On the right, the CDFs for both searches at a beamwidth of 2.64 degrees are shown. Nearly 99% of the HOOC-EM searches complete in just under 20 seconds when the location is known.

While it may be unlikely to know the exact position of the CT without GNSS, Fig. 9 highlights how performance can improve with some degree of location specificity. For example, the specific position of the CT may be unknown to the user, but knowing the general area such as the city, state, or country could still provide a notable performance boost.





Figure 8: The average time to detection of the various beam sweeping methods (left) and their cumulative density functions (right) at a beamwidth of 2.64 degrees with $T_{FAI} = 15$ seconds.





Figure 9: The average time to detection of the various beam sweeping methods (left) and their cumulative density functions (right) at a beamwidth of 2.64 degrees with $T_{FAI} = 15$ seconds and a known position.

V. CONCLUSIONS

This paper presented HOOC-EM, a novel method for fast CT beam sweeping in LEO mega-constellation networks. HOOC-EM consists of multiple algorithms: (1) partitioning the celestial sphere into hexagonal and pentagonal cells, (2) assigning probabilities to these cells based on satellite geometry, and (3) performing a maximum likelihood search of the partitioned celestial sphere that is dynamically updated according to satellite movement and the fixed assignment interval of the network. Simulations were conducted on an exemplar communications constellation modeled after the Amazon Kuiper-Ka constellation. HOOC-EM was evaluated against an elevation-first raster scan search, a uniform maximum-likelihood search, and an elevation-based maximum likelihood search.

The results showed that HOOC-EM outperformed the other methods at narrow beamwidths while performing comparably to the elevation-based search at wider beamwidths. Furthermore, the simulations also indicated that providing the CT with some information on its position may improve HOOC-EM search speeds. Analyses of the CDFs of the various search methods show that the HOOC-EM search has much shorter tails than the elevation-based search.

ACKNOWLEDGEMENTS

This work was supported by Amazon's Project Kuiper as an affiliate of the 6G@UT center within the Wireless Networking and Communications Group at The University of Texas at Austin.

REFERENCES

- [3GPP, 2017] 3GPP (2017). Study on NR to support non-terrestrial networks. Tr 38.811, 3rd Generation Partnership Project (3GPP). Version 15.0.0.
- [3GPP, 2023] 3GPP (2023). Solutions for NR to support non-terrestrial networks. Tr 38.821, 3rd Generation Partnership Project (3GPP). Version 16.0.0.
- [Andrews et al., 2014] Andrews, J. G., Buzzi, S., Choi, W., Hanly, S. V., Lozano, A., Soong, A. C., and Zhang, J. C. (2014). What will 5G be? *IEEE Journal on selected areas in communications*, 32(6):1065–1082.
- [Artiga and Vázquez, 2023] Artiga, X. and Vázquez, M. Á. (2023). Evaluation of new radio beam management framework for LEO satellites. In *2023 IEEE Conference on Standards for Communications and Networking (CSCN)*, pages 241–246. IEEE.
- [Baykas et al., 2011] Baykas, T., Sum, C.-S., Lan, Z., Wang, J., Rahman, M. A., Harada, H., and Kato, S. (2011). IEEE 802.15.3c: The first IEEE wireless standard for data rates over 1 Gb/s. *IEEE Communications Magazine*, 49(7):114–121.
- [Blázquez-García et al., 2023] Blázquez-García, R., Cristallini, D., Ummenhofer, M., Seidel, V., Heckenbach, J., and O'Hagan, D. (2023). Experimental comparison of Starlink and OneWeb signals for passive radar. In *2023 IEEE Radar Conference (RadarConf23)*, pages 1–6.
- [Chan, 1980] Chan, F. (1980). A quadrilateralized spherical cube earth data base. In NASA. *Goddard Space Flight Center Fifth Ann. Flight Mech.(Estimation Theory Symp.)*.
- [Curtis, 2009] Curtis, H. D. (2009). *Orbital Mechanics for Engineering Students*. Elsevier, Burlington, MA, 2nd edition.
- [Dimitrijević et al., 2016] Dimitrijević, A. M., Lambers, M., and Rančić, D. (2016). Comparison of spherical cube map projections used in planet-sized terrain rendering. *Facta Universitatis, Series: Mathematics and Informatics*, 31(2):259–297.
- [Euler, 1758] Euler, L. (1758). *Elementa doctrinae solidorum. Novi commentarii academiae scientiarum Petropolitanae*, pages 109–140.
- [Ghasempour et al., 2017] Ghasempour, Y., Da Silva, C. R., Cordeiro, C., and Knightly, E. W. (2017). IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine*, 55(12):186–192.
- [Heng et al., 2021] Heng, Y., Andrews, J. G., Mo, J., Va, V., Ali, A., Ng, B. L., and Zhang, J. C. (2021). Six key challenges for beam management in 5.5 G and 6G systems. *IEEE Communications Magazine*, 59(7):74–79.
- [Humphreys et al., 2023] Humphreys, T. E., Iannucci, P. A., Komodromos, Z. M., and Graff, A. M. (2023). Signal structure of the Starlink Ku-band downlink. *IEEE Transactions on Aerospace and Electronic Systems*, pages 1–16.
- [Iannucci and Humphreys, 2022] Iannucci, P. A. and Humphreys, T. E. (2022). Fused low-earth-orbit GNSS. *IEEE Transactions on Aerospace and Electronic Systems*, pages 1–1.
- [Kim et al., 2020] Kim, J., Yun, M. Y., You, D., and Lee, M.-S. (2020). Beam management for 5G satellite systems based on NR. In *2020 International Conference on Information Networking (ICOIN)*, pages 32–34. IEEE.

- [Komodromos et al., 2023] Komodromos, Z. M., Qin, W., and Humphreys, T. E. (2023). Signal simulator for Starlink Ku-Band downlink. In *Proceedings of the ION GNSS+ Meeting*, pages 2798–2812.
- [Kummu and Varis, 2011] Kummu, M. and Varis, O. (2011). The world by latitudes: A global analysis of human population, development level and environment across the north–south axis over the past half century. *Applied geography*, 31(2):495–507.
- [Kutkov, 2023] Kutkov, O. (2023). Connecting external gps antenna to the starlink terminal. <https://olegkutkov.me/2023/11/07/connecting-external-gps-antenna-to-the-starlink-terminal/>. Accessed: 2023-03-01.
- [Maine et al., 1995] Maine, K., Devieux, C., and Swan, P. (1995). Overview of Iridium satellite network. In *WESCON'95 Conference*, page 483. IEEE.
- [Maral et al., 2020] Maral, G., Bousquet, M., and Sun, Z. (2020). *Satellite communications systems: systems, techniques and technology*. John Wiley & Sons.
- [McCollum, 2001] McCollum, J. M. (2001). Honeycombing the icosahedron and icosahedroning the sphere¹. In *Proceedings of the Second Annual Forest and Inventory Symposium: Salt Lake City, Utah, October 17-18, 2000*, volume 47, page 25. USDA Forest Service, Southern Research Station.
- [Morales et al., 2019] Morales, J. J., Khalife, J., Cruz, U. S., and Kassas, Z. M. (2019). Orbit modeling for simultaneous tracking and navigation using LEO satellite signals. In *Proceedings of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*, pages 2090–2099.
- [Nitsche et al., 2014] Nitsche, T., Cordeiro, C., Flores, A. B., Knightly, E. W., Perahia, E., and Widmer, J. C. (2014). IEEE 802.11 ad: directional 60 GHz communication for multi-gigabit-per-second Wi-Fi. *IEEE Communications Magazine*, 52(12):132–141.
- [Qin et al., 2023] Qin, W., Komodromos, Z. M., and Humphreys, T. E. (2023). An agile, portable antenna system for LEO megaconstellation-based PNT. In *Proceedings of the ION GNSS+ Meeting*.
- [Rappaport et al., 2014] Rappaport, T. S., Heath Jr, R. W., Daniels, R. C., and Murdock, J. N. (2014). *Millimeter wave wireless communications*. Pearson Education.
- [Rinaldi et al., 2020] Rinaldi, F., Määttänen, H.-L., Torsner, J., Pizzi, S., Andreev, S., Iera, A., Koucheryavy, Y., and Araniti, G. (2020). Broadcasting services over 5G NR enabled multi-beam non-terrestrial networks. *IEEE Transactions on Broadcasting*, 67(1):33–45.
- [Robinson, 2002] Robinson, R. C. (2002). Perron frobenius theorem. *mij*, 10:0.
- [Space Exploration Holdings, 2020] Space Exploration Holdings (2020). SpaceX Gen2 NGSO Satellite System, Attachment Waiver Requests. https://licensing.fcc.gov/myibfs/download.do?attachment_key=2378667. SAT-LOA-20200526-00055.
- [Systems, 2019] Systems, K. (2019). Application for fixed satellite service by Kuiper Systems LLC. https://licensing.fcc.gov/myibfs/download.do?attachment_key=1773885. SAT-LOA-20190704-00057.
- [Tegmark, 1996] Tegmark, M. (1996). An icosahedron-based method for pixelizing the celestial sphere. *The Astrophysical Journal*, 470(2):L81.
- [Wei et al., 2023] Wei, W., Zheng, Z., Wang, Y., Su, Y., Wang, X., Liu, L., Wang, B., Sun, C., and Wang, D. (2023). Random access control and beam management scheme for the 5G NR based beam hopping LEO satellite communication systems. *Mobile Networks and Applications*, pages 1–12.
- [Xue et al., 2024] Xue, Q., Ji, C., Ma, S., Guo, J., Xu, Y., Chen, Q., and Zhang, W. (2024). A survey of beam management for mmwave and THz communications towards 6G. *IEEE Communications Surveys & Tutorials*.