Copyright
by
Zachary L. Clements
2026

The Dissertation Committee for Zachary L. Clements
certifies that this is the approved version of the following dissertation:

GNSS Interference Detection and Geolocation

Committee:

Todd E. Humphreys, Supervisor

Brandon A. Jones

Jingyi Chen

Srinivas Bettadpur

Patrick B. Ellis

GNSS Interference Detection and Geolocation

by

Zachary L. Clements

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

May 2026

To the relentless pursuit of ground truth.

Acknowledgments

I would like to express my sincere thanks and gratitude to a number of people who helped me complete my Ph.D. I would first like to thank my parents Suli and Joe, my brothers Jacob and Sean, and Daisy for their endless support and encouragement.

I would like to thank my advisor Prof. Todd Humphreys for his supervision over the last six years. His technical prowess and guidance have provided the foundation upon which this dissertation was built. I am grateful for the invaluable lessons I have learned about navigating research and life. His keen eye for detail and unwavering commitment to excellence have made me a better researcher and person. There certainly was never a dull moment and no shortage of interesting projects.

I would like to thank Dr. Patrick Ellis for providing me the opportunity to intern at Spire Global, an experience which helped shape my present career goals. He has been an incredibly reliable mentor who has guided me through both technical and personal endeavors. I would also like to thank him for serving on my Ph.D. committee.

Very many thanks to the rest of my Ph.D. committee members Prof. Brandon Jones, Prof. Jingyi Chen, and Prof. Srinivas Bettadpur for reviewing my Ph.D. research and serving on my dissertation committee.

During my Ph.D. I have met countless individuals who have made this a much more enjoyable process. I would like to thank my current labmates Robert, Zack, Tucker, Wenkai, Alperen, Sam, Jeremy, and Dimitry, and former labmates Andrew, Hailey, and James, for their camaraderie over the last six years—I could not have asked for better labmates. I would also like to thank all alumni of the UT Radionavigation Lab for laying the foundation for all ongoing research efforts in the lab.

I would like to thank everyone that I went through this program with, especially Bryn, Sofia, Andrew, Agrim, Ian, and everyone on the third floor of ASE. I would like to thank Sung Gyung for her constant encouragement and humor. I would also like to thank everyone in the Institute of Navigation community for their support.

I would like to acknowledge and thank the U.S. Department of Transportation for their generous funding of the work in this dissertation under the CARMEN and CARMEN+ University Transportation Centers. I would also like to thank Spire Global for the internship opportunity and our collaboration. Finally, I would like to also thank all others who have sent us data, which has lead to many valuable insights.

GNSS Interference Detection and Geolocation

Zachary L. Clements, Ph.D.
The University of Texas at Austin, 2026

Supervisor: Todd E. Humphreys

This dissertation develops the theory and provides experimental verification of GNSS interference detection and localization algorithms with data from terrestrial- and space-based receivers. The first study offers an experimental demonstration of single-satellite geolocation of broadcast GNSS spoofers and an analysis of various error sources. The second study explores dual-satellite geolocation of GNSS interference from Low Earth Orbit (LEO) and presents an analysis of interference signals captured in the GNSS frequency bands by two time-synchronized LEO receivers over the Eastern Mediterranean, with the emitters being geolocated. The third study investigates a transient space-based GNSS interference source that has caused continental-wide GNSS interference and presents detection and estimation algorithms to identify the source. Together, these contributions strengthen GNSS security by providing the theory and experimental verification of GNSS interference detection and geolocation algorithms across several platforms.

Table of Contents

Acknowledgments **v**

Abstract **vii**

List of Tables **xi**

List of Figures **xii**

Chapter 1. Introduction **1**

1.1 The GNSS Interference Problem 1
1.2 GNSS Interference Defenses and PNT Alternatives 5
1.3 Protecting the GNSS Spectrum 7
1.4 Thesis Statement and Expected Contributions 10
1.5 Organization 12
1.6 List of Publications 13
1.6.1 Journal Publications 13
1.6.2 Conference Publications 14
1.6.3 Other Major Publications 15

Chapter 2. Single-Satellite-Based Geolocation of Broadcast GNSS Spoofers from Low Earth Orbit **16**

2.1 Introduction 17
2.2 Signal Models 22
2.2.1 GNSS Spoofing Signals 22
2.2.2 Received Doppler Model 24
2.3 Conceptual Overview of Broadcast GNSS Spoofer Geolocation 26
2.4 Spoofer Geolocation with γ 30
2.4.1 Measurement Model 30
2.4.2 Effects of Estimated γ 33

2.4.3 Range-rate Nonlinear Least-Squares 35
2.5 Experimental Results 37
2.5.1 Experimental Design 38
2.5.2 Experimental Spoofer Geolocation with γ 39
2.5.3 Experimental Spoofer Geolocation with GNSS Observables 42
2.6 Spoofer Clock Instability Error Analysis 45
2.7 Controlling Spoofing Detection While Degrading Geolocation Accuracy 50
2.7.1 Optimal Spoofing Detection via Clock Drift Monitoring 50
2.7.2 Expression for Geolocation Error 54
2.7.3 Jointly Optimized Spoofer Clock Drift Selection 56
2.8 Conclusion 59

Chapter 3. Dual-Satellite Geolocation of Terrestrial GNSS Interference from Low Earth Orbit 60
3.1 Introduction 61
3.2 Measurement Model 65
3.2.1 TDOA and FDOA Measurement Model 66
3.2.2 The Generalized Cross-Correlation Function 69
3.3 Emitter Geolocation 72
3.3.1 Two-step Geolocation 72
3.3.2 Direct Geolocation 75
3.3.3 Connection between the CAF Domain and Position Domain 78
3.4 A Real-world Capture 84
3.4.1 Capture Platform 84
3.4.2 Spectrum Analysis 86
3.5 Emitter Geolocation 87
3.5.1 Two-Step Geolocation 88
3.5.2 Direct Geolocation 91
3.6 Conclusion 96

Chapter 4. Transient Space-Based GNSS Interference Detection and Estimation **97**
4.1 Introduction 98
4.2 Measurement Models and Detection 101
4.2.1 Measurement Models 102
4.2.2 Detection 104
4.3 Interference Properties 109
4.3.1 Temporal Patterns 110
4.3.2 Spatial Patterns 112
4.3.3 Spectral Properties 114
4.3.4 Comparison with Solar Radio Burst 117
4.4 Elevation-Mask-Based Interference Source Identification 118
4.5 Multi-Hypothesis CNR-Based Interference Source Identification 121
4.5.1 Generalized Likelihood Ratio Test 121
4.5.2 Simulation Study 122
4.5.2.1 Simulation 126
4.5.2.2 Results and Discussion 127
4.5.3 GNSS Satellite Identification with Real CNR Data . . . 129
4.5.4 Implications for CINR-Based Interference Source Identification 131
4.6 TDOA-Based Interference Source Identification 133
4.6.1 TDOA Measurement Models 133
4.6.2 TDOA-Based Association Framework 135
4.6.3 A Wideband Raw-Signal Capture 137
4.6.4 TDOA Measurement and Method Validation 139
4.6.5 Satellite Association 141
4.6.6 Further Analysis and Discussion 143
4.7 Conclusion 145

Chapter 5. Conclusions **147**

Bibliography **150**

List of Tables

2.1 Final geolocation error using the observed Doppler time history of each spoofed PRN. 44

2.2 Theoretical marginal contribution of transmitter clock instability to the 95% horizontal error ellipse for the capture scenario specified in Section 2.5 in the absence of all other error sources. 46

4.1 The number of days per year that saw at least one wide-area transient GNSS interference event on the GPS L1 frequency with at least one station suffering a drop of 5 dB or greater. . 111

4.2 The minimum altitude and number of satellites that satisfy various elevation masks for the event on day 160 of 2021. 120

List of Figures

2.1 Shown here are the Doppler components in single-satellite spoofer geolocation. The Doppler components corresponding to (2.5) are shown on the left. The Doppler components for each spoofing signal corresponding to (2.4) are shown in red to the right. 25

2.2 The Pearson correlation coefficient between sequential estimation errors $w_{\mathrm{a}}[i]$ 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 between estimation errors decreases. 35

2.3 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. 39

2.4 Left: UT RNL's GRID receiver display when processing the spoofing signals. Right: 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. 40

2.5 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. 41

2.6 Left: Final spoofer position estimate (white) based on $\hat{\gamma}(t)$. Shown in red is the true spoofer location. The error of the final estimate is 68 m. The true emitter is contained within the 95% horizontal error ellipse, derived from (2.27), which has a semi-major axis of 6.7 km. Right: Post-fit range-rate residuals of $\hat{\gamma}(t)$ time history with respect to the estimated spoofer position (top) and true spoofer position (bottom). The residuals with respect to the estimated position are unbiased and have a standard deviation of 0.12 m/s. 42

2.7 Left: Geolocation using the observed Doppler time history of each spoofed PRN. Each individual spoofer position estimate is biased due to the unmodeled frequency component. Right: Range-rate residuals with respect to the estimated spoofer position (top). Range-rate residuals with respect to the true spoofer position (Bottom). 44

2.8 Left: Monte Carlo sample RMSE as a function of h_{-2} with the capture geometry specified in Section 2.5, for an estimator applying the correct (R) and incorrect (R_{a}) measurement covariance. Right: Percentage increase in sample RMSE when R_{a} is applied rather than R. 48

2.9 Left: Monte Carlo containment percentage when the theoretical 95% error ellipse has been calculated with R_{a} alone [setting $R_{\mathrm{b}} = 0$ in (2.19)], for various values of the underlying parameter σ_{a}. Right: Area of the corresponding theoretical 95% error ellipse as a function of σ_{a}. The horizontal line shows the area of the 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. . 48

2.10 The attack trajectory $\mathrm{v}[i]^* - \mathrm{v}^*[1]$ that maximizes e_{h} for the LEO-based receiver geometry shown in Fig. 2.3. 58

2.11 Worst-case geolocation error for a spoofer that optimizes $c\boldsymbol{\delta \dot{t}}_{\bar{\mathrm{r}}}$ for receivers performing 1-Hz spoofing detection tests with σ_{m} = 0.05 m/s and for the LEO-based receiver geometry shown in Fig. 2.3. The geolocation error is shown over a range of $\bar{P}_{\mathrm{D}}$ for two representative victim receiver clock quality levels and three representative values of P_{F}. 59

3.1 On the left is an example CAF for a single wideband emitter without any cyclostationary properties. The $\Delta\tau$ and Δf that correspond to the maximum magnitude are the maximum likelihood estimates. On the right is the CAF of a single wideband emitter with cyclostationary properties. There is an ambiguity because there are peaks with the same magnitude. 71

3.2 Example position-domain correlation grid from direct geolocation for a single wideband emitter without any cyclostationary properties. The maximum magnitude corresponds to the MLE of emitter's position. 77

3.3 Shown here is an example mapping of $|S^{\mathrm{CAF}}|$ and $|S^{\mathrm{PD}}|$. The position-domain correlation is effectively a homeomorphic mapping of the CAF domain. The blue dots represent the nearest sample and frequency bin pair. 79

3.4 Shown here are $|S^{\mathrm{CAF}}|$ and $|S^{\mathrm{PD}}|$ from the simulation across three distinct observation intervals for a cyclostationary chirp jammer. The red markings denote the geometric manifold mapping between the CAF and PD, while the white diamonds indicate the true $\Delta\tau$ and Δf and the true position. 81

3.5 Position-domain correlation after noncoherently accumulating all 58 snapshots of the simulation with the 3D view (left) and from above (right). Likelihood accumulates at the true source location, whereas the shifting ambiguities are suppressed, demonstrating that direct geolocation provides a powerful mechanism for geolocating cyclostationary interference. 82

3.6 Simulation setup with multiple interference sources. 84

3.7 Position domain correlation at the first snapshot (top) and the final snapshot (bottom). 84

3.8 Position domain correlation after noncoherently accumulating all 58 snapshots with the 3D view (left) and from above (right). The four distinct sharp peaks represent the MLE of each GNSS interference source's position. 84

3.9 Ground tracks of a pair of LEO-based receivers during the two 60-second captures over the Eastern Mediterranean. Both satellites are moving from north to south. 85

3.10 The average power spectral densities of the data during both captures at both frequencies. 87

3.11 Spectograms centered at the GPS L1 (left) and L2 (right) frequencies from data captured in April 2022. 87

3.12 Shown on the left are TDOA (top) and FDOA (bottom) measurements corresponding to the measurements from the interference source located at the Khmeimim Air Base. Shown on the right are the TDOA measurement residuals (top) and the FDOA measurement residuals (bottom). 88

3.13 The final two-step T/FDOA geolocation estimate converges to Khmeimim Air Base on the coast of Syria, which has been previously discovered to have transmitted GPS jamming signals [100]. The semi-major and semi-minor axes of the 95% ellipse are 632 and 57 meters, respectively when a conservative $\sigma_{\Delta\tau} = 50$ m and $\sigma_{\Delta f} = 5$ m/s are assumed. 89

3.14 Cost surface for the T/FDOA measurements. The global minimum corresponds to the Khmeimim Air Base. 90

3.15 Contours of the direct geolocation position-domain correlation grids over Khmeimim Air Base for 20-ms (top), 100-ms (middle), and 1-second (bottom) integration times during the beginning (left) and end(right) of the 60-second capture. The impact of integration time and geometry is apparent. 92

3.16 Direct geolocation on the chirp jammer in the beginning (left) and end (middle) of the capture. A 100 ms accumulation was taken once every 5 seconds. The eleven position-domain correlations were noncoherently combined (right). Noncoherent integration is a powerful tool that suppresses the spurious false peaks from waveform structure and cyclostationarity. 93

3.17 Direct geolocation over Syria with two different views of the position domain. A 100-ms accumulation was taken once per second, over 58 seconds. Shown here are 58 position-domain correlations noncoherently combined. Each peak corresponds to an emitter position estimate. 94

3.18 Final emitter position estimates on L1 and L2 across both simultaneous captures from direct geolocation. All emitters were assumed to be stationary and constrained to the surface of the Earth. 95

3.19 Jammers found in Ukraine (left) and Israel (right). 96

4.1 The distributions of Λ_i under H_0 (no interference present) for stations METG ($i = 1$) and MATE ($i = 2$). 107

4.2 Reported CNR (left) and the detection statistic Λ_i (right) for IGS stations METG ($i = 1$, Finland), MATE ($i = 2$, Italy), and THU2 ($i = 3$, Greenland) over a 15-minute interval on day 160 of year 2021. The dashed red line is the detection threshold with a 10^{-4} probability of false alarm. 108

4.3 The number of stations that detected interference during day 160 of year 2021. The expanded view on the right shows a lower-power event detected by 21 stations, followed by a higher-power event detected by 58 stations. 109

4.4 Heat map of the test statistic (in dB) at triggering stations for the lower-power event (left) and the higher-power event (right). The drop in GPS L1 C/A CNR during the more powerful interference event was as large as 6 dB, centered near the Baltic region. 109

4.5 Distribution of the day of the week and hour of day (with respect to UTC) during which interference events with at least one station suffering a drop of 5 dB or greater occurred. Clearly, the high-power interference events typically occur during business days and business hours. 111

4.6 The number of stations that detected interference during three different days. Interference events can occur several times a day. The left column shows the entire day, while the right column is an expanded sub-interval. 112

4.7 Detection statistic heat maps, in dB, for day 146 of year 2021 and day 014 of year 2025. GNSS receivers in the Baltic region are impacted most, which is generally representative of nearly all of the interference events. 113

4.8 Uncalibrated average PSD near the GPS L1 band during nominal operation (thick black line) and 1-Hz PSD estimates during 48 transient interference events, as recorded in Gdynia, Poland. 115

4.9 Power spectrum derived from raw wideband samples captured in Amsterdam, Netherlands, during an interference event on February 11, 2026. An initial burst of interference at 1577.5 MHz was followed by a burst at 1558.5 MHz. 116

4.10 Left: CNR time history of tracked GPS L1, L2, and L5 signals produced by IGS station SUTM in South Africa during the solar radio burst on November 11, 2025. Right: Heat map of maximum CNR degradation at GPS L2 during the solar radio burst. 118

4.11 Left: The position of all tracked objects during the high-power interference burst on day 160 year 2021. Right: The position of all satellites that satisfy a 0° elevation mask, excluding debris and rocket bodies. The feasible region in which the interference source could have been positioned is interior to the red surface. For reference, the colored spherical shell corresponds to medium Earth orbit (20,000 km altitude). 121

4.12 Simulation study setup. Left: The geometry of the true source GNSS satellite over 16 epochs at 15-minute intervals, with terrestrial reference stations shown as black dots. Right: Transmitter (top) and receiver (bottom) gain patterns. The receiver gain patterns are from a model provided by the manufacturer and verified empirically. 124

4.13 The number of satellites that, for at least 5% of the Monte Carlo trials, remained viable candidates after association testing with $P(J_j[k, \boldsymbol{x}_s^*] > \nu_j[k]) < 10^{-3}$. The legends indicate which parameters are assumed unknown. The number of satellites satisfying all elevation masks at each epoch, $|\mathcal{S}_e[k]|$, is shown in light blue. All reference stations modeled with Leica AR20 antennas (Left)or Leica AR25 antennas (right). 128

4.14 Satellites that remained viable candidates on more than 5% of Monte Carlo trials for epoch $k = 1$ and epoch $k = 16$ when the reference stations are modeled as using the Leica AR25 antenna. The legend gives symbols for the true satellite and for valid candidates under scenarios with an increasing number of unknown transmitter parameters. 129

4.15 Normalized estimates of $G^i_{\mathrm{R0}}(\theta)$ for stations MEDI (Leica AR20, $i = 4$) and NICO (Leica AR25, $i = 5$). The blue dots represent collected values of $\mathrm{CNR}_{ij} - G^j_{\mathrm{T}} - L_{ij}$, while the red line indicates a weighted third-order polynomial fit, which becomes $G^i_{\mathrm{R0}}(\theta)$, all normalized to $G^i_{\mathrm{R0}}(0) = 0$. As expected, measurement noise increases significantly as the elevation angle decreases. 130

4.16 CNR time history of tracked GPS L1 C/A (left) and BeiDou B1I (right) signals at the IGS station METG in Finland. The GPS L1 C/A signals are first affected by each interference event, followed by the BeiDou B1I signals. The black line indicates the 2.3-second time-overlapped raw IQ capture interval. 138

4.17 Interference heat maps for day 42 of year 2026 from tracked GPS L1 C/A (left) and BeiDou B1I (right) signals. 138

4.18 Left: Example normalized cross correlation at the optimal FDOA. Right: TDOA measurement time history. 140

4.19 TDOA (left) and FDOA (right) residuals for Cosmos 2546 (NORAD ID 45608). The effects of nearest-sample quantization and Doppler search quantization are evident in the respective plots. 143

Chapter 1

Introduction

Global Navigation Satellite Systems (GNSS) such as GPS provide meter-accurate positioning while offering global accessibility and all-weather, radio-silent operation. Its open access nature has facilitated low-cost hardware development and compatibility, allowing it to emerge as the standard for applications requiring precise position, navigation, and timing (PNT). GNSS is a key enabling technology for all aspects of autonomous systems, location-based services, environmental monitoring, and precision agriculture. Additionally, modern national infrastructure is critically dependent on GNSS. This includes all modes of transportation, telecommunications, power distribution, and financial networks.

1.1 The GNSS Interference Problem

GNSS is revolutionary technology critical to modern civilian and military operations, however, it is incredibility vulnerable to deliberate interference and due to two main reasons. The first reason is that the GNSS signals are incredibly weak: by the time they reach the surface of Earth, they have no more flux density than light received from a 50 W bulb at a distance of

2000 km. Because of this, GNSS receivers are vulnerable to signal jamming, a blunt radio frequency (RF) attack where an an attacker transmits strong interference signals at the same frequency. This type of attack is aimed at denial of service and can easily achieved due to the weak received signal power of GNSS signals.

The second reason is that GNSS signals are highly predictable due to their open-access interface specifications. This high signal predictability renders civil users vulnerable to spoofing attacks, which is when an attacker transmits counterfeit signals that mimic the known GNSS signal structure. This type of attack is aimed at deception, as the attacker emits an ensemble of false GNSS signals intending that the victim receiver(s) accept them as authentic GNSS signals, inferring a false position fix and/or clock offset. When a GNSS receiver cannot confidently differentiate between authentic and spoofing signals, it is rendered useless—or worse—hazardously misleading. The same feature that allowed for global low-cost adoption serves as a zero-day exploit.

To compound this, the cost of deploying these disruptive interference sources heavily favors the attacker. Consider that nearly all necessary components for a 1 kW GNSS jammer are present in a $100 microwave oven, whereas a jam-hardened GNSS receiver costs more than $10,000. Furthermore, the decreasing cost of sophisticated software-defined radios and publicly accessible code mean that potent spoofing attacks can be achieved at a low financial cost. This fundamental cost asymmetry and decreasing technical barrier to entry further exacerbates the issue.

The escalating threat of GNSS interference has become a primary focus of concern for academic researchers, industrial stakeholders, and United States government agencies. The foundational document for this discourse is the 2001 United States Department of Transportation assessment, frequently cited as the "Volpe Report," which initially identified the profound vulnerabilities of the national transportation infrastructure to civil GPS interference [73]. As the first publicly accessible technical evaluation of its kind, the Volpe Report established the baseline for understanding how GNSS spoofing could fundamentally compromise the integrity of critical infrastructure. Since the report has been released, GNSS interference has evolved rapidly. GNSS interference is no longer just an academic topic, or even strictly limited to the military battlegrounds. The civilian aviation and maritime industries are seeing significant electronic warfare spillover from nearby conflict zones. The author of this dissertation was the first to notice widespread GPS spoofing capturing the receivers onboard aircraft [7] and his research on the topic has garnered attention from the media [4, 15, 46, 47].

In 2024 the OPSGROUP, an independent, international organization designed for aviation professionals, released a report detailing the GPS spoofing threat in aviation based on the collective findings of 950 industry participants [139]. There was a staggering increase in GNSS interference compromising the integrity of international airspace. The report revealed a 500% increase in spoofing incidents during 2024, with the volume of affected aircraft surging from a daily average of 300 flights in the first quarter to approximately

1,500 flights per day by August. This period of heightened activity resulted in an estimated 310,000 total spoofing encounters throughout the year. The report concluded that GPS is deeply interwoven into critical systems, and beyond simple navigational displacement, the report emphasizes the cascading failures that occur when a single corrupted data point triggers a systemic collapse of non-navigational components.

In 2026, the Royal Institute of Navigation released a similar report focusing specifically on the effects of GNSS interference on maritime systems [52]. The report concluded that GNSS interference is no longer an isolated regional anomaly but a global cyber-physical security crisis. In the second quarter of 2025 alone, over 10,000 vessels were officially reported as affected by GNSS interference, an eight-fold increase over the previous quarter. Just as in aviation, GNSS is integrated into nearly every facet of bridge and engine room operations, meaning GNSS is a critical single point of failure.

In response to these systemic vulnerabilities, modern strategic discourse has shifted toward the "Protect, Toughen, and Augment" (PTA) framework as a definitive methodology for GNSS resilience [48]. The framework, championed by GPS pioneer Dr. Brad Parkinson, is designed to eliminate the single point of failure inherent in current GNSS-dependent systems. By diversifying the PNT portfolio, the PTA framework aims to mitigate the projected economic impact of a total GNSS collapse, which a 2019 United States Department of Transportation estimate places at over $1 billion per day. The goal of this dissertation is to develop capabilities to protect the GNSS spectrum.

1.2 GNSS Interference Defenses and PNT Alternatives

The publication of the Volpe Report catalyzed a fundamental shift in the navigation community, establishing GNSS resilience and security as a primary research thrust and initiating an extensive body of scholarly work. This academic interest intensified significantly following the development of the first publicly acknowledged GNSS spoofer by Professor Todd Humphreys and the Radionavigation Laboratory at the University of Texas at Austin [9,10,68,76]. By demonstrating the inherent vulnerabilities of unauthenticated signals, this development served as a critical inflection point for global PNT security. Furthermore, the laboratory's contributions have provided the research community with indispensable high-fidelity datasets and standardized test batteries, which remain the industry benchmark for verifying the efficacy of novel spoofing detection techniques [64, 66]. Countermeasures to GNSS interference can generally be considered in one of three main categories: (1) receiver hardening, (2) cryptographic authentication, and (3) alternate PNT.

The objective of GNSS hardening is to enhance a receiver's resilience against interference through multi-layered defensive strategies implemented across the entire radio frequency (RF) processing chain. These measures encompass specialized antenna architectures, robust front-end implementations, sophisticated signal quality monitoring (SQM), and GNSS observables monitoring for interference detection. Spatial processing techniques, for instance, utilize multi-antenna arrays to monitor the direction of arrival of incoming signals [114, 116]. Reliable spoofing detection techniques even exist for chal-

lenging environments such as dynamic platforms in urban areas where strong multipath and in-band noise are common [50,51,105,106,116,135]. Consistency checks between the estimated signal and onboard inertial sensors can provide quick and reliable spoofing detection [28, 29, 85, 127]. Multi-frequency and multi-constellation GNSS receivers offer redundancy so that single-frequency or single-constellation failures can be traversed [27, 102].

Systematic cryptographic authentication techniques provide a robust mechanism for securing civil GNSS navigation and timing by embedding unpredictable, verifiable features directly into the signal structure [63,136,137]. A landmark development in this field is Galileo's Open Service Navigation Message Authentication (OSNMA), which represents the first operational GNSS service to implement a civil cryptographic authentication framework [44]. By interleaving digital signatures within the broadcast navigation message, OSNMA enables receivers to deterministically verify the provenance and integrity of each data packet. This cryptographic verification ensures that any attempt by a malicious actor to alter the ephemeris or timing data will result in a signature mismatch, allowing the receiver to autonomously reject the corrupted information before it can propagate through critical onboard systems.

Because of the vulnerability to interference, alternate forms of navigation has been an emerging research area. Exploiting signals of opportunity for alternate navigation is a well-studied topic with a growing literature [70, 75]. Mega Low Earth Orbit (LEO) communications constellations can be opportunistically exploited for navigation and 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. But because PNT is not the original purpose of these signals, tremendous effort has gone into understanding the signal structure and timing properties [67, 79, 117]. Once these properties are discovered, these signals can be exploited for PNT [80,97]. Other forms of alternate navigation include terrestrial radio-based positioning [128, 129], alternate sensors such as cameras [17, 130], optimal trajectory design [41, 42], and terrain navigation [59, 60, 60].

Despite recent breakthroughs in GNSS interference mitigation, spoofing detection, and alternative navigation, a significant portion of the deployed GNSS receiver continues to operate on legacy platforms that lack the hardware architecture to support these modern defenses and alternatives. Adding the requisite additional hardware to be able to perform such techniques might not be an option due to size, weight, and power (SWaP) constraints. Furthermore, in heavily regulated sectors like civil aviation and commercial shipping, the lengthy and cost-prohibitive certification cycles for new avionics mean that even validated technologies may take years to reach widespread deployment. There is an urgent demand for immediate-impact defenses to help GNSS users protect themselves against harmful interference.

1.3 Protecting the GNSS Spectrum

Given the widespread deployment of legacy GNSS receivers that lack the architectural capacity to detect even rudimentary spoofing or integrate

alternative PNT sources, there is an urgent need alert GNSS users of challenging environments. This can be achieved through the strategic or serendipitous deployment of GNSS receivers and software-defined radios as they can be specifically configured or repurposed to monitor the GNSS spectrum. By leveraging these distributed sensors to detect and geolocate interference in real-time, authorities can provide the comprehensive spectral situational awareness necessary to safeguard vulnerable PNT systems from disruptions.

It has been shown that dedicated terrestrial-based GNSS interference monitoring systems provide valuable spectrum situational awareness to enhance navigation security [54, 55, 82, 123]. Some of these static terrestrial systems have demonstrated the ability to geolocate the interference source [11, 95]. The more general case of localizing an emitter transmitting an arbitrary wideband signal with terrestrial and airborne receivers has been extensively studied [5, 101, 107, 134]. However, because the receivers were either at fixed locations or tactically deployed in the nearby airspace, only emitters in the immediate area could be geolocated. There remains a need for global, persistent, low-latency, and accurate GNSS interference detection and localization.

GNSS receivers in LEO are a proven asset for detecting, classifying, and geolocating terrestrial GNSS interference [86, 100]. GNSS interference monitoring from LEO offers worldwide coverage with a frequent refresh rate, making it possible to maintain a common operating picture of terrestrial interference. The economic feasibility of this has been transformed by the dramatic reduction in orbital launch costs over the last decade.

Accurate single-satellite geolocation of emitters whose 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 [25, 26, 30, 36–38, 100]. 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.

Multiple time-synchronized receivers can exploit time- and frequency-difference-of-arrival (T/FDOA) measurements to estimate the emitter location. The author this dissertation was able to geolocate over 30 GNSS interference sources across the Near East from a dual-satellite time-synchronized capture [23, 24]. One of the potential downsides to this approach is its computational burden due to the large search space, however, it can be significantly alleviated by making use of a graphic processing unit (GPU) [22].

A critical evolution in the GNSS threat landscape is that GNSS interference is no longer confined to terrestrial or near-terrestrial sources; it is can be emanating from space-based platforms. Just as space-based receivers can detect and geolocate terrestrial interference, terrestrial-based receivers can detect and geolocate space-based interference. The author of this dissertation was able to investigate and identify a space-based GNSS interference source responsible for multiple transient wide-area GNSS outage over Europe, Greenland, and Canada [31, 32]. GNSS interference monitoring is essential for protecting GNSS users against evolving threats.

1.4 Thesis Statement and Expected Contributions

This dissertation defends the following thesis statement:

Pairing optimal detection and estimation algorithms with networks of terrestrial- and space-based GNSS receivers can provide unprecedented levels of global and persistent GNSS interference monitoring.

The contributions of this dissertation are summarized as follows:

- Chapter 2: Single-Satellite-Based Geolocation of Broadcast GNSS Spoofers from Low Earth Orbit
  - A single-satellite Doppler-based geolocation technique crafted specifically for indiscriminate broadcast-type GNSS spoofing signals and its experimental verification in which a LEO-based receiver captures broadcast GNSS spoofing signals transmitted from a known ground station on a non-GNSS frequency band.
  - A derivation of the measurement covariance matrix that incorporates the time-correlation between the measurements due to the transmitter's clock drift and an analysis of how unmodeled oscillator instability degrades the geolocation estimate.
  - A presentation of an optimal spoofing detection strategy based on receiver clock drift monitoring.

  - A derivation and analysis of a spoofing 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.

- Chapter 3: Dual-Satellite Geolocation of Terrestrial GNSS Interference from Low Earth Orbit
  - A presentation of estimators to geolocate radio frequency (RF) interference sources from raw in-phase and quadrature (IQ) samples from multiple time-synchronized receivers based on time- and frequency-difference-of-arrival (T/FDOA).
  - An analysis of the homeomorphic mapping between the CAF domain and the position domain of an interference signal and showcasing how direct geolocation is better-suited for interference geolocation with LEO-based receivers.
  - Signals from narrowband, matched-code, and continuous-wave jammers captured in the GNSS frequency bands by two time-synchronized LEO receivers over the Eastern Mediterranean are analyzed, with the emitters being geolocated.

- Chapter 4: Space-Based GNSS Interference Detection and Estimation
  - An investigation into the spatial, temporal, and spectral properties of space-based GNSS interference and the presentation of a wide-

area interference detection framework using carrier-to-noise ratio (CNR) measurements from a network of GNSS reference stations.

- A presentation of multiple received-power-based satellite identification strategies based on elevation masking and the generalized likelihood ratio test (GLRT) to provide performance insights and limitations of received-power-based association.
- An analysis to identify the offending satellite based on T/FDOA measurements generated from IQ samples recorded by two receivers in Europe, along with a ephemeris and timing error sensitivity analysis.

1.5 Organization

Chapter 2 offers an experimental demonstration of single-satellite geolocation of broadcast GNSS spoofers and an analysis of various error sources. Chapter 3 explores dual-satellite geolocation of GNSS interference from LEO and presents an analysis of interference signals captured in the GNSS frequency bands by two time-synchronized LEO receivers over the Eastern Mediterranean, with the emitters being geolocated. Chapter 4 investigates a transient space-based GNSS interference source that has caused continental-wide GNSS interference and presents detection and estimation algorithms to identify the source. Chapter 5 concludes the dissertation. The mathematical notation is independent for each chapter.

1.6 List of Publications

1.6.1 Journal Publications

[J1] **Zachary L. Clements**, Argyris Kriezis, and Todd E. Humphreys. Detection and Estimation of Transient Space-Based GNSS Interference. *NAVIGATION*, 2026. In Preparation.

[J2] Zacharias .M. Komodromos, Samuel C. Morgan, **Zachary L. Clements**, Wenkai Qin, W. Jeremy Morrison, Todd E. Humphreys. Network-Aided LEO PNT from OneWeb: Pseudorange and Doppler Results and Insights. *NAVIGATION*, 2026. In Preparation.

[J3] Samuel C. Morgan, Zacharias .M. Komodromos, Wenkai Qin, Andrew .M. Graff, **Zachary L. Clements**, W. Jeremy Morrison, Todd E. Humphreys. Towards Pseudorange-Based Positioning with Starlink: A Mock Fused LEO GNSS Implementation. *npj Wireless Technology*, 2026. In Preparation.

[J4] **Zachary L. Clements**, Patrick B. Ellis, Iain Goodridge, Matthew J. Murrian, Mark L. Psiaki, and Todd E. Humphreys. Single-Satellite-Based Geolocation of Broadcast GNSS Spoofers from Low Earth Orbit. *NAVIGATION*, 2026.

[J5] Wenkai Qin, Andrew M. Graff, **Zachary L. Clements**, Zacharias .M. Komodromos, Todd E. Humphreys. Timing Properties of the Starlink Ku-Band Downlink. *IEEE Transactions on Aerospace and Electronic Systems*, 2026.

1.6.2 Conference Publications

[C1] **Zachary L. Clements** and Todd E. Humphreys. Transient Space-Based GNSS Interference: Observations and Analysis. In *Proceedings of the ION GNSS+ Meeting*, 2025.

[C2] Samuel C. Morgan, Zacharias .M. Komodromos, Wenkai Qin, **Zachary L. Clements**, Andrew .M. Graff, W. Jeremy Morrison, Todd E. Humphreys. A Mock Implementation of Fused LEO GNSS. In *Proceedings of the IEEE/ION PLANS Meeting*, 2025.

[C3] Zacharias .M. Komodromos, Samuel C. Morgan, **Zachary L. Clements**, Wenkai Qin, W. Jeremy Morrison, Todd E. Humphreys. Network-Aided Pseudorange-Based LEO PNT from OneWeb. In *Proceedings of the IEEE/ION PLANS Meeting*, 2025.

[C4] **Zachary L. Clements**, Iain Goodridge, Patrick Ellis, Matthew J. Murrian, and Todd E. Humphreys. Demonstration of Single-Satellite GNSS Spoofer Geolocation. In *Proceedings of the ION International Technical Meeting*, 2024.

[C5] **Zachary L. Clements**, Patrick Ellis, and Todd E. Humphreys. Dual-Satellite Geolocation of Terrestrial GNSS Jammers from Low Earth Orbit. In *Proceedings of the IEEE/ION PLANS Meeting*, 2023.

[C6] **Zachary L. Clements**, Patrick Ellis, Mark L. Psiaki, and Todd E. Humphreys. Geolocation of GNSS Spoofing Signals from Low Earth

Orbit. In *Proceedings of the ION GNSS+ Meeting*, 2022.

[C7] **Zachary L. Clements**, James E. Yoder, and Todd E. Humphreys. Carrier-phase and IMU based GNSS Spoofing Detection for Ground Vehicles. In *Proceedings of the ION International Technical Meeting*, 2022.

[C8] **Zachary L. Clements**, Peter A. Iannucci, Todd E. Humphreys, and Thomas Pany. Optimized Bit-Packing for Bit-Wise Software-Defined GNSS Radio. In *Proceedings of the ION GNSS+ Meeting*, 2021.

1.6.3 Other Major Publications

[P1] **Zachary L. Clements** and Todd E. Humphreys. Orbital Outlaw: GNSS Interference from Space. *Royal Institute of Navigation: Navigation News Jan/Feb Edition*, 2026.

[P2] Jacob S. Clements and **Zachary L. Clements**. GPU-accelerated Direct Geolocation of GNSS Interference. *University of Maryland M.S. Scholarly Paper*, 2024.

[P3] **Zachary L. Clements**, Patrick Ellis, and Todd E. Humphreys. Pinpointing GNSS Interference from Low Earth Orbit. *Inside GNSS*, 18(5), 2023.

[P4] **Zachary L. Clements**, James E. Yoder, and Todd E. Humphreys. GNSS Spoofing Detection: An Approach for Ground Vehicles Using Carrier-Phase and Inertial Measurement Data. *GPS World*, 34(2), 2023.

Chapter 2

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

Abstract

This chapter 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 chapter 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.

2.1 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 [71,114]. A successful spoofing attack may lead to serious consequences.

The academic community has long warned the public about the threat of GNSS spoofing [62,68,120]. Within the past decade, significant progress has been made in GNSS spoofing detection and mitigation [65, 71, 114, 115, 118]. 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 [50, 51, 105, 106, 116, 135]. Consistency checks between the estimated signal and onboard inertial sensors can provide quick and reliable spoofing detection [28,29,85,127]. Monitoring the clock state can also be used to detect spoofing [69, 72, 78]. Cryptographic authentication techniques are currently being developed and implemented to verify received signals [2,3,44, 63,77,93].

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 [4,16,43,45,108,126,139]. 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 [76]. 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 [21, 23, 24, 86, 92, 100]. 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 [8]. 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 [23, 24, 61, 101, 121]. 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 chapter were able to geolocate over 30 GNSS interference sources across the Near East from a dual-satellite time-synchronized capture [23, 24]. 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 chapter 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 [36–38, 100]. 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 [76]. A key contribution of the current chapter 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. [20] also presented a single-receiver spoofer geolocation technique based on counterfeit clock observables. However, [20] 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 chapter'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 chapter 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.

2.2 Signal Models

2.2.1 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 $\tilde{\boldsymbol{x}} = [\, \boldsymbol{r}_{\tilde{\mathrm{r}}}^{\mathsf{T}},\ \delta t_{\tilde{\mathrm{r}}},\ \boldsymbol{v}_{\tilde{\mathrm{r}}}^{\mathsf{T}},\ \delta \dot{t}_{\tilde{\mathrm{r}}}\,]^{\mathsf{T}}$, where $\boldsymbol{r}_{\tilde{\mathrm{r}}}$ is the spoofed position in Earth-centered-Earth-fixed (ECEF) coordinates, $\delta t_{\tilde{\mathrm{r}}}$ is the spoofed clock bias increment, $\boldsymbol{v}_{\tilde{\mathrm{r}}}$ is the spoofed velocity, and $\delta \dot{t}_{\tilde{\mathrm{r}}}$ 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

$$x(t) = \sum_{n=1}^{N} s_n(t) \tag{2.1}$$

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

$$s_n(t) = A_n D_n[t - \tau_n(t)] C_n[t - \tau_n(t)] \exp[j2\pi\theta_n(t)] \tag{2.2}$$

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 [114]. The Doppler of the nth spoofing signal is related to $\theta_n(t)$ by

$$\tilde{f}_n(t) = \frac{d}{dt}\theta_n(t) \tag{2.3}$$

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 Doppler of the nth transmitted spoofing signal may be modeled as

$$\tilde{f}_n(t) = -\frac{1}{\lambda}\hat{\boldsymbol{r}}_{\tilde{\mathrm{r}}n}^{\mathsf{T}}(t)(\boldsymbol{v}_{\tilde{\mathrm{r}}}(t) - \boldsymbol{v}_{\tilde{\mathrm{s}}n}(t)) - \frac{c}{\lambda}\big(\delta\dot{t}_{\tilde{\mathrm{r}}}(t) - \delta\dot{t}_{\tilde{\mathrm{s}}n}(t)\big) \tag{2.4}$$

where λ is the carrier wavelength, c is the speed of light, $\hat{\boldsymbol{r}}_{\tilde{\mathrm{r}}n}$ is the unit vector pointing from the nth spoofed navigation satellite to the spoofed position, both in ECEF coordinates, $\boldsymbol{v}_{\tilde{\mathrm{r}}}$ is the spoofed receiver velocity, $\boldsymbol{v}_{\tilde{\mathrm{s}}n}$ is the nth spoofed navigation satellite velocity, and $\delta\dot{t}_{\tilde{\mathrm{s}}n}$ is the spoofed clock drift of the nth 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 (2.4) that compensates for the relative motion between the victim and spoofer, but in the case of broadcast spoofing, this term is zero.

2.2.2 Received Doppler Model

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

$$f_{\mathrm{D}}(t) = -\frac{1}{\lambda}\hat{\boldsymbol{r}}^{\mathsf{T}}(t)(\boldsymbol{v}_{\mathrm{r}}(t) - \boldsymbol{v}_{\mathrm{t}}(t)) - \frac{c}{\lambda}\big(\delta\dot{t}_{\mathrm{r}}(t) - \delta\dot{t}_{\mathrm{t}}(t)\big) \tag{2.5}$$

where $\hat{\boldsymbol{r}}$ is the unit vector pointing from the transmitter to the receiver, $\boldsymbol{v}_{\mathrm{r}}$ is the velocity of the receiver, $\boldsymbol{v}_{\mathrm{t}}$ is the velocity of the transmitter, $\delta\dot{t}_{\mathrm{r}}$ is the clock drift of the receiver, and $\delta\dot{t}_{\mathrm{t}}$ is the clock drift of the transmitter. Note that this is a simplified Doppler model that neglects higher-order terms. [113] presented a complete Doppler model. For the purposes of this chapter, 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 ($\boldsymbol{v}_{\mathrm{t}}(t) = 0$), as shown in Fig. 2.1. [25] provided an analysis of how spoofer motion affects the geolocation solution. 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 chapter.

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.

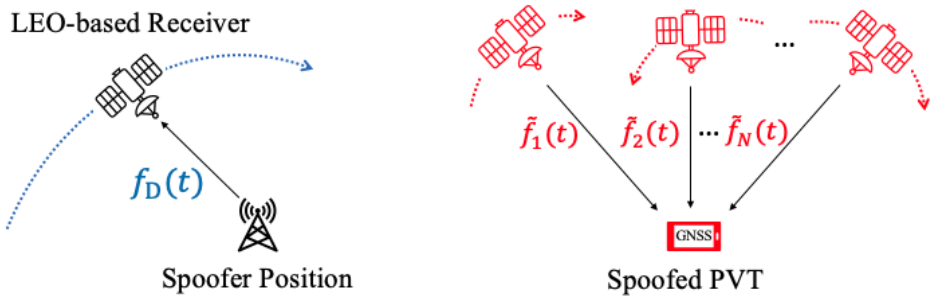


Figure 2.1: Shown here are the Doppler components in single-satellite spoofer geolocation. The Doppler components corresponding to (2.5) are shown on the left. The Doppler components for each spoofing signal corresponding to (2.4) are shown in red to the right.

Dropping time indices for clarity, the observed Doppler of the nth spoofing signal at the moving receiver, f_n, may be written as

$$\begin{aligned} f_n &= f_{\mathrm{D}} + \tilde{f}_n \\ &= -\frac{1}{\lambda}\hat{\boldsymbol{r}}^{\mathsf{T}}\boldsymbol{v}_{\mathrm{r}} - \frac{c}{\lambda}\left(\delta\dot{t}_{\mathrm{r}} - \delta\dot{t}_{\mathrm{t}}\right) - \frac{1}{\lambda}\hat{\boldsymbol{r}}_{\tilde{\mathrm{r}}n}^{\mathsf{T}}(\boldsymbol{v}_{\tilde{\mathrm{r}}} - \boldsymbol{v}_{\tilde{\mathrm{s}}n}) - \frac{c}{\lambda}\left(\delta\dot{t}_{\tilde{\mathrm{r}}} - \delta\dot{t}_{\tilde{\mathrm{s}}n}\right) \end{aligned} \tag{2.6}$$

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 by [100], $\tilde{f}_n = 0$. 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 (2.6) 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 $\hat{\boldsymbol{r}}^{\mathsf{T}}(t)\boldsymbol{v}_{\mathrm{r}}(t)$, the range-rate time history between transmitter and receiver, which can be exploited for geolocation.

2.3 Conceptual Overview of Broadcast GNSS Spoofer Geolocation

This section presents an overview of the technique for spoofer geolocation originally presented by [25, 26]. The common Doppler components across all spoofing signals from (2.6) are indicated below:

$$f_n = \underbrace{-\frac{1}{\lambda}\hat{\boldsymbol{r}}^{\mathsf{T}}\boldsymbol{v}_{\mathrm{r}} - \frac{c}{\lambda}\big(\delta\dot{t}_{\mathrm{r}} - \delta\dot{t}_{\mathrm{t}}\big)}_{\text{common}} - \frac{1}{\lambda}\hat{\boldsymbol{r}}_{\tilde{\mathrm{r}}n}^{\mathsf{T}}(\boldsymbol{v}_{\tilde{\mathrm{r}}} - \boldsymbol{v}_{\tilde{\mathrm{s}}n}) - \frac{c}{\lambda}\big(\underbrace{\delta\dot{t}_{\tilde{\mathrm{r}}}}_{\text{common}} - \delta\dot{t}_{\tilde{\mathrm{s}}n}\big) \tag{2.7}$$

All common Doppler terms can be lumped into a single term

$$\gamma(t) = \frac{1}{c}\hat{\boldsymbol{r}}^{\mathsf{T}}(t)\boldsymbol{v}_{\mathrm{r}}(t) + \delta\dot{t}_{\mathrm{r}}(t) - \delta\dot{t}_{\mathrm{t}}(t) + \delta\dot{t}_{\tilde{\mathrm{r}}}(t) \tag{2.8}$$

so that (2.6) may be written

$$f_n = -\frac{1}{\lambda}\hat{\boldsymbol{r}}_{\tilde{\mathrm{r}}n}^{\mathsf{T}}(\boldsymbol{v}_{\tilde{\mathrm{r}}} - \boldsymbol{v}_{\tilde{\mathrm{s}}n}) - \frac{c}{\lambda}\big(\gamma - \delta\dot{t}_{\tilde{\mathrm{s}}n}\big) \tag{2.9}$$

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

$$\hat{\boldsymbol{x}}(t) = [\,\hat{\boldsymbol{r}}_{\tilde{\mathrm{r}}}^{\mathsf{T}}(t),\ \hat{\xi}(t),\ \hat{\boldsymbol{v}}_{\tilde{\mathrm{r}}}^{\mathsf{T}}(t),\ \hat{\gamma}(t)\,]^{\mathsf{T}} \tag{2.10}$$

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 $\hat{\gamma}(t)$ [56]. [26, 104] provide a brief review of PVT estimation from pseudorange and Doppler measurements. Note that the estimated receiver clock bias $\hat{\xi}(t)$ will include $\delta t_{\tilde{\mathrm{r}}}$ as a component but will not in general equal $\delta t_{\tilde{\mathrm{r}}}$.

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 $\tilde{f}_n(t)$ for all $n \in \{1, 2, ..., N\}$.

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 $\hat{\boldsymbol{r}}^{\mathsf{T}}(t)\boldsymbol{v}_{\mathrm{r}}(t)$, 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.

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

affected by the spoofing, and and thus $\delta\dot{t}_{\mathrm{r}}(t)$ 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 chapter, 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 by [100]. 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. PVT can be trivially maintained by a multi-constellation receiver when only single-constellation spoofing signals are present. In the event that all GNSS signals are unavailable due to terrestrial interference, knowledge of the receiver's position and velocity can be maintained by using orbital propagation models such as simplified general perturbations 4 (SGP4). Over short periods, the orbit is stable enough so that the receiver will be able to maintain sufficient PVT accuracy from the onset of GNSS denial.

The terms $\delta\dot{t}_{\mathrm{t}}(t)$ and $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ originate from the spoofer. Specifically, $\delta\dot{t}_{\mathrm{t}}(t)$ originates from the spoofer's hardware, while $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ 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 [100]. The

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

Additionally, 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 time-varying term. If this term were to vary rapidly with time, it would cause trouble for this chapter's technique. Relatedly, if the targeted victim's position and velocity were somehow accurately known to the LEO-based receiver, this chapter'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 chapter's technique can be applied without modification.

Section 2.6 explores the consequences for geolocation of cases where $\delta\dot{t}_{\mathrm{t}}(t)$ departs from a constant model. It also presents an analysis of how

aggressively an attacker can ramp $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ 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 $\delta\dot{t}_{\mathrm{t}}(t) + \delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ translates to geolocation error.

2.4 Spoofer Geolocation with γ

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

2.4.1 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, atmospheric and relativistic effects, ephemeris errors, 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 $i \in \mathcal{I} = \{1, 2, ..., I\}$ is the solution index within a given data capture interval. Let $z[i]$ denote the ith measurement to be used for

spoofer geolocation, modeled as

$$z[i] = c\hat{\gamma}[i] = c\gamma[i] + w_{\mathrm{a}}[i], \quad i \in \mathcal{I} \tag{2.11}$$

The velocity-equivalent estimation error $w_{\mathrm{a}}[i]$, which has units of m/s, is a discrete-time noise process with $\mathbb{E}[w_{\mathrm{a}}[i]] = 0$ and $\mathbb{E}[w_{\mathrm{a}}[i]w_{\mathrm{a}}[j]] = \sigma_{\mathrm{a}}^2\delta_{ij}$, for all $i, j \in \mathcal{I}$. Section 2.4.2 will justify this model's assumption that $w_{\mathrm{a}}[i]$ 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, $\delta\dot{t}_{\mathrm{r}}$ is assumed to be known and fully compensated; accordingly, it will be neglected hereafter. Additionally, $\delta\dot{t}_{\tilde{\mathrm{r}}}$ 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 $\delta\dot{t}_{\mathrm{t}}(t)$. Let $\delta\dot{t}_{\mathrm{t}}[i] = \delta\dot{t}_{\mathrm{t}}(i\Delta t), i \in \mathcal{I}$. Over a capture interval, $\delta\dot{t}_{\mathrm{t}}[i]$ is modeled as

$$c\delta\dot{t}_{\mathrm{t}}[i] = c\delta\dot{t}_{\mathrm{t}}[0] + b[i], \quad i \in \mathcal{I} \tag{2.12}$$

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

$$b[i] = \sum_{k=1}^{i} v[k], \quad i \in \mathcal{I} \tag{2.13}$$

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 $\mathbb{E}[w_{\mathrm{a}}[k]v[j]] = 0$ for all $k, j \in \mathcal{I}$, and $b[0] = 0$. Using

the model in [14, Chap. 8], σ_v^2 can be characterized as

$$\sigma_v^2 = 2\pi^2 h_{-2} \Delta t c^2 \tag{2.14}$$

where h_{-2} is the first parameter of the standard clock model based on the fractional frequency error power spectrum [100]. Scaling by c^2 converts to units of $(\mathrm{m/s})^2$.

Note that $\delta\dot{t}_{\mathrm{t}}[0]$ and $\delta\dot{t}_{\tilde{\mathrm{r}}}$ 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

$$b_0 = -c\delta\dot{t}_{\mathrm{t}}[0] + c\delta\dot{t}_{\tilde{\mathrm{r}}} \tag{2.15}$$

$$w[i] = w_{\mathrm{a}}[i] + b[i], \quad i \in \mathcal{I} \tag{2.16}$$

Given all of this, (2.11) is rewritten so that the final measurement model takes the form

$$z[i] = \hat{\boldsymbol{r}}_i^{\mathsf{T}} \boldsymbol{v}_{\mathrm{r},i} + b_0 + w[i], \quad i \in \mathcal{I} \tag{2.17}$$

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

$$\begin{aligned} R[i,j] &= \mathbb{E}[w[i]w[j]] \\ &= \mathbb{E}\left[\left(w_{\mathrm{a}}[i] + \sum_{k=1}^{i} v[k]\right)\left(w_{\mathrm{a}}[j] + \sum_{l=1}^{j} v[l]\right)\right] \\ &= \mathbb{E}[w_{\mathrm{a}}[i]w_{\mathrm{a}}[j]] + \mathbb{E}\left[\left(\sum_{k=1}^{i} v[k]\right)\left(\sum_{l=1}^{j} v[l]\right)\right] \\ &= \mathbb{E}[w_{\mathrm{a}}[i]w_{\mathrm{a}}[j]] + \sum_{k=1}^{i}\sum_{l=1}^{j} \mathbb{E}[v[k]v[l]] \\ &= \sigma_{\mathrm{a}}^2\delta_{ij} + \sigma_v^2 \min\{i,j\} \end{aligned} \tag{2.18}$$

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

$$R = R_{\mathrm{a}} + R_{\mathrm{b}} \quad \text{where} \quad R_{\mathrm{a}} = \sigma_{\mathrm{a}}^2\, \mathbb{I}_{I\times I} \quad \text{and} \quad R_{\mathrm{b}} = \sigma_v^2\, M_{I\times I} \tag{2.19}$$

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.

2.4.2 Effects of Estimated γ

One might question the choice to model the estimation error process $w_{\mathrm{a}}[i] = c(\hat{\gamma}[i] - \gamma[i])$ 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 $\tilde{\boldsymbol{x}}[i]$ 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 [6, Chap. 5]

$$\mathbb{E}\left[\tilde{\boldsymbol{x}}[i+1]\tilde{\boldsymbol{x}}^{\mathsf{T}}[i]\right] = (\mathbb{I} - W[i+1]H[i+1])F[i]P[i] \tag{2.20}$$

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

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 (2.10). Assume the receiver's PVT estimator applies a dynamics model consistent with a static position and the clock process noise model from [14]. 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 (2.14). Fig. 2.2 shows the Pearson correlation coefficient for $w_{\mathrm{a}}[i]$ 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. 2.2 indicates that, for $h_{-2} \geq 3 \times 10^{-19}$,

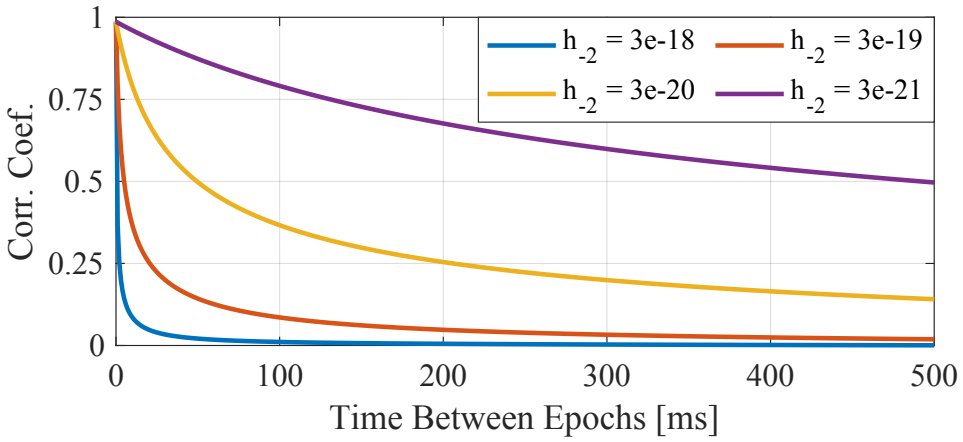


Figure 2.2: The Pearson correlation coefficient between sequential estimation errors $w_{\mathrm{a}}[i]$ as a function of time between estimation epochs for various values of h_{-2}. As the receiver's modeled process noise intensity increases, the time correlation between between estimation errors decreases.

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 $w_{\mathrm{a}}[i]$ 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.

2.4.3 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 $\boldsymbol{x}$

$$\boldsymbol{x} = \begin{bmatrix} \boldsymbol{r}_{\mathrm{t}} \\ b_0 \end{bmatrix} \tag{2.21}$$

where $\boldsymbol{r}_{\mathrm{t}}$ is the transmitter's ECEF position and b_0 is the unknown measurement bias. Let $\boldsymbol{z}$ represent the $I \times 1$ stacked measurement vector. The standard weighted nonlinear least-squares cost function is

$$J(\boldsymbol{x}) = \frac{1}{2}(\boldsymbol{z} - \boldsymbol{h}(\boldsymbol{x}))^{\mathsf{T}} R^{-1} (\boldsymbol{z} - \boldsymbol{h}(\boldsymbol{x})) \tag{2.22}$$

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

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

$$H = \begin{bmatrix} \frac{dh_1}{d\boldsymbol{r}_{\mathrm{t}}} & 1 \\ \vdots & \vdots \\ \frac{dh_I}{d\boldsymbol{r}_{\mathrm{t}}} & 1 \end{bmatrix} \tag{2.23}$$

where

$$\frac{dh_i(\boldsymbol{x})}{d\boldsymbol{r}_{\mathrm{t}}} = \boldsymbol{v}_{\mathrm{r}}^{\mathsf{T}} \frac{\left(\hat{\boldsymbol{r}}_i \hat{\boldsymbol{r}}_i^{\mathsf{T}} - \mathbb{I}_{3\times 3}\right)}{\rho_i} \tag{2.24}$$

is the 1×3 Jacobian of the ith range-rate measurement. The range between the receiver and the transmitter at the ith measurement is denoted ρ_i. This measurement model Jacobian is equivalent to columns 1, 2, 3, and 8 of the Jacobian presented by [113], 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

$$z[I+1] = h_{\text{alt}}(\boldsymbol{x}) + w_{\text{alt}} \tag{2.25}$$

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

$$H_{\text{alt}} = \big[\cos(\phi_{\text{lat}})\cos(\lambda_{\text{lon}}),\ \cos(\phi_{\text{lat}})\sin(\lambda_{\text{lon}}),\ \sin(\phi_{\text{lat}}),\ 0\big] \tag{2.26}$$

where ϕ_{lat} and λ_{lon} are the latitude and longitude of $\boldsymbol{r}_{\text{t}}$, respectively. The measurement vector $\boldsymbol{z}$, vector-valued function $\boldsymbol{h}(\boldsymbol{x})$, 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

$$P_{\boldsymbol{x}\boldsymbol{x}} = \left(H^{\mathsf{T}} R^{-1} H\right)^{-1} \tag{2.27}$$

2.5 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 by [26], which contains a comprehensive description of the special adaptations made to deal with the spoofer's non-GNSS carrier frequency.

2.5.1 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 [66]. 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



Figure 2.3: 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.

could be decoded. Fig. 2.3 shows locations relevant to the demonstration. In the context of this chapter, 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.

2.5.2 Experimental Spoofer Geolocation with γ

The transmitted spoofing signals captured in LEO were processed with the UT RNL's GRID software-defined GNSS receiver [27, 102, 109]. Fig. 2.4 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 [26]. On GRID's display, the 4,810 m/s clock drift (labeled δ`tRdot`) is immediately noticeable. Of course, no oscillator on a

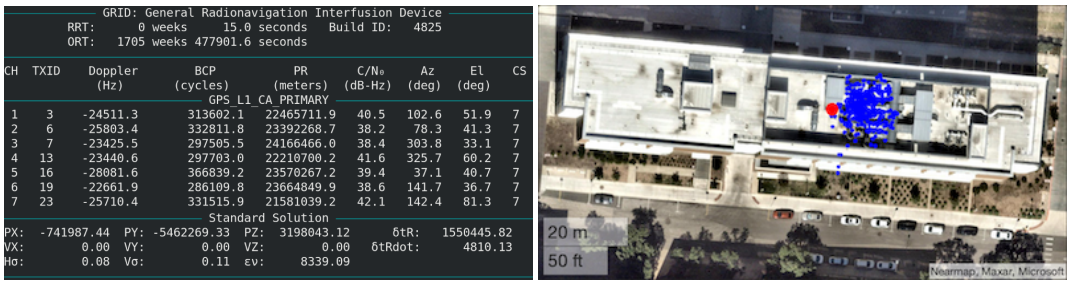


Figure 2.4: Left: UT RNL's GRID receiver display when processing the spoofing signals. Right: A scatter of GRID-derived position solutions. The red dot is the spoofed position. The 3D bias is 45.9 m, mostly concentrated in the vertical direction. This error is attributed to the S-band carrier.

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 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 by [26].

A Doppler-equivalent time-history $\hat{\gamma}(t)$ over 17.75 seconds is shown as the black trace in Fig. 2.5 along with the raw measured Doppler of each

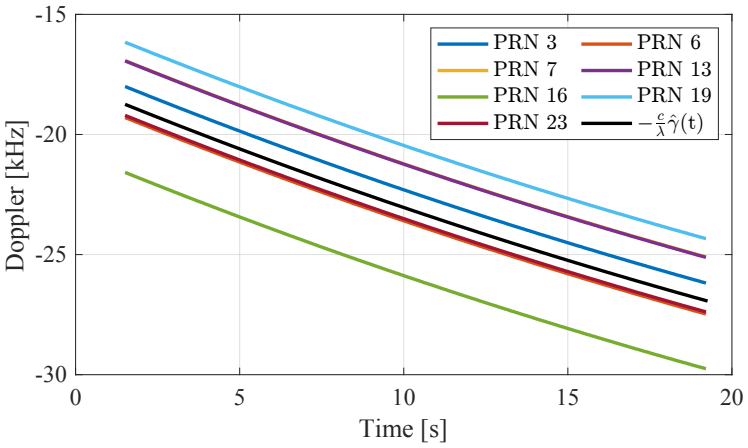


Figure 2.5: Measured Doppler time history of each received spoofing signal. Also shown is the Doppler-equivalent time history $\hat{\gamma}(t)$ (black trace), which is used for geolocation.

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 2.3. The measured Doppler time history of each spoofing signal, as given in (2.6), 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 2.4. The final position fix, shown in Fig. 2.6, was within 68 meters of the true location. 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 eccentric-

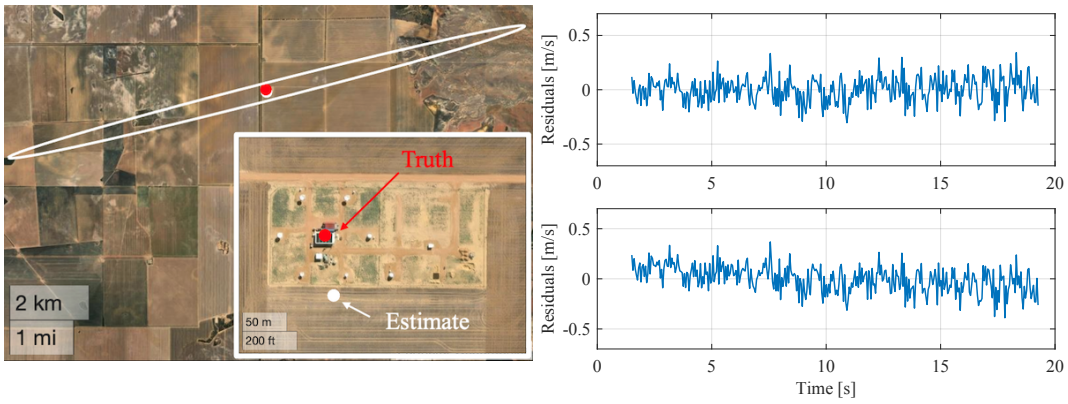


Figure 2.6: Left: Final spoofer position estimate (white) based on $\hat{\gamma}(t)$. Shown in red is the true spoofer location. The error of the final estimate is 68 m. The true emitter is contained within the 95% horizontal error ellipse, derived from (2.27), which has a semi-major axis of 6.7 km. Right: Post-fit range-rate residuals of $\hat{\gamma}(t)$ time history with respect to the estimated spoofer position (top) and true spoofer position (bottom). The residuals with respect to the estimated position are unbiased and have a standard deviation of 0.12 m/s.

ity is dictated by the receiver-transmitter geometry. Shown in Fig. 2.6 are the Doppler post-fit residuals, with respect to the estimated spoofer position and the true spoofer position. The residuals with respect to the estimated spoofer position 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 chapter's geolocation technique.

2.5.3 Experimental Spoofer Geolocation with GNSS Observables

This chapter's advocated technique requires a means of getting ephemerides and clock models of the spoofed navigation satellites implied in the spoof-

ing. 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 [100] can be applied to estimate the spoofer's location. Of course, as shown earlier, this will yield a biased estimate of the spoofer's position because the time-varying frequency term $\tilde{f}_n(t)$ is unmodeled. However, if the spoofing signals induce a static terrestrial 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 spoofer is the dominant term in $f_n(t)$. Shown in Fig. 2.7 are the biased position fixes and corresponding error ellipses when each $f_n(t)$ time history is fed as measurements to the nonlinear least-squares estimator as described in Section 2.4. Only two of the seven 95% error ellipses contain the true spoofer position. The spread of the spoofer 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 spoofer's induced trajectory were dynamic rather than static, the spread of the geolocation estimates would be larger, as shown by [25].

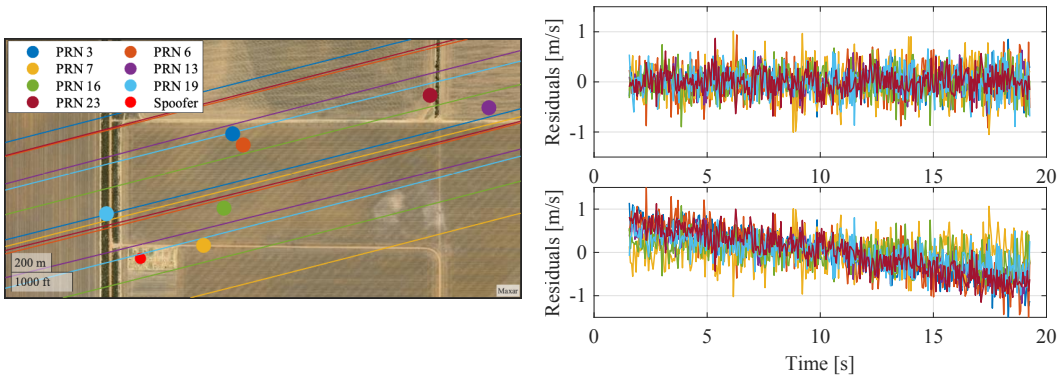


Figure 2.7: Left: Geolocation using the observed Doppler time history of each spoofed PRN. Each individual spoofer position estimate is biased due to the unmodeled frequency component. Right: Range-rate residuals with respect to the estimated spoofer position (top). Range-rate residuals with respect to the true spoofer position (Bottom).

PRN	3	6	7	13	16	19	23
Error [m]	772	763	322	1,897	487	281	1,661

Table 2.1: Final geolocation error using the observed Doppler time history of each spoofed PRN.

Also shown in Fig. 2.7 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.

2.6 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 $\delta\dot{t}_{\tilde{\mathrm{r}}} = 0$ 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 (2.14). This general result applies to any clock quality and any capture geometry.

As an example consider the capture scenario in Section 2.5 for a 20-second capture over Perth. Table 2.2 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 2.2 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 2.2: Theoretical marginal contribution of transmitter clock instability to the 95% horizontal error ellipse for the capture scenario specified in Section 2.5 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 2.5 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 covariance R as specified in (2.19), and then with an incorrect measurement covariance equal to R_{a} (i.e., R_{b} in (2.19) 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. 2.8.

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 conducted, again with the capture geometry specified in Section 2.5. For each trial, both a real-

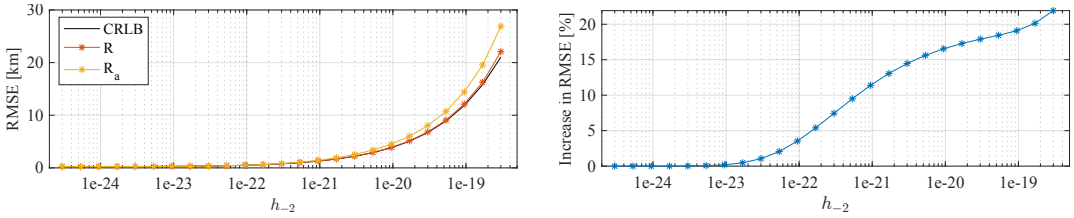


Figure 2.8: Left: Monte Carlo sample RMSE as a function of h_{-2} with the capture geometry specified in Section 2.5, for an estimator applying the correct (R) and incorrect (R_a) measurement covariance. Right: Percentage increase in sample RMSE when R_a is applied rather than R.

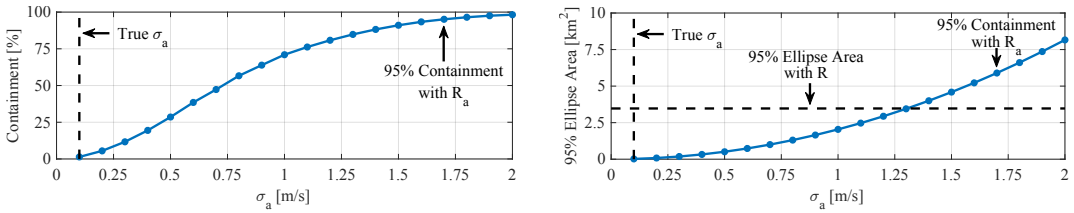


Figure 2.9: Left: Monte Carlo containment percentage when the theoretical 95% error ellipse has been calculated with R_a alone [setting $R_\mathrm{b} = 0$ in (2.19)], for various values of the underlying parameter σ_a. Right: Area of the corresponding theoretical 95% error ellipse as a function of σ_a. The horizontal line shows the area of the 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.

ization of AWGN with $\sigma_\mathrm{a} = 0.1$ 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 km^2.

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 $\sigma_{\mathrm{a}} = 0.1$ m/s, the containment percentage fell to 1.38%. Fig. 2.9 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. 2.9 also shows the area of the theoretical 95% error ellipse for various values of σ_{a}. The area of the 95% error ellipse for $\sigma_{\mathrm{a}} = 1.7$ m/s is 5.90 km^2, 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.

The spoofer's oscillator quality is typically unknown, which makes the selection of the estimator's modeled h_{-2} a design parameter. The range of plausible h_{-2} is likely limited to TCXO quality or better, otherwise, the transmitted spoofing signals are easily detectable by victims. A multi-model approach can be taken where the estimator of each model assumes a different h_{-2}. The convergence of the state estimate can be tested through a goodness of fit test on the weighted sum of squared errors as presented by [12].

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.

2.7 Controlling Spoofing Detection While Degrading Geolocation Accuracy

Researchers have developed formidable defenses against spoofing based on receiver clock state monitoring [69, 72, 78]. A would-be spoofer has little flexibility to meddle with the spoofed clock drift $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ if intending to avoid detection by such defenses. It follows that a stealthy spoofer is scarsely 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 $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ with the aim of either (1) inflating victim receivers' timing error, or (2) confounding geolocation based on this chapter'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.

2.7.1 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 $k \in \mathcal{K} = \{1, 2, ..., K\}$ uniformly sampled navigation epochs. At the kth epoch, the distribution of a GNSS receiver's measured clock drift $\delta\dot{t}_{\mathrm{u}}$ is modeled as

$$c\delta\dot{t}_{\mathrm{u}}[k] \sim \mathcal{N}\big(c\delta\dot{t}_{\mathrm{u}}[k-1],\, \sigma_{\mathrm{u}}^2\big) \quad \text{where} \quad \sigma_{\mathrm{u}}^2 = \sigma_{\mathrm{m}}^2 + q \tag{2.28}$$

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 $c\delta\dot{t}_{\mathrm{u}}$, which is related to the time between navigation epochs Δt and the GNSS receiver clock parameter h_{-2}^{u} by [14]

$$q = 2\pi^2 h_{-2}^{\mathrm{u}} \Delta t c^2 \tag{2.29}$$

Let

$$\eta_k = \frac{c\delta\dot{t}_{\mathrm{u}}[k] - c\delta\dot{t}_{\mathrm{u}}[k-1]}{\sigma_{\mathrm{u}}} \sim \mathcal{N}(0,1) \tag{2.30}$$

be the normalized increment in measured receiver clock drift at the kth 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 $\delta\dot{t}_{\tilde{\mathrm{r}}}[k]$, the spoofed clock drift increment, while assuming that the spoofer's transmitter clock drift $\delta\dot{t}_{\mathrm{t}}(t) = 0$, 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

$$\mu_k = \frac{c\delta\dot{t}_{\tilde{\mathrm{r}}}[k] - c\delta\dot{t}_{\tilde{\mathrm{r}}}[k-1]}{\sigma_{\mathrm{u}}} \tag{2.31}$$

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

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

$$H_0 : \theta_k = \eta_k \ , \quad k \in \mathcal{K} \tag{2.32}$$

$$H_1 : \theta_k = \eta_k + \mu_k \ , \quad \mu_k \neq 0, \quad k \in \mathcal{K} \tag{2.33}$$

Under H_1, the value of μ_k is unknown to the receiver and belongs to the set $(-\infty, 0) \cup (0, \infty)$. A uniformly most powerful test does not exist for this hypothesis test because the critical regions corresponding to $\mu_k < 0$ and $\mu_k > 0$ are different [112]. Instead, a locally most powerful (LMP) test is applied. The LMP design problem is nearly the same as the Neyman-Pearson design problem, so that the probability of detection is maximized while maintaining a fixed probability of false alarm P_{F}. For a single epoch, the detection statistic $\Lambda^*(\theta_k)$ is

$$\Lambda^*(\theta_k) = \theta_k^2 \tag{2.34}$$

and has the following distributions under H_0 and H_1

$$H_0 : \Lambda^*(\theta_k) \sim \chi_1^2 \tag{2.35}$$

$$H_1 : \Lambda^*(\theta_k) \sim \chi_1^2(\lambda), \quad \lambda = \mu_k^2 \tag{2.36}$$

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 $\boldsymbol{\theta} = [\theta_1, \theta_2, ..., \theta_K]^{\mathsf{T}} \in \mathbb{R}^K$ and $\boldsymbol{\mu} = [\mu_1, \mu_2, ..., \mu_K]^{\mathsf{T}} \in \mathbb{R}^K$. The joint test statistic then becomes

$$\Lambda^*(\boldsymbol{\theta}) = \sum_{k \in \mathcal{K}} \theta_i^2 = \boldsymbol{\theta}^{\mathsf{T}} \boldsymbol{\theta} \tag{2.37}$$

with the following distributions under H_0 and H_1:

$$H_0 : \Lambda^*(\boldsymbol{\theta}) \sim \chi^2_K \tag{2.38}$$

$$H_1 : \Lambda^*(\boldsymbol{\theta}) \sim \chi^2_K(\lambda), \quad \lambda = \sum_{k \in \mathcal{K}} \mu_i^2 = \boldsymbol{\mu}^{\mathsf{T}} \boldsymbol{\mu} \tag{2.39}$$

An optimal-decision constant false alarm rate threshold ν^* for P_{F} can be calculated from

$$P_{\mathrm{F}} = P(\Lambda^*(\boldsymbol{\theta}) > \nu^* | H_0) = 1 - F(\nu*; K) \tag{2.40}$$

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

$$P_{\mathrm{D}}(\boldsymbol{\mu}) = P(\Lambda^*(\boldsymbol{\theta}) > \nu^* | H_1) = 1 - F(\nu*; K, \lambda) \tag{2.41}$$

$$= Q_{K/2}\left(\sqrt{\lambda}, \sqrt{\nu^*}\right) \tag{2.42}$$

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

$$\Lambda^*(\boldsymbol{\theta}) \underset{H_0}{\overset{H_1}{\gtrless}} \nu^* \tag{2.43}$$

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

2.7.2 Expression for Geolocation Error

One of the assumptions made when developing the estimator presented in Section 2.4 was that $\delta\dot{t}_{\tilde{\mathrm{r}}}$ is constant. If instead $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ 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 $c\delta\dot{t}_{\tilde{\mathrm{r}}}[i]$ for all $i \in \mathcal{I}$. Then at the ith measurement epoch, $c\hat{\gamma}[i] = c\gamma[i] + \epsilon[i]$. Let the vector of measurement perturbations over the capture interval be represented as $\boldsymbol{\epsilon} = [\epsilon[1],\ \epsilon[2],\ ...,\ \epsilon[I]]^{\mathsf{T}} \in \mathbb{R}^I$, and let $\tilde{\boldsymbol{x}} = [\tilde{e}, \tilde{n}, \tilde{b}]^{\mathsf{T}}$ 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 $\tilde{\boldsymbol{x}}$. The error $\tilde{\boldsymbol{x}}$ can be calculated as

$$\tilde{\boldsymbol{x}} = \left(\tilde{H}^{\mathsf{T}} R^{-1} \tilde{H}\right)^{-1} \tilde{H}^{\mathsf{T}} R^{-1} \boldsymbol{\epsilon} = B\boldsymbol{\epsilon} \tag{2.44}$$

The horizontal position error vector $\boldsymbol{e}_{\mathrm{h}}$ is defined as

$$\boldsymbol{e}_{\mathrm{h}} = [\tilde{e}, \tilde{n}]^{\mathsf{T}} \tag{2.45}$$

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

$$A = \tilde{B}^{\mathsf{T}} \tilde{B} \tag{2.46}$$

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

$$e_{\mathrm{h}} = \sqrt{\boldsymbol{e}_{\mathrm{h}}^{\mathsf{T}} \boldsymbol{e}_{\mathrm{h}}} = \sqrt{\boldsymbol{\epsilon}^{\mathsf{T}} A \boldsymbol{\epsilon}} \tag{2.47}$$

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

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 chapter'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

$$\boldsymbol{\epsilon}^* = \underset{\|\boldsymbol{\epsilon}\| \leq \zeta}{\operatorname{argmax}} \, \boldsymbol{\epsilon}^{\mathsf{T}} A \boldsymbol{\epsilon} \tag{2.48}$$

To solve this problem, A is factorized as $A = QDQ^{\mathsf{T}}$, where Q is orthogonal and $D = \operatorname{diag}(d_1, d_2, ..., d_I)$ 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 $d_1 \geq d_2 \geq ... \geq d_I$. Then

$$\boldsymbol{\epsilon}^{\mathsf{T}} A \boldsymbol{\epsilon} = \boldsymbol{\epsilon}^{\mathsf{T}} Q D Q^{\mathsf{T}} \boldsymbol{\epsilon} = \boldsymbol{y}^{\mathsf{T}} D \boldsymbol{y} \tag{2.49}$$

where $\|\boldsymbol{\epsilon}\| = \|Q^{\mathsf{T}} \boldsymbol{\epsilon}\| = \|\boldsymbol{y}\|$. The value of $\boldsymbol{y}$ that respects $\|\boldsymbol{y}\| \leq \zeta$ and maximizes $\boldsymbol{y}^{\mathsf{T}} D \boldsymbol{y}$ is given by $\boldsymbol{y}^* = [\zeta, 0, 0, ..., 0]^{\mathsf{T}}$. 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

$$\boldsymbol{\epsilon}^* = \pm Q \boldsymbol{y}^* = \pm \zeta \mathbf{v}^* \tag{2.50}$$

2.7.3 Jointly Optimized Spoofer Clock Drift Selection

Now that an optimal spoofing detector based on receiver clock drift has been presented, and a perturbation $\boldsymbol{\epsilon}^*$ that maximizes horizontal geolocation error subject to the constraint $\|\boldsymbol{\epsilon}^*\| < \zeta$ has been defined, a spoofer can develop an attack configuration for $c\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$ 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 $c\boldsymbol{\delta\dot{t}}_{\tilde{\mathrm{r}}} = c\big[\delta\dot{t}_{\tilde{\mathrm{r}}}[1], \delta\dot{t}_{\tilde{\mathrm{r}}}[2], ..., \delta\dot{t}_{\tilde{\mathrm{r}}}[I]\big]^{\mathsf{T}} \in \mathbb{R}^I$ represent the spoofer's discretized time-varying attack configuration for $\delta\dot{t}_{\tilde{\mathrm{r}}}(t)$. Suppose the spoofer sets $c\boldsymbol{\delta\dot{t}}_{\tilde{\mathrm{r}}} = \boldsymbol{\epsilon}^*$. Then the vector of spoofed clock drift increments over $K = I - 1$ navigation epochs is equivalent to

$$\boldsymbol{\mu} = \frac{\zeta}{\sigma_{\mathrm{u}}} C\mathbf{v}^* \in \mathbb{R}^K \quad \text{where} \quad C = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix} \in \mathbb{R}^{K \times I} \tag{2.51}$$

The only task remaining for the spoofer is to determine the value of ζ so that $\boldsymbol{\epsilon}^*$ can be scaled appropriately. Suppose the spoofer is willing to allow a detection probability $\bar{P}_{\mathrm{D}}$ for the detection test in (2.43). Based on the parameters σ_{u}, I, and P_{F}, the parameter ζ can be chosen to maintain an expected probability of detection $\bar{P}_{\mathrm{D}}$. Given the functional form of the probability of detection in (2.42), ζ must satisfy the equation

$$Q_{K/2}\left(\frac{\zeta}{\sigma_{\mathrm{u}}}\|C\mathbf{v}^*\|, \sqrt{\nu^*}\right) = \bar{P}_{\mathrm{D}} \tag{2.52}$$

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

$$c\boldsymbol{\delta}\dot{\boldsymbol{t}}_{\tilde{\mathrm{r}}}^{*} = \pm\,\zeta(\mathbf{v}^{*} - \mathbf{1}\mathrm{v}^{*}[1]) \tag{2.53}$$

where $\mathbf{1}$ is the appropriately sized vector of all ones and $\mathrm{v}^{*}[1]$ is the first element of $\mathbf{v}^{*}$. Note that subtracting $\mathbf{1}\mathrm{v}^{*}[1]$ ensures $c\boldsymbol{\delta}\dot{\boldsymbol{t}}_{\tilde{\mathrm{r}}}[1] = 0$, consistent with initialization of the spoofing attack. This subtraction does not change the optimization 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 $\boldsymbol{\delta}\dot{\boldsymbol{t}}_{\tilde{\mathrm{r}}}^{*}$ to maximally degrade geolocation by a LEO-based receiver capturing its signals over 21 seconds with the geometry shown in Fig. 2.3. Further suppose the LEO-based receiver computes measurements at 1 Hz, so that $I = 21$, and sets R with $\sigma_{\mathrm{a}} = 0.1$ m/s and σ_v consistent with a TCXO. The attack trajectory $\mathbf{v}^{*} - \mathbf{1}\mathrm{v}^{*}[1]$ that maximizes horizontal geolocation error is shown in Fig. 2.10. 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 $\sigma_{\mathrm{m}} = 0.05$ m/s. Shown in Fig. 2.11 is the maximum horizontal geolocation error given a triad of $\bar{P}_{\mathrm{D}}$, P_{F}, and affected

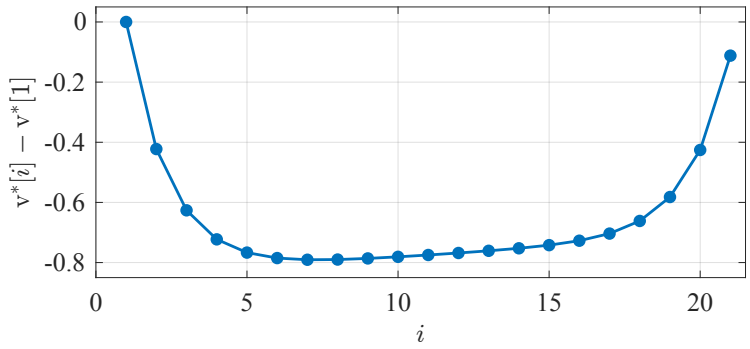


Figure 2.10: The attack trajectory $\mathrm{v}[i]^* - \mathrm{v}^*[1]$ that maximizes e_h for the LEO-based receiver geometry shown in Fig. 2.3.

receiver clock quality. For example, if the spoofer accepts a detection rate of $\bar{P}_\mathrm{D} = 0.5$ by receivers equipped with a TCXO having their spoofing detector set with $P_\mathrm{F} = 10^{-3}$, the maximum e_h due to $c\boldsymbol{\delta \dot{t}}^*_{\tilde{\mathrm{r}}}$ 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 chapter'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.

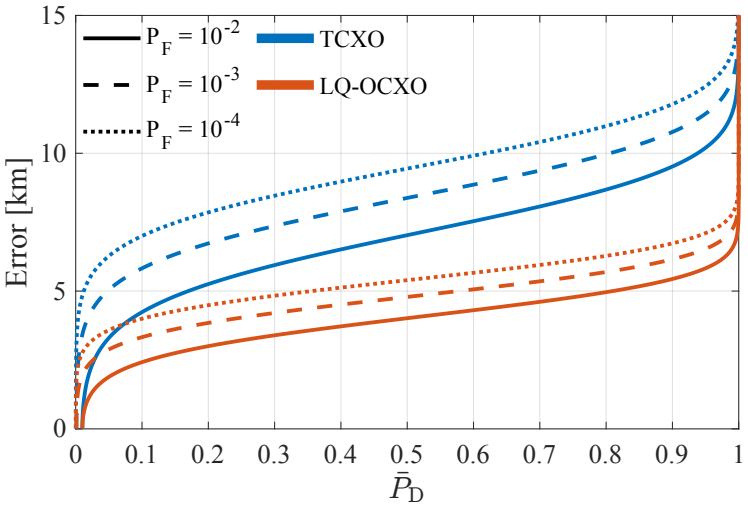


Figure 2.11: Worst-case geolocation error for a spoofer that optimizes $c\boldsymbol{\delta \dot{t}}_{\tilde{\mathrm{r}}}$ for receivers performing 1-Hz spoofing detection tests with $\sigma_{\mathrm{m}} = 0.05$ m/s and for the LEO-based receiver geometry shown in Fig. 2.3. The geolocation error is shown over a range of $\bar{P}_{\mathrm{D}}$ for two representative victim receiver clock quality levels and three representative values of P_{F}.

2.8 Conclusion

This chapter 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.

Chapter 3

Dual-Satellite Geolocation of Terrestrial GNSS Interference from Low Earth Orbit

Abstract

This chapter explores two-step and direct geolocation of terrestrial Global Navigation Satellite System (GNSS) jammers from Low Earth Orbit (LEO). Within the past decade, there has been a sharp increase in GNSS outages due to deliberate GNSS jamming. Receivers in LEO are uniquely situated to detect, classify, and geolocate terrestrial GNSS jammers. The conventional two-step geolocation method first estimates the differential delay and differential Doppler, then uses a time history of these to estimate the transmitter location. By contrast, direct geolocation is a single-step search over a geographical grid that enables estimation of the transmitter location directly from the observed signals. Signals from narrowband, matched-code, and continuous-wave jammers recently captured in the GNSS frequency bands by two time-synchronized LEO receivers over the Eastern Mediterranean are analyzed and the emitters geolocated. It is demonstrated that the direct approach is effective even for low signal-to-noise ratio interference signals based on short captures with multiple emitters. Moreover, the direct approach enables geolocation of multiple emitters with cyclostationary signals (e.g., chirp

jammers), whereas the two-step method struggles in such cases to associate emitters with their corresponding structures in differential delay and Doppler space.

3.1 Introduction

Global Navigation Satellite Systems (GNSS) such as GPS provide meter-accurate positioning while offering global accessibility and all-weather, radio-silent operation. However, GNSS is fragile: its service is easily denied by jammers or deceived by spoofers [28, 62, 114]. GNSS signals are especially vulnerable to jamming because they are extremely weak: near the surface of Earth, they have no more flux density than light received from a 50 W bulb at a distance of 2000 km [65]. Furthermore, GNSS jammers are easily accessible and low cost, threatening GNSS-reliant systems [13, 29]. Without proper countermeasures, victim GNSS receivers can be rendered useless.

The civilian maritime and airline industries frequently encounter GNSS jamming and spoofing. Corrupted Automatic Identification System (AIS) and Automatic Dependent Surveillance-Broadcast (ADS-B) messages from surface vessels and aircraft are often reported. Irregularities in AIS and ADS-B reports are often indicative of GNSS interference. Geolocation of GNSS jammers with ADS-B data is possible, but only coarse jammer position estimates are achievable [33, 88].

Recently, the German Aerospace Center (DLR) performed a data collection flight over the Eastern Mediterranean to study the behavior of regular

avionics and aviation-grade GNSS receivers under jamming conditions [108]. The DLR also conducted an international maritime measurement campaign to detect GNSS interference [89]. In both studies, the recorded data showed evidence of high-power GNSS jammers, including a chirp jammer centered at the GPS L1 frequency in the Eastern Mediterranean.

GNSS jamming waveforms can take several forms, including tones, chirps, pulses, and matched-code interference [94]. A first step to developing situational awareness and eliminating GNSS interference is geolocating the emitters involved. It was shown that a network of ground-based receivers could track and geolocate chirp-style jamming signals in [95], and matched-code jamming signals in [11]. The more general case of localizing an emitter transmitting an arbitrary wideband signal with terrestrial and airborne receivers has been extensively studied [5, 101, 107, 138]. However, because the receivers were either at fixed locations or tactically deployed in the nearby airspace, only emitters in the immediate area could be geolocated. There remains a need for global, persistent, low-latency, and accurate GNSS interference detection and localization.

Receivers based in Low Earth Orbit (LEO) are a proven asset for detecting, classifying, and geolocating terrestrial GNSS interference [22, 25, 86, 100]. 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. Moreover, LEO satellites' stand-off distance from terrestrial interference sources typically permits tracking authentic GNSS

signals, enabling precise time-tagged data captures from time-synchronized LEO-based receivers and precise orbit determination. LEO constellations with distributed time-synchronized receivers can provide unprecedented emitter geolocation. Unsurprisingly, several commercial enterprises have seized the opportunity to provide spectrum monitoring and emitter geolocation as a service (e.g., Spire Global and Hawkeye360).

Accurate single-satellite-based emitter geolocation is possible from Doppler measurements alone, provided the emitter's carrier can be extracted [25, 37, 100]. Performance bounds and error characterization for Doppler-based single-satellite geolocation are presented in [38] [36]. However, accurately locating emitters with arbitrary waveforms using a single LEO receiver is impossible in general: if the signal's carrier cannot be tracked, only coarse received-signal-strength techniques can be applied for geolocation.

On the other hand, geolocation of emitters producing arbitrary wideband signals is possible and has been extensively studied [61, 101, 121]. Multiple time-synchronized receivers can exploit time- and frequency-difference of arrival (T/FDOA) measurements to estimate the emitter location. Geolocation based on T/FDOA is typically a two-step process. First, a time series of T/FDOA measurements is produced by correlating captured signals against another. Second, the time series is fed to a nonlinear estimation algorithm to geolocate the source. Two-step T/FDOA has been previously applied for terrestrial emitter localization from geostationary orbit [58].

One weakness of two-step geolocation is that it ignores the constraint

that all measurements must be consistent with a single position in the case of a stationary emitter, or a single trajectory in the case of a moving emitter [134]. This means that the T/FDOA measurements obtained in the first step are not guaranteed to intersect at a single location or along a single trajectory. A second weakness of the two-step approach is that interference signals exhibiting cyclostationarity—such as chirp jammers—give rise to structures in the T/FDOA measurement domain that make it harder to track individual emitters. Identification and tracking becomes especially challenging when there are multiple cyclostationary emitters with overlapping frequency content and a wide range of received power, in which case the T/FDOA measurement domain becomes highly structured with features ambiguously related to the emitters involved.

Another multi-receiver technique is direct geolocation, which is a single-step search over a geographical grid to estimate a transmitter's location directly from the observed signals [9, 87, 119, 131, 134]. In direct geolocation, the TDOA and FDOA are directly parameterized for a single geographical point, given knowledge of the receivers' position, velocity, and clock states. Direct geolocation outperforms the two-step method in low signal-to-noise ratio (SNR) environments and in short-data-capture scenarios, making it ideal for LEO-based geolocation. Furthermore, as will be shown in this chapter, direct geolocation is better suited for processing captures with cyclostationary signals from multiple emitters, because rather than searching in the T/FDOA measurement domain cluttered by overlapping structures, it searches in the

position domain, where individual emitters are separated by their physical distance, irrespective of any time correlation in their signals.

This chapter makes three primary contributions. First, it develops and presents algorithms—from raw samples to final emitter position estimate—for two-step and direct geolocation. Secondly, it investigates how the time-domain properties of the transmitted signal manifest in the complex ambiguity function (CAF) of the two-step approach, and explores the relationship between the CAF domain and the position domain. Finally, this chapter demonstrates two-step and direct geolocation on raw intermediate frequency (IF) samples recorded from Spire Global's LEO constellation. For the first time in the open literature, real-world GNSS narrowband, matched-code, and chirp jamming signals captured by two time-synchronized LEO receivers are characterized and their emitters geolocated.

3.2 Measurement Model

Consider the following model for a signal transmitted by a terrestrial stationary emitter where $s(t)$ is the baseband signal, f_c is the center frequency, and the bandwidth of $s(t)$ is $B \ll f_c$:

$$x(t) = s(t)e^{j2\pi f_c t} \tag{3.1}$$

This analytic signal model is valid when the complex envelope varies slowly with respect to the center frequency, which applies to the GPS L1 (1575.42 MHz) and L2 (1227.6 MHz) frequencies. When a moving receiver receives this

signal, there is a time delay τ due to the propagation time, as well as a Doppler shift f due to the relative motion between emitter and receiver. The line-of-sight velocity, receiver clock drift, and transmitter clock drift are components of f. Over short intervals, τ and f can be modeled as constants. The received complex signal at the ith receiver after downconversion by a known carrier frequency f_c at time t can be modeled as

$$y_i(t) = s(t - \tau_i)e^{j2\pi f_i t} + w_i(t) \tag{3.2}$$

where $w_i(t)$ corresponds to AWGN.

3.2.1 TDOA and FDOA Measurement Model

Let $\boldsymbol{p}_i(t)$ and $\boldsymbol{v}_i(t)$ denote the position and velocity vector for the ith receiver, and $\boldsymbol{p}$ denote the emitter's position vector, all in a common rectangular coordinate frame. In this model, the emitter is assumed to be stationary. Ionospheric and tropospheric delay are ignored because they are minor and will be substantially canceled in the differencing to follow. The time of arrival (TOA), in seconds, at the ith receiver of the signal transmitted from the emitter at time t is modeled as

$$\begin{aligned} \tau_i(t) &= \frac{1}{c}\sqrt{(\boldsymbol{p}_i(t) - \boldsymbol{p})^{\mathsf{T}}(\boldsymbol{p}_i(t) - \boldsymbol{p})} \\ &= \frac{1}{c}\sqrt{\boldsymbol{r}_i^{\mathsf{T}}(t)\boldsymbol{r}_i(t)} \end{aligned} \tag{3.3}$$

where the range vector $\boldsymbol{r}_i(t)$ between the emitter and the ith receiver is

$$\boldsymbol{r}_i(t) = \boldsymbol{p}_i(t) - \boldsymbol{p} \tag{3.4}$$

The range between the emitter and the ith receiver is related to τ_i and $\boldsymbol{r}_i$ by

$$\rho_i(t) = c\tau_i(t) = \sqrt{\boldsymbol{r}_i^{\mathsf{T}}(t)\boldsymbol{r}_i(t)} \tag{3.5}$$

Finally, the unit vector from emitter position to the ith receiver position is defined as

$$\hat{\boldsymbol{r}}_i(t) = \boldsymbol{r}_i(t)/\rho_i(t) \tag{3.6}$$

Different LEO-based receivers will receive the same signal at different times due to the differing geometry between receivers. Assuming the receivers are synchronized to GPS time, the time difference of arrival (TDOA) of the same signal between the first and second receiver is defined as

$$\Delta\tau(t) = \tau_2(t) - \tau_1(t) \tag{3.7}$$

which can be converted to a range difference by multiplying by the speed of light.

The frequency of arrival (FOA) measurement is synonymous with the received Doppler of a signal. For a stationary emitter, the FOA on the ith moving receiver is composed of three components: (1) the range-rate between the emitter and receiver $\hat{\boldsymbol{r}}_i^{\mathsf{T}}(t)\boldsymbol{v}_i(t)$, (2) the clock offset rate of the receiver $\delta\dot{t}_i(t)$, and (3) the clock offset rate of the emitter $\delta\dot{t}_{\mathrm{e}}(t)$. The FOA at the ith receiver is modeled as

$$f_i(t) = -\frac{1}{\lambda}\hat{\boldsymbol{r}}_i^{\mathsf{T}}(t)\boldsymbol{v}_i(t) - \frac{c}{\lambda}\big\{\delta\dot{t}_i(t) - \delta\dot{t}_{\mathrm{e}}(t)\big[1 - \delta\dot{t}_i(t)\big]\big\} \tag{3.8}$$

Different LEO receivers will also receive the same signal at different frequencies due to the differing instantaneous range-rates and receiver clock offset rates. Let $\Delta\tilde{f}(t)$ denote the frequency difference of arrival that includes the receivers' clock offset rates.

$$\begin{aligned}\Delta\tilde{f}(t) &= f_2(t) - f_1(t) \\ &\approx -\frac{1}{\lambda}\big\{\hat{\boldsymbol{r}}_2^{\mathsf{T}}(t)\boldsymbol{v}_2(t) - \hat{\boldsymbol{r}}_1^{\mathsf{T}}(t)\boldsymbol{v}_1(t)\big\} - \frac{c}{\lambda}\big\{\delta\dot{t}_2(t) - \delta\dot{t}_1(t)\big\}\end{aligned} \tag{3.9}$$

An advantageous feature of $\Delta\tilde{f}(t)$ is that the clock offset rate from the emitter gets canceled out in the differencing. The approximation disregards the $\delta\dot{t}_{\mathrm{e}}(t)$ and $\delta\dot{t}_i(t)$ cross terms, as they are negligible. The frequency difference of arrival (FDOA) $\Delta f(t)$ with compensation for the receiver clock offset rates is defined as

$$\begin{aligned}\Delta f(t) &= \Delta\tilde{f}(t) + \underbrace{\frac{c}{\lambda}\big\{\delta\dot{t}_2(t) - \delta\dot{t}_1(t)\big\}}_{\text{known}} \\ &= -\frac{1}{\lambda}\big\{\hat{\boldsymbol{r}}_2^{\mathsf{T}}(t)\boldsymbol{v}_2(t) - \hat{\boldsymbol{r}}_1^{\mathsf{T}}(t)\boldsymbol{v}_1(t)\big\}\end{aligned} \tag{3.10}$$

If the front end performing the raw capture and the onboard GNSS receiver are driven by the same oscillator, then the receiver clock drift can be compensated and the sample streams can be synchronized. The onboard GNSS navigation solution of each receiver outputs an estimate of the receiver clock offset δt_i and clock offset rate $\delta\dot{t}_i$ at each GNSS measurement epoch. This means that the clock offset rate at each receiver is known and can be compensated for, making the FDOA measurement the difference between range-rates scaled by the negative reciprocal of wavelength.

3.2.2 The Generalized Cross-Correlation Function

Assume that all receivers are synchronized to GPS time and clock errors have been compensated for. The generalized cross-correlation function (GCCF) [9] for a pair of received complex baseband signals $y_1(t)$ and $y_2(t)$ is

$$S^{\mathrm{GCCF}}(y_1(t), y_2(t), \tau_1(t), \tau_2(t)) = \int_0^T y_1(t - \tau_1(t)) y_2^*(t - \tau_2(t)) e^{j2\pi f_c(\tau_1(t) - \tau_2(t))} dt \tag{3.11}$$

where T is the integration interval. The more familiar complex ambiguity function (CAF) [124] with constant $\Delta\tau$ and Δf is

$$S^{\mathrm{CAF}}(y_1(t), y_2(t), \Delta\tau, \Delta f) = \int_0^T y_1(t) y_2^*(t + \Delta\tau) e^{j2\pi\Delta f t} dt \tag{3.12}$$

The discrete-time CAF can be efficiently evaluated with the Fast Fourier Transform (FFT), parallelization, and vectorization.

If there is only a single emitter, it is shown in [124, 125] that the $\Delta\tau$ and Δf that maximize the magnitude S^{CAF} are the corresponding maximum likelihood estimates (MLE) of $\Delta\tau$ and Δf between a pair of receivers. The theoretical Cramér-Rao lower bound (CRLB) of these estimates are derived in [124] and are

$$\sigma_{\Delta\tau} = \frac{1}{\beta_{\mathrm{rms}}} \frac{1}{\sqrt{BT\mathrm{SNR}}} \tag{3.13}$$

$$\sigma_{\Delta f} = \frac{1}{T_{\mathrm{rms}}} \frac{1}{\sqrt{BT\mathrm{SNR}}} \tag{3.14}$$

where B is the noise bandwidth at receiver input (assumed same for both receivers), β_{rms} is the rms radian frequency in the received signal spectrum,

T_{rms} is the rms integration time, and SNR is the effective input signal to noise ratio [124]. For a signal with a rectangular spectrum and a constant energy signal with ambiguity processing carried out only over a segment length T, then $\sigma_{\Delta\tau}$ and $\sigma_{\Delta f}$ simplify to

$$\sigma_{\Delta\tau} = \frac{\sqrt{3}}{\pi B}\frac{1}{\sqrt{BT\text{SNR}}} \tag{3.15}$$

$$\sigma_{\Delta f} = \frac{\sqrt{3}}{\pi T}\frac{1}{\sqrt{BT\text{SNR}}} \tag{3.16}$$

Over short observation intervals, the errors introduced by assuming a constant $\Delta\tau$ and Δf are negligible [9]. Maximizing the coherent integration time, T, is desirable as it enhances the post-processing SNR, thereby reducing $\sigma_{\Delta\tau}$ and $\sigma_{\Delta f}$. Notably, the frequency estimate $\sigma_{\Delta f}$ benefits more significantly from extended integration due to its inverse relationship with T. However, the maximum value of T is practically constrained by receiver dynamics and clock instabilities. For LEO platforms utilizing temperature-compensated crystal oscillators (TCXOs), the coherent integration time is typically capped at 40 ms to prevent peak smearing.

The presence of emitters is visualized in the CAF domain as shown in Fig. 3.1. The left plot illustrates a single emitter without cyclostationary properties, manifesting as a unique, unambiguous peak. Conversely, the right plot demonstrates the CAF of a cyclostationary signal, which presents multiple peaks of equivalent magnitude. This phenomenon arises from the periodic structure of the transmitted waveform, which is a common characteristic of

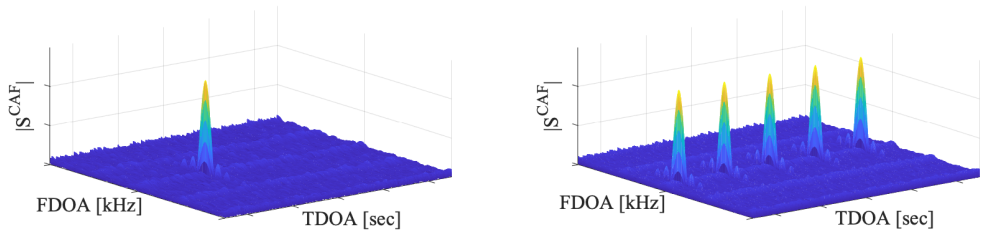


Figure 3.1: On the left is an example CAF for a single wideband emitter without any cyclostationary properties. The $\Delta\tau$ and Δf that correspond to the maximum magnitude are the maximum likelihood estimates. On the right is the CAF of a single wideband emitter with cyclostationary properties. There is an ambiguity because there are peaks with the same magnitude.

interference. Since these peaks represent equivalent local maxima in the likelihood surface, they must all be tracked over time. This introduces a significant combinatorial association problem, especially when there are many peaks close that are closely spaced.

The peak tracking challenge is further exacerbated in multi-emitter environments, particularly when multiple sources exhibit cyclostationary characteristics. The auto-ambiguity terms generated by each emitter's waveform may interfere with each other, leading to biases in the T/FDOA estimate [9]. Moreover, the height of an emitter's peak in the CAF is determined by the SNR. Weaker emitters will have smaller peaks, leading to possible missed detections in the presence of stronger emitters.

3.3 Emitter Geolocation

3.3.1 Two-step Geolocation

In the traditional two-step geolocation approach, a time series of T/FDOA measurements is first obtained by repeated CAF generation and peak tracking. The CAF is computed at each time epoch a T/FDOA measurement pair is desired, for a total of N T/FDOA pairs. For model simplicity, the TDOAs can be converted to range difference in meters by scaling them by the speed of light, and the FDOAs can be converted to range-rate difference in meters per second by scaling by $-\lambda$. In this section, TDOAs and FDOAs are assumed to be converted into m and m/s respectively.

The T/FDOA measurements are beset with errors, modeled as AWGN, arising from electromagnetic interference, atmospheric and relativistic effects, ephemeris errors, and other minor effects. The kth TDOA and FDOA measurements are modeled as

$$z_{\Delta\tau}[k] = c\Delta\tau[k] + w_{\Delta\tau}[k] \tag{3.17}$$

$$z_{\Delta f}[k] = -\lambda\Delta f[k] + w_{\Delta f}[k] \tag{3.18}$$

where $w_{\Delta\tau}$ represents the TDOA AWGN and $w_{\Delta f}$ represents the FDOA AWGN. It is assumed that $w_{\Delta\tau}$ and $w_{\Delta f}$ are independent. Let $\boldsymbol{z}_{\Delta\tau}$ and $\boldsymbol{z}_{\Delta f}$ represent the stacked TDOA and FDOA measurement vectors, respectively. The $2N \times 1$ T/FDOA measurement vector $\boldsymbol{z}$ is represented

$$\boldsymbol{z} = \begin{bmatrix} \boldsymbol{z}_{\Delta\tau}^{\mathsf{T}} & \boldsymbol{z}_{\Delta f}^{\mathsf{T}} \end{bmatrix}^{\mathsf{T}} \tag{3.19}$$

A nonlinear least-squares estimator is used to solve for the the state $\boldsymbol{x}$. In this scenario, the state is comprised only of the transmitter's position $\boldsymbol{x} = \boldsymbol{p}$. The standard weighted nonlinear least-squares cost function is

$$J(\boldsymbol{x}) = \frac{1}{2}(\boldsymbol{z} - \boldsymbol{h}(\boldsymbol{x}))^{\mathsf{T}} R^{-1}(\boldsymbol{z} - \boldsymbol{h}(\boldsymbol{x})) \tag{3.20}$$

where $\boldsymbol{h}(\boldsymbol{x})$ is the $2N \times 1$ nonlinear measurement model function

$$\boldsymbol{h}(\boldsymbol{x}) = \begin{bmatrix} c\Delta\bar{\tau}[1] & \cdots & c\Delta\bar{\tau}[N] & -\lambda\Delta\bar{f}[1] & \cdots & -\lambda\Delta\bar{f}[N] \end{bmatrix}^{\mathsf{T}} \tag{3.21}$$

Here, $\Delta\bar{\tau}[k]$ and $\Delta\bar{f}[k]$ are the expected values of $\Delta\tau[k]$ and $\Delta f[k]$ at the state $\boldsymbol{x}$. R is the $2N \times 2N$ measurement covariance matrix with the variance of the TDOA measurements $\sigma^2_{\Delta\tau}$ along the first N diagonal elements and the variance of the FDOA measurements $\sigma^2_{\Delta f}$ along the second N diagonal elements. The optimal estimate of $\boldsymbol{x}$ minimizes the cost J.

The linearized measurement model H is a $2N \times 3$ matrix that takes the form

$$H = \begin{bmatrix} c\frac{d\Delta\tau[1]}{d\boldsymbol{x}} \\ \vdots \\ c\frac{d\Delta\tau[N]}{d\boldsymbol{x}} \\ -\lambda\frac{d\Delta f[1]}{d\boldsymbol{x}} \\ \vdots \\ -\lambda\frac{d\Delta f[N]}{d\boldsymbol{x}} \end{bmatrix} \tag{3.22}$$

where the Jacobians for the kth TDOA measurement and kth FDOA mea-

surement are

$$c\frac{d\Delta\tau}{d\boldsymbol{x}} = -[\hat{\boldsymbol{r}}_2 - \hat{\boldsymbol{r}}_1]^{\mathsf{T}} \tag{3.23}$$

$$-\lambda\frac{d\Delta f}{d\boldsymbol{x}} = \boldsymbol{v}_2^{\mathsf{T}}\frac{\left(\hat{\boldsymbol{r}}_2\hat{\boldsymbol{r}}_2^{\mathsf{T}} - \mathbb{I}\right)}{\rho_2} - \boldsymbol{v}_1^{\mathsf{T}}\frac{\left(\hat{\boldsymbol{r}}_1\hat{\boldsymbol{r}}_1^{\mathsf{T}} - \mathbb{I}\right)}{\rho_1} \tag{3.24}$$

where $\mathbb{I}$ is the identity matrix and the indexing being dropped for notational clarity.

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

$$z[2N+1] = h_{\text{alt}}(\boldsymbol{x}) + w_{\text{alt}} \tag{3.25}$$

where the pseudo-measurement error $w_{\text{alt}} \sim \mathcal{N}(0,\, \sigma_{\text{alt}}^2)$ is assumed to be independent of those for $\boldsymbol{z}$. The pseudo-measurement's 1×3 Jacobian is

$$H_{\text{alt}} = \left[\cos(\phi_{\text{lat}})\cos(\lambda_{\text{lon}}),\ \cos(\phi_{\text{lat}})\sin(\lambda_{\text{lon}}),\ \sin(\phi_{\text{lat}})\right] \tag{3.26}$$

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

One of the most common ways to solve the standard nonlinear least-squares problem is with the Gauss-Newton method, which starts with an initial guess for the state $\boldsymbol{x}$, linearizes the measurement model function $\boldsymbol{h}(\boldsymbol{x})$ around

the current best guess, and solves the linear least-squares problem. An increment to the current best estimate $\Delta \boldsymbol{x}$ is computed, and then the previous estimate is replaced by the new best estimate $\boldsymbol{x} + \Delta \boldsymbol{x}$ and iterates again until convergence is achieved. Finally, the estimation error's CRLB can be approximated as

$$P_{\boldsymbol{xx}} = \left(H^{\mathsf{T}} R^{-1} H\right)^{-1} \tag{3.27}$$

This two-step T/FDOA geolocation model can also be reduced to a TDOA-only solution or a FDOA-only solution. Those solutions can be used as a reasonableness test for the combined T/FDOA solution and to quantify the accuracy of the TDOA and FDOA measurements. Additionally, the MLE of $\sigma^2_{\Delta\tau}$ and $\sigma^2_{\Delta f}$ can be estimated from the post-fit measurement residuals of the TDOA-only and FDOA-only solutions, respectively. Finally, if the user does not have a time history of the clock offset rate at each receiver, it can be modeled as a constant over short captures (e.g., less than 60 seconds [100]). The difference between the receiver clock offset rate can be added to the state as a bias term and can be estimated.

3.3.2 Direct Geolocation

Direct geolocation is a single-step search over a geographical grid that enables estimation of the transmitter position directly from correlating raw observed signals without the need for intermediate T/FDOA measurements [9, 134]. The correlation magnitude is maximized directly by parameterizing

the delay and Doppler time histories in terms of the assumed emitter's position, as well as the known receivers' position, velocity, and clock time histories. Direct geolocation has been shown to be maximum likelihood in [9, 134] and to outperform the two-step approach in low SNR regimes and in cases limited to short captures [134].

A three-dimensional grid of candidate emitter positions must be first designated. One approach is to create a uniform grid in latitude and longitude. Then, for each latitude and longitude pair, the altitude can be retrieved from a terrain model. This constrains the emitter position to the relative terrain on the surface of the Earth. These points can then be transformed to the ECEF coordinate frame. Given that the emitter's position is assumed, and that the time history of receivers' position, velocity, and clock offset rate are known, the time history of TOAs and FOAs of a transmitted signal at each receiver can be computed with (3.3) and (3.8). It follows that a time history of TDOA and FDOA can be computed using (3.7) and (3.10). Over short intervals when the $\Delta\tau$ and Δf are assumed to be constant, the position domain (PD) correlation value S^{PD} of $\boldsymbol{p}$ is defined as

$$S^{\mathrm{PD}}(y_1(t), y_2(t), \boldsymbol{p}) = \int_0^T y_1(t) y_2^*(t + \Delta\tau(\boldsymbol{p})) e^{j2\pi\Delta f(\boldsymbol{p})t} dt \tag{3.28}$$

This is computed for every candidate position. The MLE of the emitter's position is the $\boldsymbol{p}$ that maximizes $|S^{\mathrm{PD}}|$. An example position domain correlation grid is shown in Fig. 3.2.

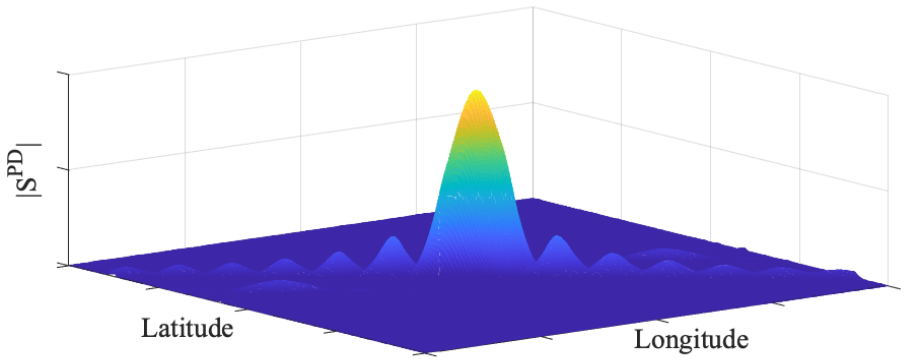


Figure 3.2: Example position-domain correlation grid from direct geolocation for a single wideband emitter without any cyclostationary properties. The maximum magnitude corresponds to the MLE of emitter's position.

Furthermore, the position-domain correlation can be extended to non-constant $\Delta\tau$ and Δf because the satellites' position and velocity are known, and the position of the emitter is being assumed. This allows for the calculation of the time-varying $\Delta\tau(t)$ and $\Delta f(t)$. In the time-varying $\Delta\tau(t)$ and $\Delta f(t)$ scenario, the position domain correlation is defined

$$S^{\mathrm{PD}}(y_1(t), y_2(t), \boldsymbol{p}) = \int_0^T y_1(t) y_2^*(t + \Delta\tau(\boldsymbol{p}, t)) e^{j2\pi\Delta\theta(\boldsymbol{p},t)} dt \tag{3.29}$$

where $\Delta\theta(\boldsymbol{p}, t)$ corresponds to the phase-difference time history [9]. This is similar to the GCCF in (3.11). This is one of the most powerful aspects of direct geolocation, because the coherent integration interval is no longer limited by constant $\Delta\tau$ and Δf, which is the case in the two-step approach. Longer coherent integration results in higher SNR and more accurate frequency resolution, which are both desirable properties. As the integration time is

increased, the peak at the true emitter position becomes sharper and more pronounced. However, one drawback of longer coherent integration intervals is that a finer grid is required. If the grid is too coarse, the peak will slip through the cracks leading to a missed detection.

When there are multiple time epochs so that multiple correlation intervals can be performed, and SNR is assumed to be constant, the MLE of the emitter's becomes

$$\hat{\boldsymbol{p}} = \underset{\boldsymbol{p}}{\operatorname{argmax}} \sum_{k} |S_k^{\mathrm{PD}}(\boldsymbol{p})| \tag{3.30}$$

The CRLB is derived in [134] and shown to be identical to (3.27) when the values of (3.13) and (3.14) are used. It has been shown in [134] that the direct approach more closely follows the CRLB in low SNR regimes compared to the two-step approach.

3.3.3 Connection between the CAF Domain and Position Domain

One can immediately observe the similarities between the CAF and position domain correlation. Over short observation intervals where $\Delta\tau$ and Δf are assumed to be constant, the two functions can be viewed as homeomorphic mappings of one another. The statistical properties and structure of a cross-correlated waveform are well-defined within the CAF domain. For instance, in the case of a pseudorandom binary code as shown in Fig. 3.3. The CAF exhibits a characteristic triangle function at the optimal Δf, with the width determined by the signal's bandwidth. Conversely, at the optimal $\Delta\tau$, the

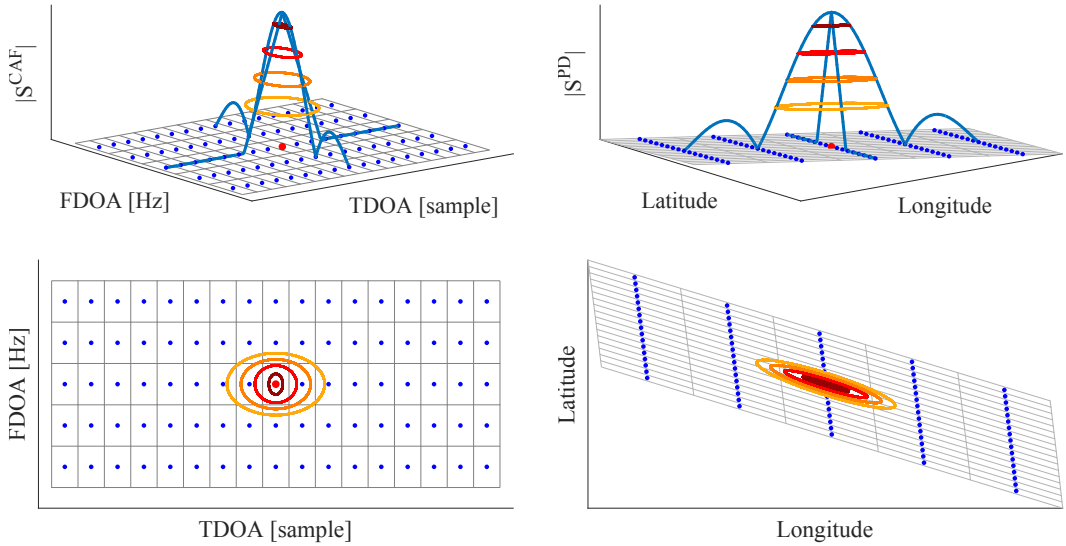


Figure 3.3: Shown here is an example mapping of $|S^{\mathrm{CAF}}|$ and $|S^{\mathrm{PD}}|$. The position-domain correlation is effectively a homeomorphic mapping of the CAF domain. The blue dots represent the nearest sample and frequency bin pair.

CAF manifests as a sinc function, where the main lobe width is proportional to $1/T$. The CAF grid resolution is typically defined by the sampling interval in the TDOA dimension and by $1/T$ in the FDOA dimension. This is shown in Fig. 3.3, where each blue dot is at the center of a T/FDOA cell. Each cell is typically modeled as statistically independent, while the errors within a cell being correlated.

When performing correlation directly in the position domain, the resulting surface represents a geometrically warped CAF. As illustrated in Fig. 3.3, each CAF cell is mapped into the position domain through a transformation dictated by the platform-emitter geometry and the associated Jacobian of the

manifold. Each cell in the position domain represents a different T/FDOA cell. While the $\Delta\tau$ and Δf resolutions are fixed by bandwidth and coherent integration length, their spatial footprint varies spatially. As the coherent integration interval T increases, the volume of the mainlobe in the ambiguity surface contracts, leading to a finer spatial resolution. Consequently, the resolution of direct geolocation can be analytically determined by projecting the CAF into the position domain.

A simulation study was conducted to demonstrate the robustness of direct geolocation when applied to signals with periodic cyclostationary properties. In this scenario, a chirp jamming signal is transmitted that repeats every 40 microseconds. Fig. 3.4 illustrates $|S^{\mathrm{CAF}}|$ and $|S^{\mathrm{PD}}|$ across three distinct observation intervals, alongside the corresponding receiver-emitter geometry. The red markings denote the geometric manifold mapping between the CAF and position domains, while the white diamonds indicate the true $\Delta\tau$ and Δf in the CAF domain and the true position in the position domain. The true $\Delta\tau$ and Δf can span a large region of the CAF because of the vast spatial search space due the LEO-based receivers. At any single epoch, the periodic nature of the waveform generates multiple ambiguous peaks in both the CAF and position domains. However, because the LEO receiver geometry evolves between epochs, the spatial projection of these "false" peaks is non-stationary. While the ambiguity peaks shift in geographic space as the platforms move, the peak corresponding to the true emitter position remains spatially consistent.

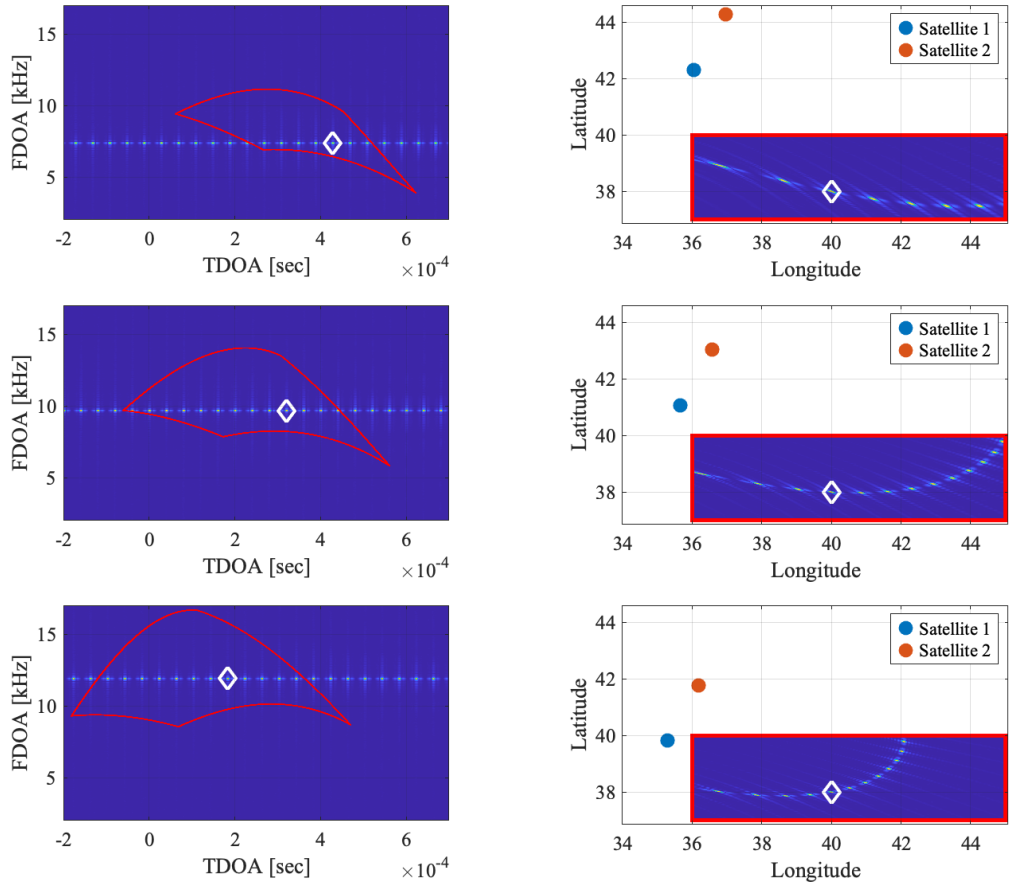


Figure 3.4: Shown here are $|S^{\mathrm{CAF}}|$ and $|S^{\mathrm{PD}}|$ from the simulation across three distinct observation intervals for a cyclostationary chirp jammer. The red markings denote the geometric manifold mapping between the CAF and PD, while the white diamonds indicate the true $\Delta\tau$ and Δf and the true position.

Fig. 3.4 displays the result of noncoherent integration of the position-domain likelihood, $|S^{\mathrm{PD}}|$. Through this process, likelihood accumulates at the true source location, whereas the shifting ambiguities are suppressed. This demonstrates that direct geolocation, by leveraging the dynamic LEO geometry, provides a powerful mechanism for resolving cyclostationary interference.

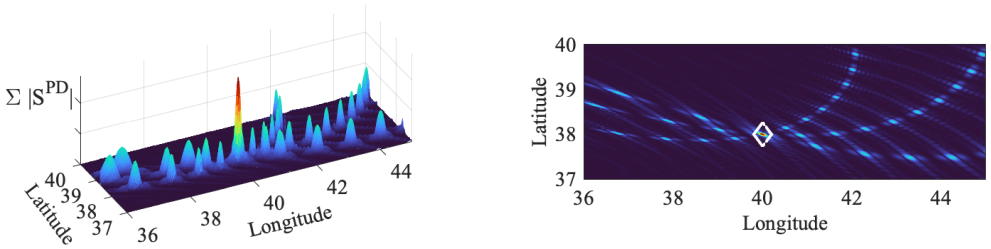


Figure 3.5: Position-domain correlation after noncoherently accumulating all 58 snapshots of the simulation with the 3D view (left) and from above (right). Likelihood accumulates at the true source location, whereas the shifting ambiguities are suppressed, demonstrating that direct geolocation provides a powerful mechanism for geolocating cyclostationary interference.

Another advantage of direct geolocation is that there will be a peak at every position where an emitter is located, provided that the emitter's signal was received at both receivers. To show this, another simulation study was conducted. Fig. 3.6 shows the geometry of the simulated capture scenario. The ground tracks of the two LEO-based receivers over 58 seconds are shown. A GPS L1 C/A spoofer, pure tone jammer, chirp jammer, and sawtooth jammer were simulated to be in a rectangular formation on the surface of the Earth. All interference sources were set to be transmitting centered at the GPS L1 frequency (1575.42 MHz) at the same power. The receivers performed 58 50 ms raw captures, each spaced by a second.

The position domain search space shown in Fig. 3.6 spans approximately 10° in longitude and approximately 4° in latitude. The position-domain correlation was computed at each of the 58 snapshots. The position-domain correlation $|S^{\mathrm{PD}}|$ for the first snapshot and the last snapshot are shown in

Fig. 3.7. One can immediately see the different correlation structures for the different waveforms. The GPS L1 C/A spoofer creates a single sharp peak, the tone creates a large streak, the chirp jammer creates multiple equally prominent peaks, and the sawtooth jammer creates an "X" shape. At a single snapshot, the pure tone and chirp jammers' positions are ambiguous. But, as seen between the first and last snapshot, the structure of each GNSS interference source changes. This is due to the changing receiver geometry. However, the peaks at the true emitter positions do not change from the start of the capture compared to the end of the capture, as shown in the prior simulation.

Shown in Fig. 3.8 is the position-domain correlation grid when the 58 position-domain grids at each snapshot are noncoherently accumulated. One can immediately appreciate four sharp peaks centered at each of the four GNSS interference sources' true position. As the receiver geometry changes over the 58 second interval, the positions of the sidelobes and other undesirable structures due to the waveform change. When position-domain grids are noncoherently accumulated, the undesirable structures are reduced to the noise floor, while the true peaks persist. This shows how powerful direct geolocation can be in crowded signal environments with waveforms that have undesirable structures. In the two-step approach, the CAF peak tracker would have been easily overwhelmed.

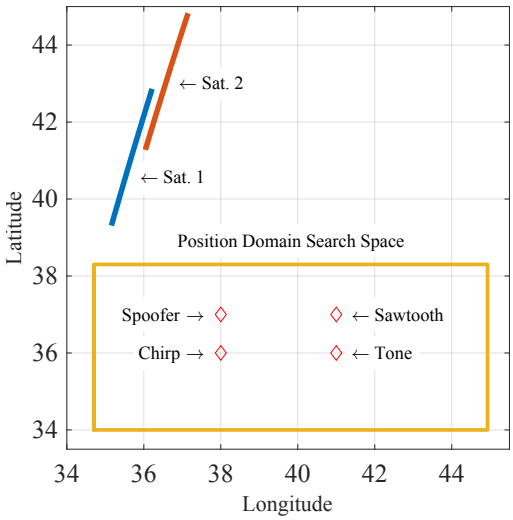


Figure 3.6: Simulation setup with multiple interference sources.

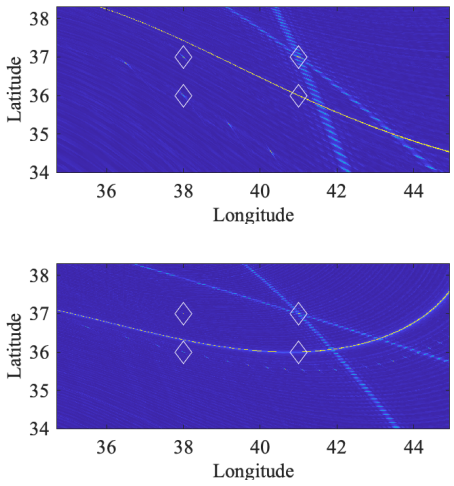
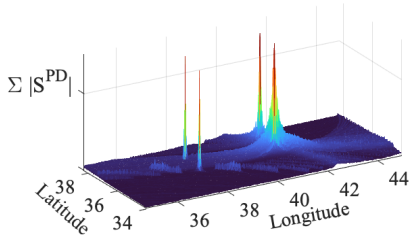


Figure 3.7: Position domain correlation at the first snapshot (top) and the final snapshot (bottom).



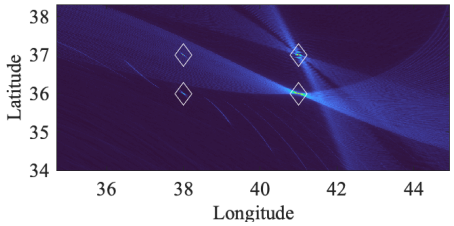


Figure 3.8: Position domain correlation after noncoherently accumulating all 58 snapshots with the 3D view (left) and from above (right). The four distinct sharp peaks represent the MLE of each GNSS interference source's position.

3.4 A Real-world Capture

3.4.1 Capture Platform

Spire Global Inc. operates a global network of over 100 satellites. Among these satellites, are *STRATOS* satellites, whose original purpose was

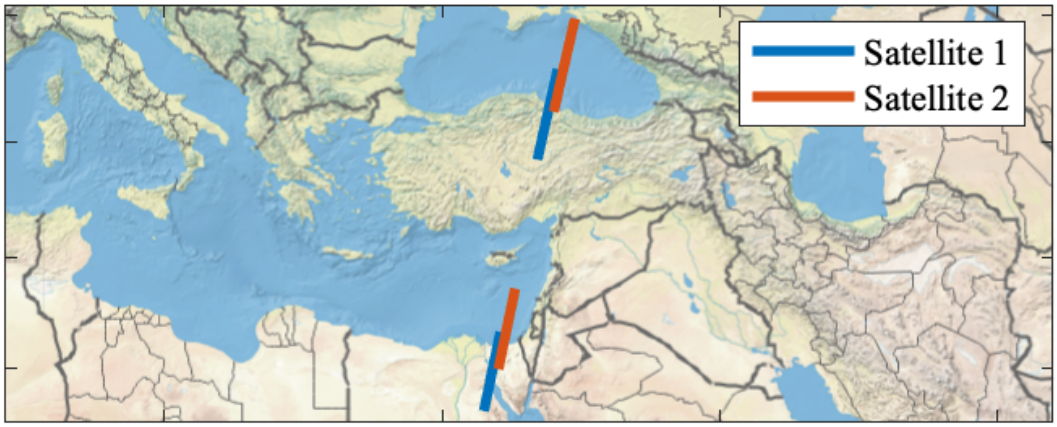


Figure 3.9: Ground tracks of a pair of LEO-based receivers during the two 60-second captures over the Eastern Mediterranean. Both satellites are moving from north to south.

GNSS radio occultation (GNSS-RO), that can be repurposed for geolocation of emitters operating in the GPS L1 and L2 bands. The *STRATOS* satellites carry one wide-field-of-view zenith-facing antenna for precise orbit determination, and one or two Earth-limb-facing narrow-field-of-view high-gain antennas. *STRATOS*'s RF circuitry has three dual-frequency channels, with each antenna connected to one of the front end channels. Digitization of each signal happens coherently based on a single sampling clock. Simultaneous collections of the 2-bit quantized 6.2 Msps raw intermediate frequency (IF) samples centered at GPS L1 and L2 can be performed. These data can be packaged and downlinked through Spire's network of dedicated ground stations.

During April 2022, two *STRATOS* satellites performed two consecutive 60-second simultaneous capture events separated by 180 seconds, while over the Eastern Mediterranean, as shown in Fig. 3.9. During each 60-second cap-

ture, the satellites had an average altitude of 524 km and an average velocity of 7,678 m/s, traveling from north to south. The forward-facing antenna was used for the raw capture in the first simultaneous event, and the backward facing antenna was used during the second simultaneous event. The zenith-facing antenna on each satellite was connected to the onboard GNSS receiver that calculated the onboard navigation solution, which includes receiver position, velocity, clock bias, and clock offset rate estimates. The raw samples from the GNSS-RO antennas and the onboard navigation solution were downlinked for post-processing.

3.4.2 Spectrum Analysis

Fig. 3.10 and Fig. 3.11 illustrate the captured signals' spectral characteristics. There is composite wideband interference on both GPS L1 and L2. Visual inspection of the spectogram indicate that L1 contains multiple CW chirp jammers as well as other wideband interference. The wider-bandwidth triangle jammer had a bandwidth of approximately 200 kHz and a period of 8 ms. The narrower-bandwidth triangle jammers had a bandwidth of approximately 20 kHz and a period of 100 ms. The 20 kHz triangle jammers have the same parameters of a jammer previously captured in the Eastern Mediterranean [89]. The power spectral densities shown in Fig. 3.10 were averaged across the entire capture. L2 contains a particularly strong narrowband jammer as well as other wideband interference. The goal is to geolocate all sources of interference that are captured by both receivers.

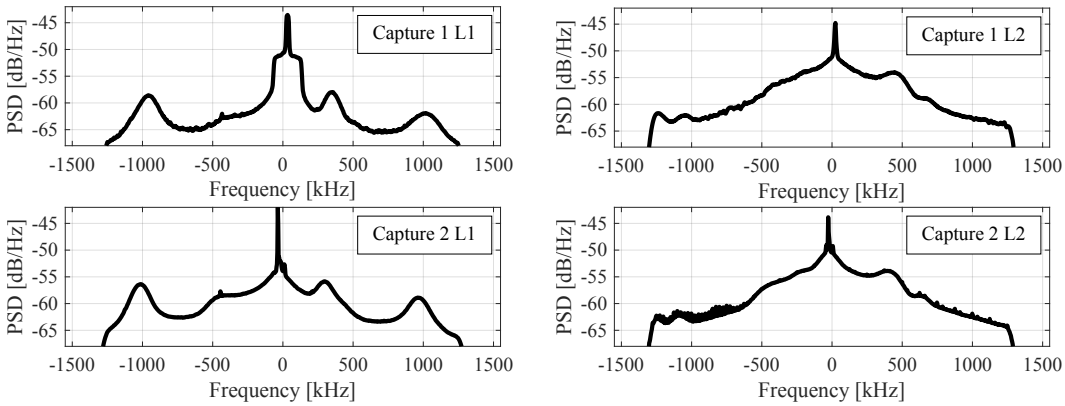


Figure 3.10: The average power spectral densities of the data during both captures at both frequencies.

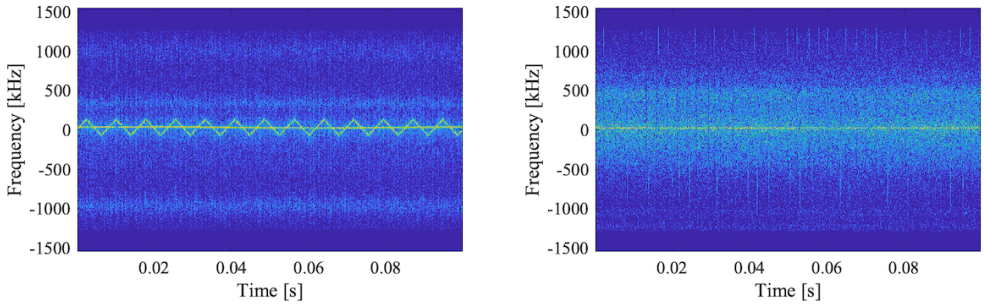


Figure 3.11: Spectograms centered at the GPS L1 (left) and L2 (right) frequencies from data captured in April 2022.

3.5 Emitter Geolocation

This section discusses two-step and direct geolocation of the emitters captured in the Spire dataset. The results focus on the geolocation of a known GNSS jammer operating out of Khmeimim Air Base on the coast of Syria. Additionally, the advantages of direct geolocation are showcased for low SNR signals, crowded signal environments, and for signals exhibiting cyclostation-

arity. Wide-area GNSS interference monitoring via direct geolocation reveals numerous jammers across Syria, Turkey, Iraq, Ukraine, and Israel.

3.5.1 Two-Step Geolocation

Fig. 3.12 shows the time history of TDOA and FDOA measurements corresponding to the largest peak in the L2 CAF over the 60-second capture. These measurements were taken at 5 Hz and then served to the nonlinear estimator. The final two-step emitter position estimate is displayed in Fig. 3.12. Shown in Fig. 3.13 is the final solution which converged to Khmeimim Air Base on the coast of Syria, which has been previously discovered to host a powerful GPS jammer [100].

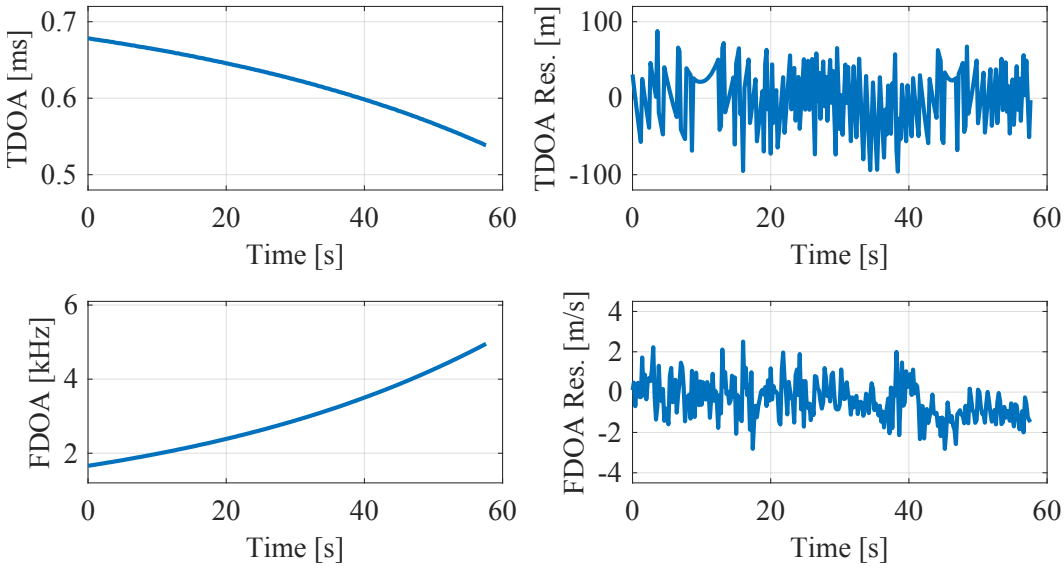


Figure 3.12: Shown on the left are TDOA (top) and FDOA (bottom) measurements corresponding to the measurements from the interference source located at the Khmeimim Air Base. Shown on the right are the TDOA measurement residuals (top) and the FDOA measurement residuals (bottom).



Figure 3.13: The final two-step T/FDOA geolocation estimate converges to Khmeimim Air Base on the coast of Syria, which has been previously discovered to have transmitted GPS jamming signals [100]. The semi-major and semi-minor axes of the 95% ellipse are 632 and 57 meters, respectively when a conservative $\sigma_{\Delta\tau} = 50$ m and $\sigma_{\Delta f} = 5$ m/s are assumed.

The TDOA and FDOA residuals from the nonlinear least-squares estimator are also shown in Fig. 3.12. They are zero-mean, indicating that the presented measurement model is accurate. The TDOA residuals have a standard deviation of 37 m and the FDOA residuals have a standard deviation of 1 m/s. The TDOA residuals time history appears to have unexpected structure; this structure arises due to the range resolution from the sampling rate. The TDOA measurement resolution is limited by the sampling rate because nearest-sample TDOA estimation was performed. Because the sampling rate at baseband is 3.1 Msps, the range resolution is 96.7 m. The TDOA residuals remain between the range resolution of $\pm$ 96.7 m, meaning the expected performance with the sampling rate was achieved. The cost surface for the

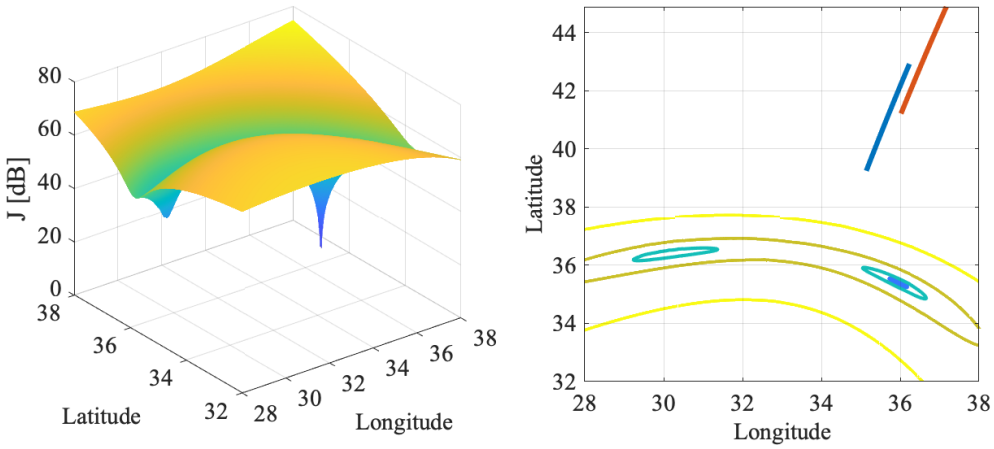


Figure 3.14: Cost surface for the T/FDOA measurements. The global minimum corresponds to the Khmeimim Air Base.

calculated set of T/FDOA measurements corresponding to the jammer at the Khmeimim Air Base is shown in Fig. 3.14. The global minimum of each cost surface aligns with Khmeimim Air Base.

The capture's specific receiver geometry gives rise to a reflection point on the opposite side of the satellites' average ground track. For a single T/FDOA measurement pair, the true emitter position and the alias are indistinguishable. The ambiguity between these points begins to resolve if the baseline between the receivers changes with each T/FDOA measurement pair. As the receiver geometry changes, each T/FDOA measurement pair in the time series will have a different alias location because of the time-varying baseline between the receivers. Conversely, each T/FDOA measurement pair in the time series will be consistent with the true emitter position, resolving the ambiguity.

3.5.2 Direct Geolocation

Fig. 3.15 shows contours of the position-domain correlation values over the Khmeimim Air Base using the L2 raw samples with various integration times. Each point in the grid was separated by .0005° in latitude and longitude. For each latitude and longitude pair, the surface altitude was retrieved from a terrain model. The global elevation model used was the Global Multi-resolution Terrain Elevation Data 2010 (GMTED2010) developed by the U.S. Geological Survey and the National Geospatial-Intelligence Agency. The elevation is measured relative to mean sea level using Earth gravitational model EGM-96. The root mean square error is approximately 26 meters.

The maximum peak in the position-domain correlation function for the various integration intervals in Fig. 3.15 is nearly equivalent to the two-step solution. This is expected because in the high SNR scenarios with long captures, the two-step and one-step solutions are equivalent. In the one-step approach, the effects of integration interval length are apparent in Fig. 3.15. As the integration time increases, the peak becomes sharper and more pronounced. This is because over longer intervals, the signal decorrelates more quickly for incorrect emitter positions. The more wideband the signal is, the more quickly it decorrelates and the peak becomes sharper. As stated earlier, the longer coherent integration interval results in better frequency resolution, which is apparent here.

A consequence of longer coherent integration is that a finer grid spacing is required. If wide-area coverage is desired, one might opt for a coarse grid

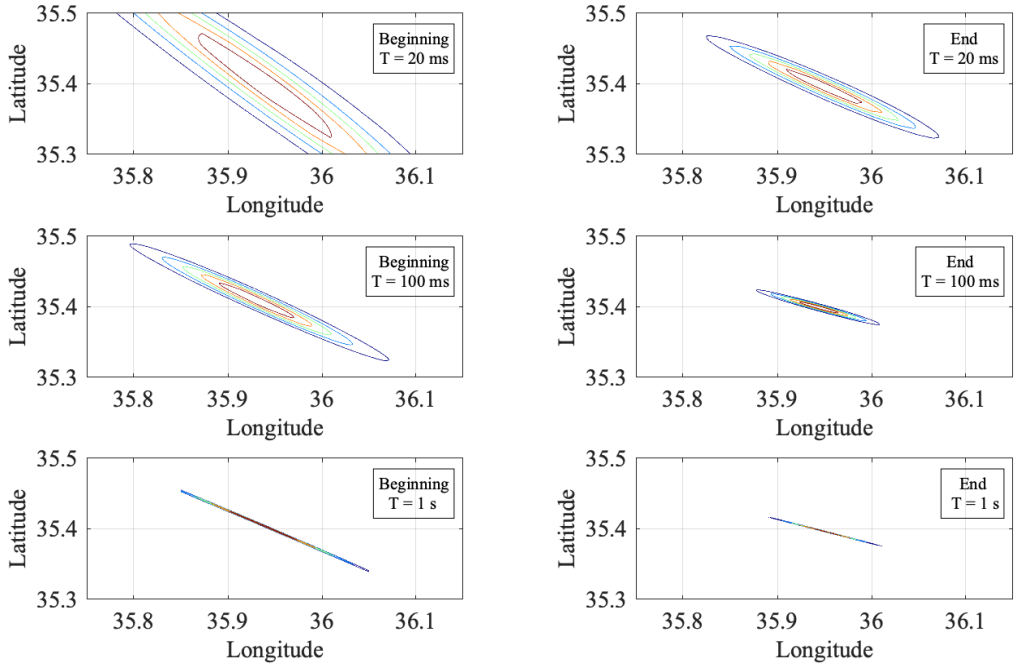


Figure 3.15: Contours of the direct geolocation position-domain correlation grids over Khmeimim Air Base for 20-ms (top), 100-ms (middle), and 1-second (bottom) integration times during the beginning (left) and end(right) of the 60-second capture. The impact of integration time and geometry is apparent.

covering a large area paired with shorter integration intervals. After getting a general sense of the emitters' locations via the course grid, a finer grid with longer integration intervals can be used to get more accurate geolocation estimates.

Fig. 3.16 demonstrates how direct geolocation is better suited for processing captures with structured, cyclostationary signals. The position-domain correlation function for the triangle jammer on L1 is shown at the beginning (top graph) and end (middle graph). There is structure in both, but that structure changes with the instantaneous receivers' geometry, which changes

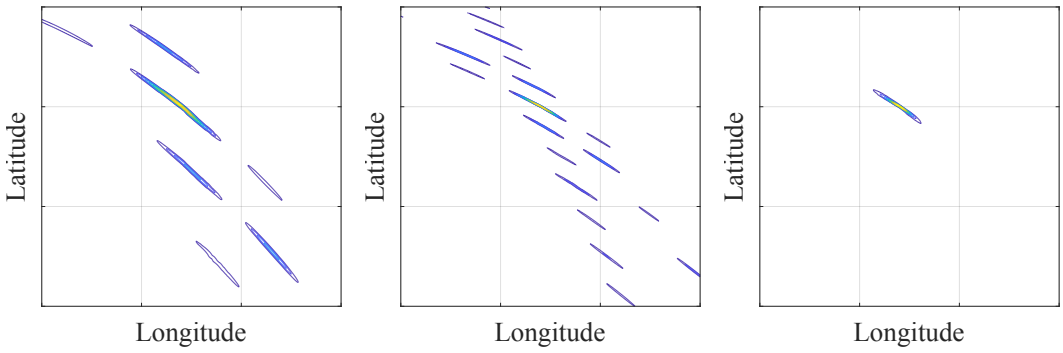


Figure 3.16: Direct geolocation on the chirp jammer in the beginning (left) and end (middle) of the capture. A 100 ms accumulation was taken once every 5 seconds. The eleven position-domain correlations were noncoherently combined (right). Noncoherent integration is a powerful tool that suppresses the spurious false peaks from waveform structure and cyclostationarity.

over time. When multiple position-domain accumulations are noncoherently combined, all of the false peaks are reduced below the noise floor, while the main peak corresponding to the true emitter position remains. The position of the higher-bandwidth triangle jammer is estimated to be in northern Israel. Noncoherent integration is a powerful tool that can suppress spurious false peaks from waveform structure and cyclostationarity.

Fig. 3.17 shows direct geolocation over Syria with the L2 data. A 100 ms accumulation was taken once per second, over 58 seconds. The 58 position-domain correlations were noncoherently combined. Each peak corresponds to an emitter position estimate, with the largest peak corresponding to Khmeimim Air Base. The receivers' geometry also heavily affects correlation in the position domain. There was better receiver geometry for the transmitters on the east-side, resulting in sharper peaks.

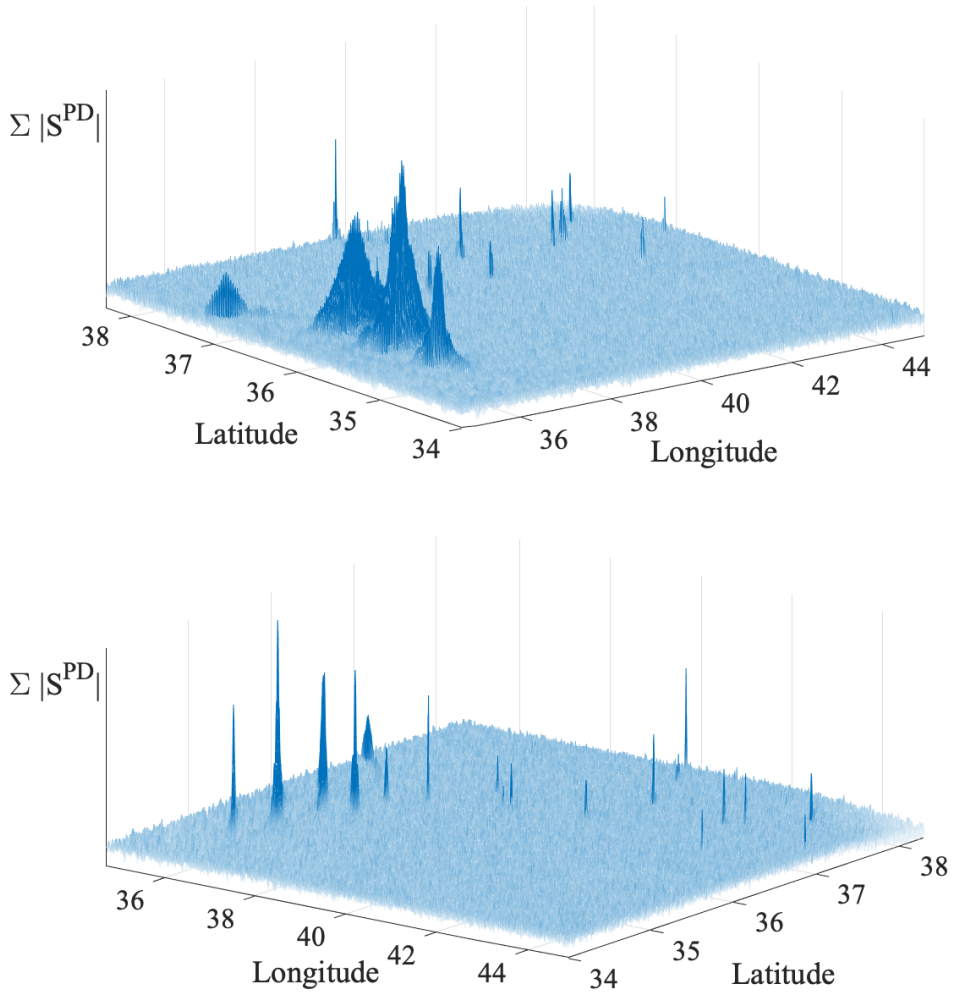


Figure 3.17: Direct geolocation over Syria with two different views of the position domain. A 100-ms accumulation was taken once per second, over 58 seconds. Shown here are 58 position-domain correlations noncoherently combined. Each peak corresponds to an emitter position estimate.

This position-domain correlation was repeated for L1 on the first capture and L1 and L2 on the second capture. The peaks from the position

domain correlations on both captures and frequencies are shown in Fig. 3.18. The final position estimate's 95% semi-major for was less than 10 km for all sources. Each one of these estimates is plausible, given the agreement between both frequencies, both captures with different geometry, and the surrounding equipment near each estimate from publicly available optical satellite imagery. Additionally, the same geolocation estimates were also achieved through the two-step approach with manual peak tracking. This showcases the superiority of the direct approach in crowded signal environments. The two-step approach would have depended on an additional complex multi-peak tracking and association algorithm to generate a time series of T/FDOA measurements, whereas in the direct approach the emitter position estimate comes directly from the raw samples.

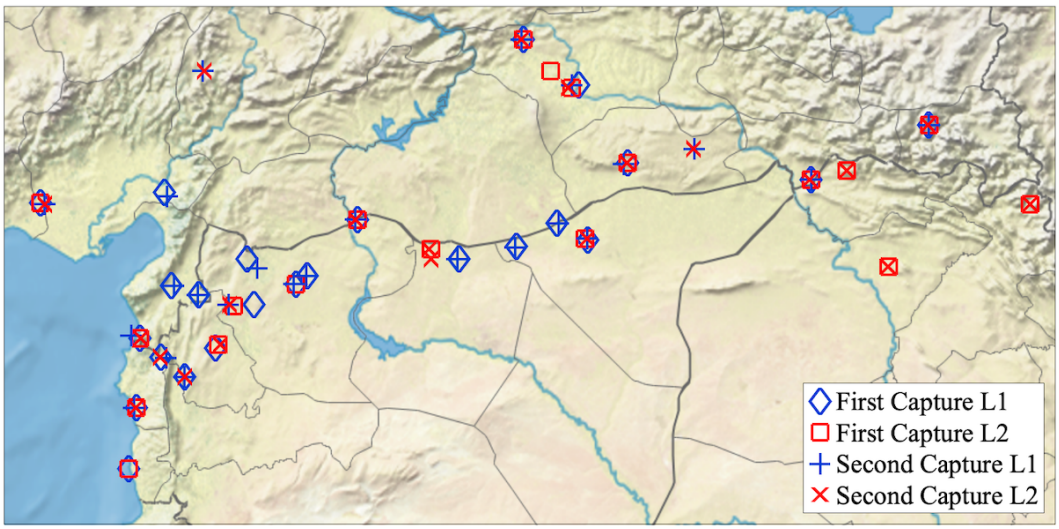


Figure 3.18: Final emitter position estimates on L1 and L2 across both simultaneous captures from direct geolocation. All emitters were assumed to be stationary and constrained to the surface of the Earth.

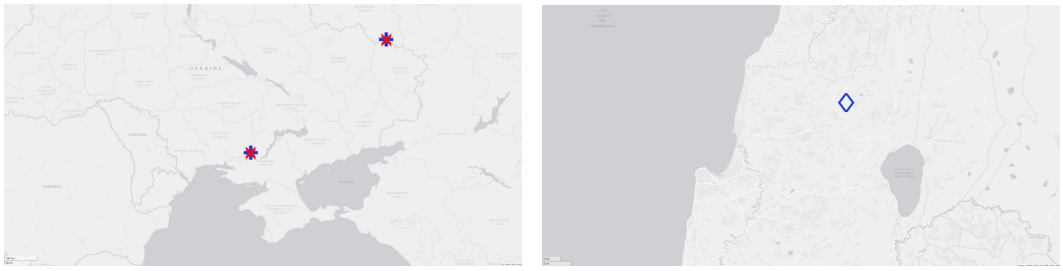


Figure 3.19: Jammers found in Ukraine (left) and Israel (right).

3.6 Conclusion

This chapter explored two-step and direct geolocation of terrestrial emitters from LEO. The full algorithms for two-step and single-step direct geolocation were presented. It was demonstrated that the direct approach is a powerful geolocation technique for low SNR signals during short captures with multiple emitters. This chapter also investigated emitter geolocation in crowded signal environments and explored geolocating cyclostationary signals. The relationship between the CAF domain and the position domain was explored and shown to be mappings of each other. Finally, this chapter demonstrated two-step and direct geolocation on raw intermediate frequency samples recorded from Spire Global's LEO constellation. Recent real-world GNSS interference signals captured by two time-synchronized LEO receivers over the Eastern Mediterranean were characterized and their emitters geolocated.

Chapter 4

Transient Space-Based GNSS Interference Detection and Estimation

Abstract

This chapter analyzes and identifies a space-based Global Navigation Satellite System (GNSS) interference source that has caused scores of powerful transient wide-area interference events over continental Europe, Greenland, and Canada since 2019. While terrestrial or near-terrestrial sources are primarily responsible for the recent uptick in GNSS interference worldwide, space-based interferers are of special concern given their potential for vast geographic reach and their portent of a qualitative escalation in GNSS interference. Based on data collected between 2019 and 2026 from a network of terrestrial GNSS reference stations, this chapter (1) develops a received-power-based detection framework; (2) details the spatial, temporal, and spectral patterns of wide-area interference events caused by the source; (3) presents and analyzes identification techniques that blend received-power and time-difference-of-arrival measurements; and (4) applies these techniques to confidently identify the GNSS interference source as a constellation of Russian early warning satellites in Molniya ("lightning") orbits.

4.1 Introduction

Global Navigation Satellite Systems (GNSS) such as GPS provide meter-accurate positioning while offering global accessibility and all-weather, radio-silent operation. However, GNSS is fragile: its service is easily degraded by both deliberate (e.g., jamming and spoofing) and naturally occurring (e.g., multipath and atmospheric) interference [62, 65, 68, 120]. The past five years have seen a crescendo of GNSS disruptions in the aviation and maritime sectors [52, 90, 108, 139]. Interference is especially problematic in these fields because GNSS is deeply interwoven into critical systems. Beyond simple navigational or timing displacement, cascading failures can occur when GNSS loss or corruption triggers follow-on problems in downstream systems [52, 139].

Fortunately, significant progress has been made over the past decades in GNSS interference mitigation and countermeasures [13, 28, 114, 115]. It has been shown that dedicated GNSS interference monitoring systems provide valuable spectrum situational awareness to enhance navigation security [54, 55, 82, 98, 122, 123], with some demonstrating the ability to geolocate the interference source [11, 95]. Furthermore, GNSS receivers situated in LEO enable terrestrial GNSS interference detection, characterization, and geolocation with worldwide coverage [23, 30, 100].

Networks of terrestrial GNSS receivers can also be enlisted for GNSS interference monitoring [1], including of space-based interference. [140] revealed that GPS PRNs 24 and 27 (SVN65 and SVN66) were transmitting leakage tones at ± 10 MHz and ± 20 MHz from the GPS L5 frequency. The anomaly

was subsequently fixed. More recently, a network of terrestrial reference stations was used to identify a BeiDou satellite (NORAD ID 40749) as a source of interference in the B3I band (1268.52 MHz) [110, 111]. The offending satellite continuously transmitted a tone centered at 1268.52 MHz and tones at ±10 MHz and ±20 MHz from this frequency. These tones are no longer present at the time of writing. Both of these interference cases were apparently due to faulty or improperly configured hardware.

This chapter reports on the detection, characterization, and source identification of powerful transient GNSS interference originating from a non-GNSS satellite. The phenomenon was briefly mentioned in [39], and more fully characterized in a preliminary conference publication by the current authors, [31]. The current chapter significantly extends both works. The interference reported here is different from that detailed in [110, 111, 140] in several respects: (1) it is transient rather than continuous—the duration of interference events is less than 10 seconds; (2) it affects signals in the widely used GPS L1 frequency band rather than the less-common L5 and B3I bands; (3) it is much more powerful, causing drops in terrestrial receiver carrier-to-noise ratios (CNRs) of up to 10 dB; and (4) it is not merely an anomaly of existing GNSS signals, nor does it originate from a GNSS satellite.

Obviously, high-powered interference with continental reach affecting the GPS L1 band—the primary band used for global aviation, shipping, and precise timing—is of serious concern. If deliberate, it portends a qualitative escalation in GNSS interference.

The effects of this interference are evident in public data from a network of terrestrial reference stations operated by the International GNSS Service (IGS) [34, 74, 81, 96]. Most useful are data from a subset of IGS stations that produce high-rate (1-Hz) GNSS observables. These data may be retrieved from the Crustal Dynamics Data Information System archive, which is made available through NASA's archive of space geodesy data [103]. On scores of occasions since 2019, all tracked signals on the GPS L1 frequency at IGS reference stations across Europe, Greenland, and Canada simultaneously saw a sudden brief drop in CNR. The onset of the disruption was synchronous to within the 1-Hz sampling resolution, suggesting a single source per event. The affected terrestrial receivers span a geographic area so large that no single ground-based or aircraft-based source could reach them all; hence the space-based origin hypothesis.

While CNR variations across a receiver network are a valuable metric for space-based interference detection, it will be shown that they are not sufficiently informative for unique source identification. By contrast, a source's position and velocity can be instantaneously estimated using time- and frequency-difference-of-arrival (T/FDOA) techniques if a spatially diverse network of four or more ground stations simultaneously captures raw broadband samples in the band affected by the interference. Moreover, by referencing a catalog of satellite ephemerides (assuming the source satellite is listed), even two stations are adequate to narrow the possibilities to a manageable few. An association framework that fuses both CNR and TDOA measurements can then allow a

unique identification, as will be shown.

This chapter makes six main contributions. First, it presents measurement models and a detection framework for transient wide-area interference. Second, it details the spatial, temporal, and spectral properties of multiple wide-area GNSS outage events from the space-based interference source and distinguishes these from a naturally occurring solar radio burst. Third, it presents a basic satellite identification strategy to narrow the candidate satellites and estimate the minimum satellite altitude at apogee. Fourth, it presents an advanced satellite associate framework using the generalized likelihood ratio test (GRLT) and applies the framework to a test scenario. Fifth, it presents a framework for instantaneously identifying an interference satellite based on a brief time history of TDOA measurements, and gives an error sensitivity analysis. Sixth, it combines IGS CNR data and raw wideband samples from two additional receivers in Europe to confidently identify the source, which is revealed to be a small constellation of Russian satellites in Molniya ("lightning") orbits.

Compared to the preliminary conference version of this chapter in [31], contributions two and three are extended and contributions four, five, and six are novel.

4.2 Measurement Models and Detection

The IGS reference station network collects and provides observables in the Receiver Independent Exchange Format (RINEX) [74]. These GPS-time-

tagged observables include carrier phase, pseudorange, Doppler, and CNR measurements for each tracked GNSS satellite. CNR observables from GPS L1 C/A signals at stations providing high-rate (1-Hz) GNSS observables are the focus of the following analysis.

4.2.1 Measurement Models

Let $\mathcal{I}$ be the set of all relevant terrestrial reference stations (primarily IGS stations), and let $\mathcal{J}$ be the set of all unique GNSS signals that can be tracked by such stations at a given center frequency (e.g., the GPS L1 frequency). By way of notation, let $|\mathcal{I}|$ denote the cardinality of $\mathcal{I}$, and similarly for all other sets. For station $i \in \mathcal{I}$ and signal $j \in \mathcal{J}$, let CNR_{ij} denote the true CNR, expressed as

$$\mathrm{CNR}_{ij} = P_{\mathrm{R}}^{ij} - N_0^i \qquad \text{(dB-Hz)} \tag{4.1}$$

$$P_{\mathrm{R}}^{ij} = P_{\mathrm{T}}^{j} + G_{\mathrm{T}}^{j}\left(\theta_{\mathrm{T}}^{ij}, \phi_{\mathrm{T}}^{ij}\right) + G_{\mathrm{R}}^{i}\left(\theta_{\mathrm{R}}^{ij}, \phi_{\mathrm{R}}^{ij}\right) + L_{ij} \qquad \text{(dBW)} \tag{4.2}$$

where P_{R}^{ij} is the received power (dBW), N_0^i is the thermal noise density (dBW/Hz), P_{T}^{j} is the transmit power (dBW), G_{T}^{j} is the transmitter antenna's gain function (dB), G_{R}^{i} is the receiver antenna's gain function (dB), θ_{R}^{ij} and θ_{T}^{ij} are the off-boresight angles at the receiver and transmitter, ϕ_{R}^{ij} and ϕ_{T}^{ij} are the azimuth angles at the receiver and transmitter, and L_{ij} is the path loss (dB), defined as

$$L_{ij} = 20\log_{10}\left(\frac{\lambda}{4\pi\rho_{ij}}\right) \qquad \text{(dB)} \tag{4.3}$$

where λ is the carrier wavelength and ρ_{ij} is the range to the satellite transmitting signal j. The noise power density N_0^i also includes the effects of multi-access interference from GNSS signals other than the jth one on the same frequency.

CNR becomes the carrier-to-interference-and-noise ratio (CINR) when there is at least one interference signal present. This chapter assumes a single interference signal and source at each epoch. Let CINR_{ij} denote the true CINR for $i \in \mathcal{I}$ and $j \in \mathcal{J}$, expressed as

$$\mathrm{CINR}_{ij} = P_{\mathrm{R}}^{ij} - 10\log_{10}\left(\tilde{N}_0^i + \tilde{I}_0^i\right) \qquad \text{(dB-Hz)} \tag{4.4}$$

where $\tilde{N}_0^i$ is the linear unit equivalent of the noise power density N_0^i, satisfying $N_0^i = 10\log_{10}(\tilde{N}_0^i)$, and $\tilde{I}_0^i$ is the linear unit equivalent of the interference power density I_0^i (dBW/Hz), satisfying $I_0^i = 10\log_{10}(\tilde{I}_0^i)$. Following [65], let $I^i(t)$ be the interference component of the product between the received signal (including the interference and thermal noise) and the local replica of the desired signal, and let $S_{\mathrm{I}}^i(f)$ be its power spectral density (PSD, dBW/Hz). Then I_0^i may be defined as

$$I_0^i = S_{\mathrm{I}}^i(0) \qquad \text{(dBW/Hz)} \tag{4.5}$$

If $\tilde{S}_{\mathrm{C}}^i(f)$ is the PSD of the local replica signal's spreading code, and $\tilde{S}_{r_{\mathrm{I}}}^i(f)$ is the PSD of the received interference signal—both in linear units—and $\hat{f}_{\mathrm{D}}^i$ is the receiver's estimate of the desired signal's apparent Doppler frequency (Hz), then $\tilde{S}_{\mathrm{I}}^i(f)$, the linear unit equivalent of $S_{\mathrm{I}}^i(f)$, is formed by the convolution

$$\tilde{S}_{\mathrm{I}}^i(f) = \tilde{S}_{\mathrm{C}}^i(f) \star \tilde{S}_{r_{\mathrm{I}}}^i(f + \hat{f}_{\mathrm{D}}^i) \tag{4.6}$$

Similar to (4.2), let P_{I}^{i} denote the received interference signal's power at the ith station, given by

$$P_{\mathrm{I}}^{i} = P_{\mathrm{I}} + G_{\mathrm{I}}\big(\theta_{\mathrm{I}}^{i}, \phi_{\mathrm{I}}^{i}\big) + G_{\mathrm{R}}^{i}\big(\theta_{\mathrm{R}}^{i\mathrm{I}}, \phi_{\mathrm{R}}^{i\mathrm{I}}\big) + L_{i\mathrm{I}} \qquad \text{(dBW)} \tag{4.7}$$

where P_{I} is the transmitted interference power and G_{I} is the interference source antenna's gain function. The arguments of G_{I} and G_{R}^{i}, and the quantity $L_{i\mathrm{I}}$ are as before except that they apply to the interference source rather than to the jth GNSS signal.

At the kth time epoch, the receiver at the ith station reports the measurement $z_{ij}[k]$, which is either the measured CNR_{ij} under nominal operating conditions (the null hypothesis H_0), or is the measured CINR_{ij} under interference (the alternate hypothesis H_1), for all $j \in \mathcal{J}$. The reported measurement $z_{ij}[k]$ is modeled under H_0 and H_1 as

$$H_0: \quad z_{ij}[k] = \mathrm{CNR}_{ij}[k] + w_{ij}[k] \tag{4.8}$$

$$H_1: \quad z_{ij}[k] = \mathrm{CINR}_{ij}[k] + w_{ij}[k] \tag{4.9}$$

where $w_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2)$ is zero-mean additive white Gaussian noise (AWGN) that models measurement error due to thermal noise, atmospheric effects, multipath, and other minor effects.

4.2.2 Detection

Determining whether transient interference is present or not at each station, and across a network of stations, is an exercise in detection theory.

The problem is to identify an interference event that typically persists for 3 to 5 seconds. Detection techniques from the nascent transient change detection literature are well suited to this problem [35, 53, 99]. During a transient interference event, the measurement $z_{ij}[k]$ transitions from H_0 to H_1—due to a sudden injection of I_0^i—and then back to H_0 as the event concludes. An optimal transient change detector, such as proposed in [53], is matched to the interference-induced pattern in $z_{ij}[k]$. But in cases like the present one in which the interference interval is unknown *a priori*, no optimal solution is available. [35] recommend a windowed solution, and derive certain performance bounds. The current chapter approximates the windowed solution with a simple differencing technique that is more amenable to analysis and becomes optimal for single-epoch change durations [99].

Let $T_{\max}$ be the maximum duration of the transient event, and Δt be the measurement epoch ($\Delta t = 1$ second in the current context). Set $l = \lceil T_{\max}/\Delta t \rceil + 1 \in \mathbb{N}$ as the detector's index stride. Then for the kth time-epoch, ith station, and jth tracked GNSS signal, let the signal-specific detection statistic be given by

$$\xi_{ij}[k] = \tfrac{1}{2}(z_{ij}[k+l] - 2z_{ij}[k] + z_{ij}[k-l]) \tag{4.10}$$

Under H_0 and H_1, $\xi_{ij}[k]$ takes on distributions

$$H_0: \quad \xi_{ij}[k] \sim \mathcal{N}(0,\ \tfrac{3}{2}\sigma_{ij}^2) \tag{4.11}$$

$$H_1: \quad \xi_{ij}[k] \sim \mathcal{N}(\mu_i,\ \tfrac{3}{2}\sigma_{ij}^2) \tag{4.12}$$

where the mean μ_i is

$$\mu_i = 10 \log_{10} \left(\frac{\tilde{N}_0^i + \tilde{I}_0^i}{\tilde{N}_0^i} \right) \qquad \text{(dB)} \tag{4.13}$$

and can be calculated if all relevant parameters are known.

Let $\mathcal{J}_i[k] \subset \mathcal{J}$ be the set of GNSS signals tracked by station i at epoch k. Then the station-specific interference detection statistic $\Lambda_i[k]$ is the across-signal average of $\xi_{ij}[k]$:

$$\Lambda_i[k] = \frac{1}{|\mathcal{J}_i[k]|} \sum_{j \in \mathcal{J}_i[k]} \xi_{ij}[k] \tag{4.14}$$

The distribution of $\Lambda_i[k]$ under H_0 and H_1 is

$$H_0 : \quad \Lambda_i[k] \sim \mathcal{N}(0,\, \sigma_i^2) \tag{4.15}$$

$$H_1 : \quad \Lambda_i[k] \sim \mathcal{N}(\mu_i,\, \sigma_i^2) \tag{4.16}$$

The quantity μ_i amounts to the average drop in observed CNR at the ith station due to the interference event, in dB. In practice, the parameters that contribute to $\tilde{I}_0^i$ are unknown *a priori*, leaving μ_i unknown. But because $\tilde{I}_0^i > 0$, it follows that $\mu_i > 0$, which allows it to be cancelled from each site's detection statistic, making a test based on (4.14) uniformly most powerful [132]. The value σ_i^2 can be estimated at the ith station using historical data, and from this estimate a detection threshold ν_i can be calculated for a constant false alarm rate (CFAR). The hypothesis test for the ith station then becomes

$$\Lambda_i[k] \underset{H_0}{\overset{H_1}{\gtrless}} \nu_i \tag{4.17}$$

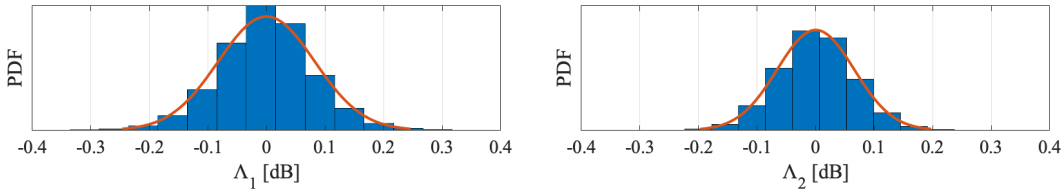


Figure 4.1: The distributions of Λ_i under H_0 (no interference present) for stations METG ($i = 1$) and MATE ($i = 2$).

Null-hypothesis (H_0) detection statistic distributions from a day's worth of data for two stations, METG and MATE, are shown in Fig. 4.1. Both appear to be zero mean and approximately Gaussian distributed. The modest difference in σ_i^2 can be attributed to local noise, different receivers and antennas, and quantization of the reported CNR.

The remainder of this section focuses on day 160 of year 2021. Fig. 4.2 shows the time history of $\{z_{ij}\}$ and $\{\Lambda_i\}$ over a 15-minute interval for three stations with $l = 3$, amounting to a 3-second stride at the 1-Hz measurement rate. The dashed red lines in the right plots indicate the detection thresholds corresponding to a 10^{-4} probability of false alarm. Around the 700-second mark, interference is detected simultaneously at all three stations. While individual interference detections are not uncommon because reference stations occasionally experience local interference, time alignment of an interference event across geographically distant sites in Finland, Italy, and Greenland is noteworthy. Moreover, the interference exhibits the same characteristic signature at each location: a CNR drop lasting approximately 3 seconds. This suggests that all three stations were affected by a common interference source.

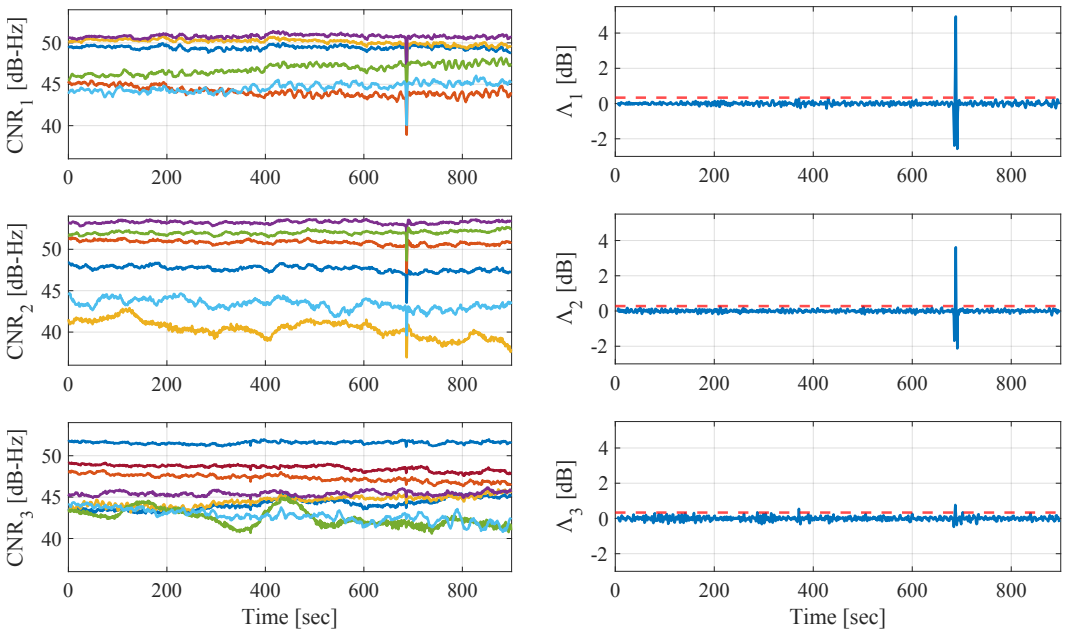


Figure 4.2: Reported CNR (left) and the detection statistic Λ_i (right) for IGS stations METG ($i = 1$, Finland), MATE ($i = 2$, Italy), and THU2 ($i = 3$, Greenland) over a 15-minute interval on day 160 of year 2021. The dashed red line is the detection threshold with a 10^{-4} probability of false alarm.

The detection test was conducted continuously across stations spanning Europe, Greenland, and Canada on day 160 of 2021. Fig. 4.3 shows the number of stations registering detections during this period with a 10^{-4} probability of false alarm. A lower-power interference event was detected by 21 stations, followed by a more powerful event detected by 58 stations. The stronger event corresponds to the 700-second mark shown in Fig. 4.2. The recorded CNR observables at each station align with the patterns depicted in Fig. 4.2, providing further evidence that these stations were affected by the same interference source. A heat map showing the spatial distribution of the

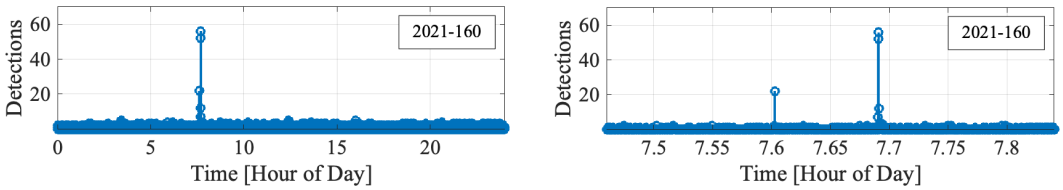


Figure 4.3: The number of stations that detected interference during day 160 of year 2021. The expanded view on the right shows a lower-power event detected by 21 stations, followed by a higher-power event detected by 58 stations.

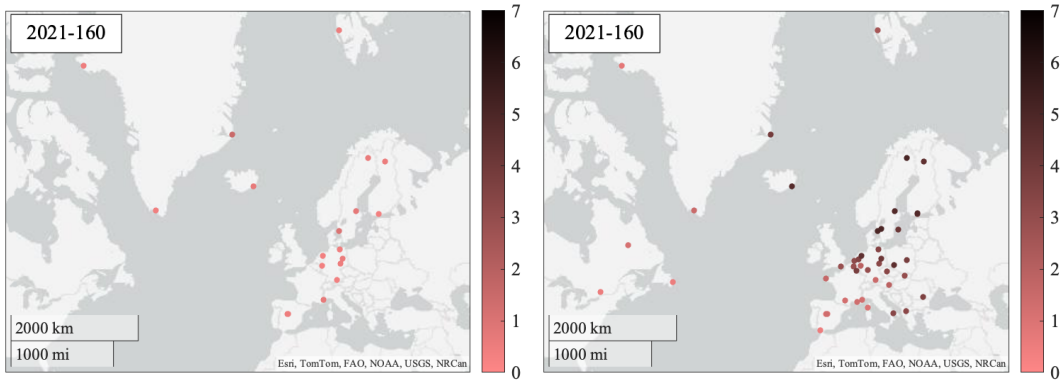


Figure 4.4: Heat map of the test statistic (in dB) at triggering stations for the lower-power event (left) and the higher-power event (right). The drop in GPS L1 C/A CNR during the more powerful interference event was as large as 6 dB, centered near the Baltic region.

CNR reduction is presented in Fig. 4.4. During the more powerful event, the tracked CNR of GPS L1 C/A signals exhibited drops as large as 6 dB, while the weaker event produced drops up to 1.5 dB.

4.3 Interference Properties

This section details the transient space-based interference's temporal, spatial, and spectral properties over the seven-year period from January 2019

to April 2026. For comparison, it also presents data from a naturally occurring solar radio burst.

4.3.1 Temporal Patterns

High-rate (1-Hz) data from 165 reference stations collected over the seven-year period were retrieved and analyzed. The detection hypothesis test described earlier was applied to every station on an epoch-by-epoch basis from January 1, 2019 to May 4, 2026. If interference was detected simultaneously at several stations, then an event detection was declared. A total of 75 days were identified with at least one wide-area transient GNSS interference event on the GPS L1 frequency during which at least one station experienced a CNR drop of 5 dB or greater. The annual number of such occurrences is summarized in Table 4.1. The earliest detection of a significant transient wide-area interference event within this period occurred in October 2019. Fig. 4.5 presents a histogram showing the distribution of high-power interference events by day of the week and hour of the week. Notably, these predominantly occurred during business days and business hours (UTC time), which suggests human involvement: a purely random phenomenon would tend to exhibit a temporally uniform distribution.

There were an 47 additional days with weaker widespread transient interference events. On these days, the largest CNR drop did not exceed 5 dB. If these days are included in Fig. 4.5, the distributions shown do not change appreciably. The weaker interference events are left out of this chapter's

Table 4.1: The number of days per year that saw at least one wide-area transient GNSS interference event on the GPS L1 frequency with at least one station suffering a drop of 5 dB or greater.

	2019	2020	2021	2022	2023	2024	2025	2026	Total
Count	8	17	5	7	15	9	12	2	75

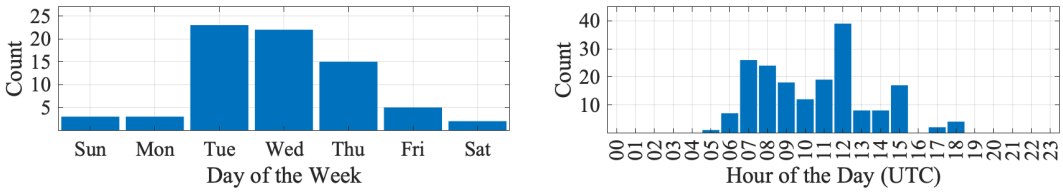


Figure 4.5: Distribution of the day of the week and hour of day (with respect to UTC) during which interference events with at least one station suffering a drop of 5 dB or greater occurred. Clearly, the high-power interference events typically occur during business days and business hours.

analysis to avoid false alarms at the periphery of the affected area.

Fig. 4.6 presents the number of daily detections across three additional example days. Unlike the case in Fig. 4.3, which shows a single high-power interference event, some days experienced multiple high-power events. Notably, day 146 of 2021 exhibited an interference signature similar to that of day 160 of the same year, with a low-power burst followed by a high-power burst. This happened twice on day 146, with the strong interference events being separated by approximately 32.6 minutes. The time delay between the low-power and high-power burst was 317 seconds on day 160, and 115 seconds for both events on day 146. This signature was observed on several other days. Day 204 of 2020 and day 014 of year 2025 recorded numerous high-power interference bursts. The timing of the bursts within a single day is typically

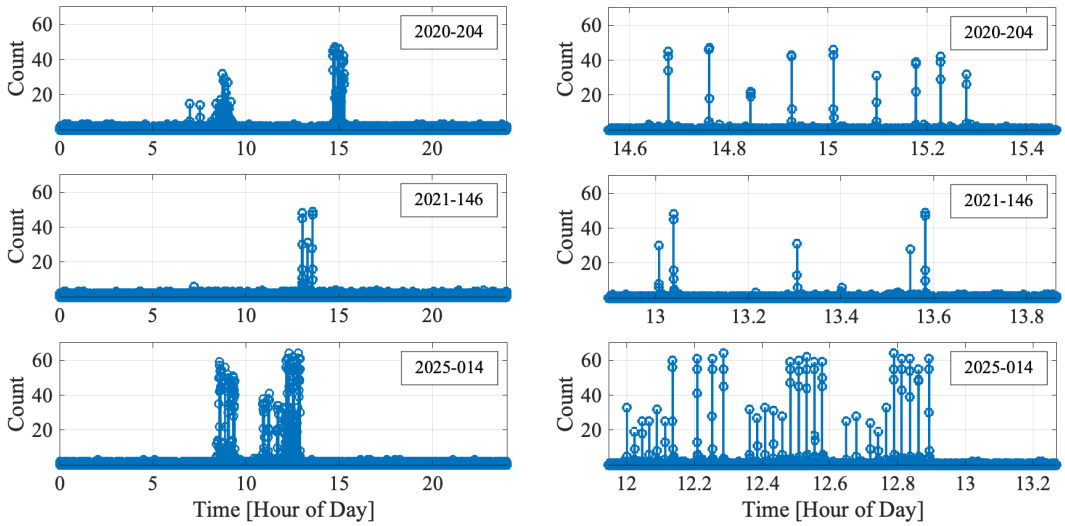


Figure 4.6: The number of stations that detected interference during three different days. Interference events can occur several times a day. The left column shows the entire day, while the right column is an expanded sub-interval.

periodic, with large bursts often spaced by an integer multiple of 150 seconds. The daily temporal patterns shown here of the high-power interference bursts are broadly representative of the phenomenon across all 75 days.

4.3.2 Spatial Patterns

Across all 75 days on which high-power events occurred, GNSS receivers in Europe were the most affected, with the Baltic region consistently experiencing the largest CNR drops. Fig. 4.7 shows example heat maps from day 146 of 2021 and day 014 of 2025. On day 146, two high-power interference events occurred, both producing nearly identical spatial patterns. Day 014 of 2025 also saw multiple high-power bursts, all with consistent patterns matching the

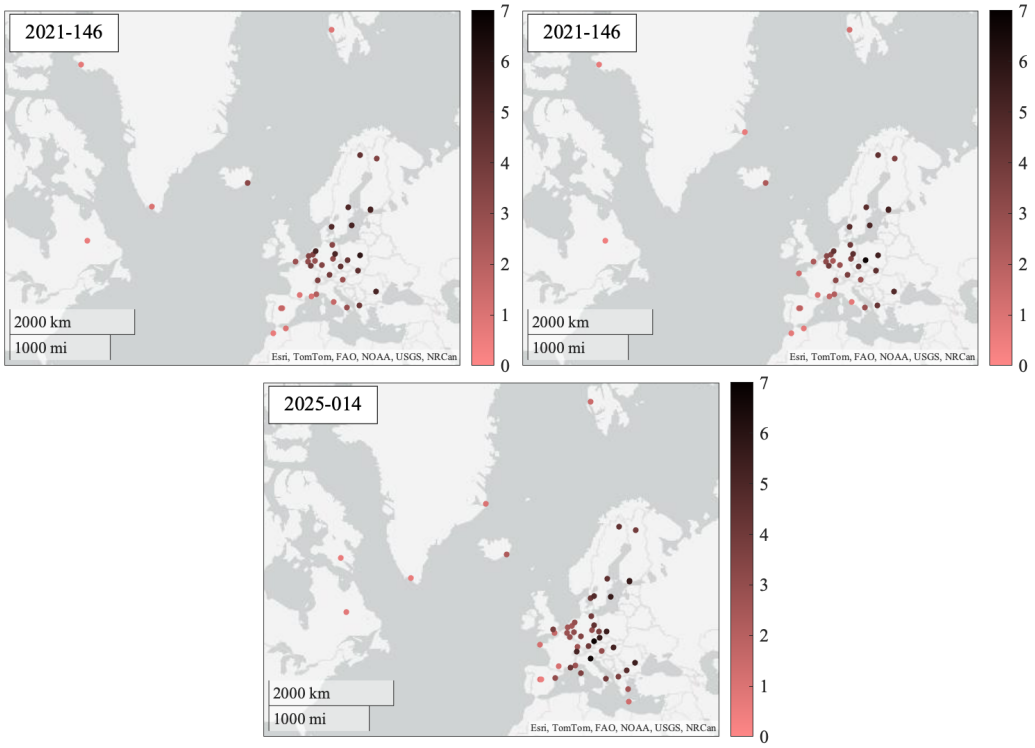


Figure 4.7: Detection statistic heat maps, in dB, for day 146 of year 2021 and day 014 of year 2025. GNSS receivers in the Baltic region are impacted most, which is generally representative of nearly all of the interference events.

example shown. The largest CNR drop across all events was 10 dB recorded at the LAMA station in Poland in 2025. Notably, during wide-area interference events in Europe, no similar disruptions were detected elsewhere in the world.

Although the overwhelming majority of interference events saw receivers in the Baltic region impacted most, day 204 of 2020 exhibited a distinct interference pattern compared to other days. On this day, there was progressive movement in the geographic center of interference over multiple events, starting in the Baltic Sea and then moving into Germany and on to the Norwe-

gian Sea, all over a 20-minute interval [31]. This deviation may be attributed to satellite motion, to a change in the interference source's beam pointing vector, or to multiple active satellite sources.

4.3.3 Spectral Properties

Spectral data were collected using a u-blox F9P GNSS receiver connected to a Trimble GNSS antenna installed in Gdynia, Poland, operating as an experimental radio frequency interference (RFI) monitoring station. Direct estimation of the absolute received interference power at the antenna reference point is not possible because of the unknown loss between the antenna and the receiver front end. Accordingly, the analysis here focuses on relative comparisons of the PSD across interference events and against interference-free baseline conditions. The u-blox receiver provides uncalibrated 1-Hz spectral observations through the "SPAN" message. Although SPAN output measurements are dimensionless, they are nonetheless suitable for comparative analysis. In previous work, a methodology was developed to transform these uncalibrated SPAN observations into a spectrally adjusted power density metric [83].

Fig. 4.8 presents the raw SPAN spectral output under nominal (RFI-free) conditions (black line) and during 48 strong interference events from 2024 to 2025. The interference spectrum exhibits a consistent spectral shape across all recorded events on days for which data are available. The interference peak is centered at 1577.5 MHz, about 2 MHz above the GPS L1 center frequency of 1575.42 MHz, and has an approximate bandwidth of 5 MHz.

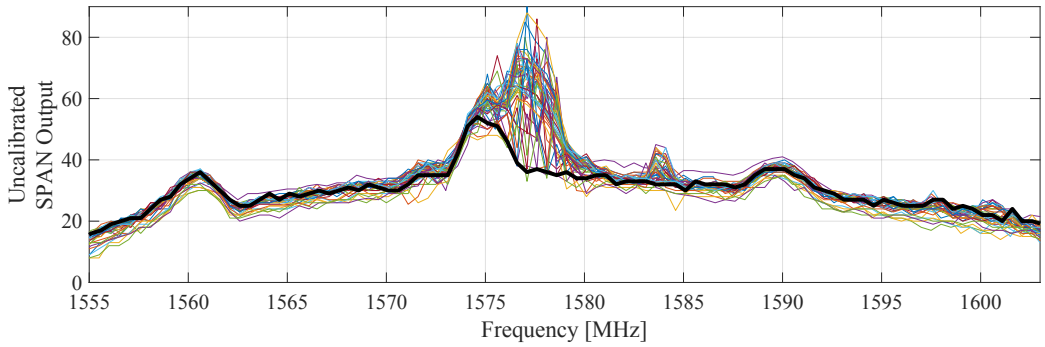


Figure 4.8: Uncalibrated average PSD near the GPS L1 band during nominal operation (thick black line) and 1-Hz PSD estimates during 48 transient interference events, as recorded in Gdynia, Poland.

In addition to tracked GPS L1 C/A signals, tracked Galileo E1 and BeiDou B1C/B1A signals also exhibited a concurrent drop in CNR during interference events. The magnitude of the decrease at each station closely matched the drop of the GPS L1 C/A signals, which might be expected because they share the same center frequency. The drop in CNR is not identical because the GPS, Galileo, and BeiDou spreading codes each have a different PSD $\tilde{S}^i_{\mathrm{C}}(f)$ which, as shown in (4.6), leads to a different interference component PSD $\tilde{S}^i_{\mathrm{I}}(f)$ and ultimately to a different interference power density I^i_0.

During strong 1577.5-MHz-centered interference events, a small but noticeable drop in CNR is also observed on tracked BeiDou B1I signals, which are centered at 1561.098 MHz. This is a puzzling because the interference spectrum shown in Fig. 4.8 has no spectral overlap with the BeiDou B1I signal, from which it follows that $I^i_0 = S^i_{\mathrm{I}}(0) = 0$. Apparently, the method by which receivers in the IGS network estimate the noise floor N^i_0 that they use to report

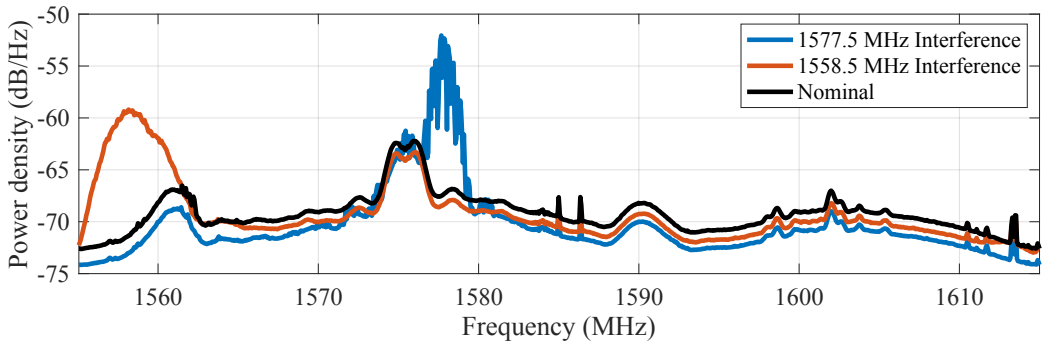


Figure 4.9: Power spectrum derived from raw wideband samples captured in Amsterdam, Netherlands, during an interference event on February 11, 2026. An initial burst of interference at 1577.5 MHz was followed by a burst at 1558.5 MHz.

CNR values is somewhat sensitive to interference near L1.

Occasionally, an initial CNR drop and recovery in tracked GPS L1 C/A signals was followed by an equal-magnitude CNR drop and recovery in tracked BeiDou B1I signals. This pattern was observed on 15 of the 75 days, the first in June 2020. This indicates that the interference source can generate signals at 1577.5 MHz and near 1561.098 MHz. Fig. 4.8 only shows the interference centered at 1577.5 MHz, but Fig. 4.9, which shows PSDs derived from raw wideband samples captured in February 2026 by a receiver in the Netherlands clearly shows both interference bands, with the lower band centered at 1558.5 MHz. Interestingly, these two interference bands never appear to be active at the same time.

No wide-area transient interference has yet been observed near the GPS L2 or L5 bands.

4.3.4 Comparison with Solar Radio Burst

It is worth noting that naturally occurring phenomena can also cause significant CNR reduction over a large geographic area. [18, 19] examined the impact of intense solar radio bursts on tracked GNSS signals using observational data from the IGS network during solar events in the mid 2000s. It was demonstrated that solar radio bursts can cause CNR degradation by as much as 25 dB across the sunlit side of Earth. But solar radio bursts are qualitatively different from the transient interference studied in this chapter: They are typically broadband, evolve more slowly than the rapid onset observed in Fig. 4.3, and the associated CNR reduction typically persists for a longer duration.

On November 11, 2025, a strong solar flare occurred with a magnitude of X5.1, causing a geomagnetic storm on Earth reaching category G4 ("severe") [40]. IGS stations positioned on the sunlit side of Earth experienced drops in CNR on all GNSS frequencies. Fig. 4.10 shows the CNR time history of tracked GPS L1, L2, and L5 signals from the SUTM reference station in South Africa. Signals on L2 and L5 were degraded as much as 17 dB for hundreds of seconds. The CNR of tracked signals at other IGS stations were degraded in a similar manner. Fig. 4.10 shows the magnitude of the CNR drop on the tracked L2 signals across the globe. Clearly, the effects of solar radio bursts manifest differently in the IGS data compared to the transient phenomenon studied here.

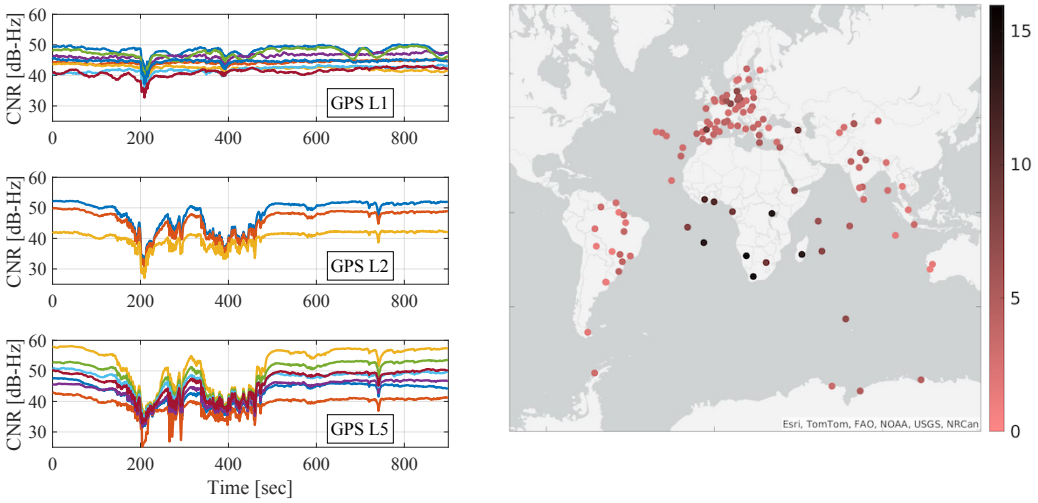


Figure 4.10: Left: CNR time history of tracked GPS L1, L2, and L5 signals produced by IGS station SUTM in South Africa during the solar radio burst on November 11, 2025. Right: Heat map of maximum CNR degradation at GPS L2 during the solar radio burst.

4.4 Elevation-Mask-Based Interference Source Identification

A simple technique for winnowing the list of possible candidate satellites causing interference is to determine which satellites were overhead the affected region during a given interference event. Two-Line Elements (TLEs) for publicly tracked orbiting objects can be obtained from space-track.org, a database maintained by the United States Space Force to support spaceflight safety. TLEs can be used to compute the approximate position of an object at any time reasonably close to the TLE epoch. While the accuracy of these estimates is roughly 1 km at epoch and degrades with the age of the TLE [80,97], their precision is more than sufficient for a preliminary candidacy test.

Let $\mathcal{S}$ be the set of all space objects with public TLEs, and $\mathcal{I}_\mathrm{d} \subseteq \mathcal{I}$ be the set of all reference stations that detected a given interference event. The elevation angle of satellite $s \in \mathcal{S}$ as seen from the GNSS antenna at station $i \in \mathcal{I}_\mathrm{d}$ is denoted α_{is}. This can be calculated assuming an ellipsoidal model for Earth such as WGS84. Assuming an elevation mask α_0, satellite $s \in \mathcal{S}$ is considered a valid candidate if $\alpha_{is} \geq \alpha_0$ for all $i \in \mathcal{I}_\mathrm{d}$.

Under this framework, one can readily determine the minimum altitude at apogee for candidate satellites. Let $\boldsymbol{r} \in \mathbb{R}^3$ be an arbitrary position in Earth-centered, Earth-fixed (ECEF) coordinates, and $\alpha_i(\boldsymbol{r})$ be its elevation angle as seen from station $i \in \mathcal{I}_\mathrm{d}$. The minimum-radius position $\boldsymbol{r}^*$ satisfying all elevation mask constraints may be found by solving

$$\boldsymbol{r}^* = \underset{\boldsymbol{r} \in \mathbb{R}^3}{\operatorname{argmin}} \|\boldsymbol{r}\| \quad \text{such that} \quad \alpha_i(\boldsymbol{r}) \geq \alpha_0 \quad \text{for all} \quad i \in \mathcal{I}_\mathrm{d} \tag{4.18}$$

A solution, which is unique for $\alpha_0 \geq 0$, can be obtained straightforwardly using a numerical optimizer. The minimum satellite altitude at apogee is then given by $\zeta^* = \|\boldsymbol{r}^*\| - r_\mathrm{E}$, where r_E is Earth's equatorial radius.

One can also construct the feasible region of satellite positions, which corresponds to the intersection of all interior ($\alpha_0 > 0$) or exterior ($\alpha_0 < 0$) points of each reference station's feasibility cone. In the special case of $\alpha_0 = 0$, feasible region construction becomes a nonlinear optimization problem with linear constraints formulated by enforcing $\boldsymbol{r} \in \mathbb{R}^3$ to be above the tangent plane at each reference station. The general case is similar but with nonlinear constraints—those in (4.18).

Table 4.2: The minimum altitude and number of satellites that satisfy various elevation masks for the event on day 160 of 2021.

α_0	-5°	-3°	0°	3°	5°	10°
ζ^* [km]	800	957	1,212	1,468	1,655	2,173
Number of valid satellites	381	303	201	145	113	87

One might assume that the reference stations not detecting interference could be used to further constrain the feasible region. However, this assumption does not hold when the antenna gain pattern of the interference source is unknown. For example, if the interference source is equipped with a narrow-beam antenna, many reference stations may not observe the interference even if the satellite satisfies the elevation mask constraint.

Fig. 4.11 illustrates the positions of all publicly tracked objects during the high-power interference burst on day 160 of 2021. Its right figure shows the subset of satellites satisfying $\alpha_0 = 0$, excluding known debris and rocket bodies. The interior of the quasi-conical red surface represents the feasible region in which the interference satellite could have been located. Table 4.2 shows ζ^* for various elevation mask angles for this event, as well as the number of satellites that satisfy the elevation mask.

If one assumes that a single satellite is responsible for the observed interference events, the list of candidate satellites can be further narrowed by eliminating those that do not satisfy the elevation mask across multiple events. As shown later, however, the single-satellite assumption does not hold for this chapter's transient interference.

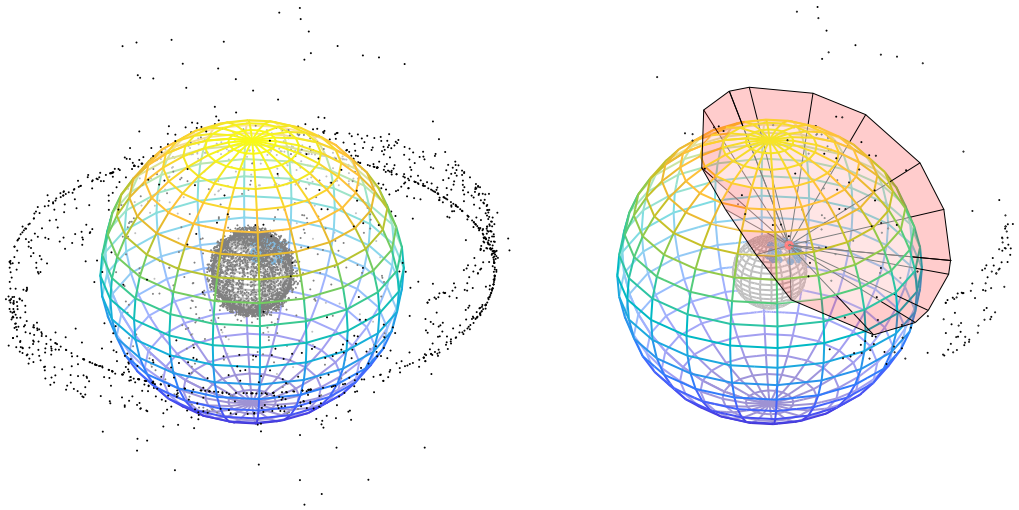


Figure 4.11: Left: The position of all tracked objects during the high-power interference burst on day 160 year 2021. Right: The position of all satellites that satisfy a 0° elevation mask, excluding debris and rocket bodies. The feasible region in which the interference source could have been positioned is interior to the red surface. For reference, the colored spherical shell corresponds to medium Earth orbit (20,000 km altitude).

4.5 Multi-Hypothesis CNR-Based Interference Source Identification

This section develops a more rigorous CNR-based technique to associate measurements to a satellite than the simple elevation-mask-based method of the previous section.

4.5.1 Generalized Likelihood Ratio Test

A binary Neyman-Pearson hypothesis testing framework provides the optimal test for simple cases in which the relevant probability density functions under both the null (H_0) and alternative (H_1) hypotheses are fully specified.

The test is optimal in the sense of maximizing the probability of detection for a fixed probability of false alarm. In real-world scenarios, however, some parameters are often unknown, resulting in composite tests. If the unknown parameters' *a priori* probability distributions are also unknown, then no known optimal test exists [132]. Such conditions apply to this chapter's attempt to identify the source of space-based interference using received-power measurements because parameters such as the source's transmit power are unknown.

The Generalized Likelihood Ratio Test (GLRT) is often employed for composite detection problems with unknown prior distributions. It amounts to a two-step procedure. First, the unknown parameter vector $\boldsymbol{x}$ is estimated from the measurements $\boldsymbol{z}$ via maximum likelihood estimation under each hypothesis. Second, these estimates are used to evaluate likelihoods as if they were the correct values of $\boldsymbol{x}$ under the respective hypotheses [132]. In the current context, the hypothesis test is M-ary rather than binary, with H_s being the hypothesis that satellite $s \in \mathcal{S}$ is the interference source. The vector $\boldsymbol{x}_s$ contains parameters assumed unknown *a priori* that are used to model the candidate satellite's transmit power, antenna gain pattern, and pointing direction (beam vector), or some subset of these.

4.5.2 Simulation Study

The performance of a multi-hypothesis satellite association technique may be illustrated by a simulation study. This section's study applies an M-ary version of the GLRT not to the interference source identification problem itself

but to a closely related problem: identifying which satellite is the source of a GNSS signal whose CNR measurements are taken from a network of reference stations. This related problem is chosen because (1) it yields insights pertinent to the interference identification problem, and (2) the technique developed and demonstrated in simulation can be tested on real-world CNR measurements from the IGS network to demonstrate identification of a known source GNSS satellite.

For the ith station and jth GNSS signal at epoch k, let $z_{ij}[k]$ be an interference-free CNR measurement, modeled as in (4.8). The measurement noise variance $\sigma_{ij}^2[k] = f(\theta_{\mathrm{R}}^{ij}[k])$ is modeled as a nonlinear function of the received zenith angle, with $\sigma_{ij} = 0.25$ dB at zenith ($\theta_{\mathrm{R}}^{ij}[k] = 0$) and $\sigma_{ij} = 1.25$ dB at the horizon ($\theta_{\mathrm{R}}^{ij}[k] = 90°$). Both receiver and transmitter antenna gain patterns are modeled as azimuthally symmetric—$G_{\mathrm{R}}^{i}(\theta, \phi) = G_{\mathrm{R}}^{i}(\theta)$, $G_{\mathrm{T}}^{j}(\theta, \phi) = G_{\mathrm{T}}^{j}(\theta)$—which is approximately true for both GNSS receiver and transmitter antennas. Fig. 4.12 shows the assumed reference station locations (black dots) and the position of an exemplar GNSS satellite for 16 time epochs spaced by 15-minute intervals. It also shows the corresponding transmitter and receiver gain patterns. In this study, all transmitter antennas are nadir-pointing, and all reference stations are assumed to utilize either Leica AR20 or AR25 antennas.

At epoch k, let $\mathcal{I}_j[k] \subseteq \mathcal{I}$ be the set of participating stations tracking the jth GNSS signal, and $\mathcal{S}_{\mathrm{e}}[k]$ be the set of candidate source satellites that have all passed elevation mask vetting. Let the measurement vector $\boldsymbol{z}_j[k]$ contain

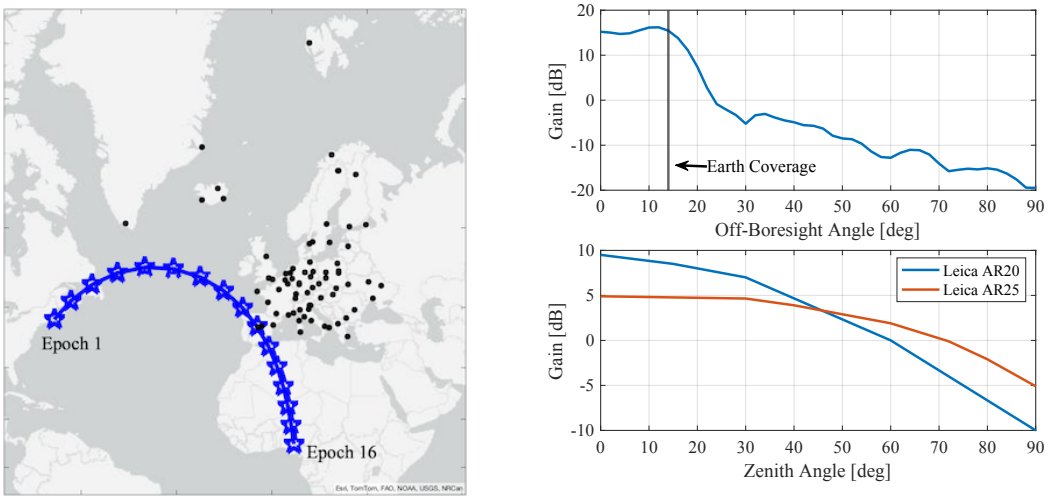


Figure 4.12: Simulation study setup. Left: The geometry of the true source GNSS satellite over 16 epochs at 15-minute intervals, with terrestrial reference stations shown as black dots. Right: Transmitter (top) and receiver (bottom) gain patterns. The receiver gain patterns are from a model provided by the manufacturer and verified empirically.

ordered CNR measurements of signal $j \in \mathcal{J}$ from those receivers tracking the signal at epoch k, i.e., measurements from the set $\{z_{ij}[k] \mid i \in \mathcal{I}_j[k]\}$.

A GLRT-based association test is used to assess whether a candidate satellite $s \in \mathcal{S}_{\mathrm{e}}[k]$ is consistent with $\boldsymbol{z}_j[k]$ given varying degrees of prior knowledge about the satellite's transmitter. The unknown parameter vector $\boldsymbol{x}_s[k]$ may include any combination of transmit power $P_s[k]$, beamwidth $\beta_s[k]$, and beam vector $\boldsymbol{v}_s[k] = [\theta_s[k],\ \phi_s[k]]^{\mathsf{T}}$, where $\theta_s[k]$ and $\phi_s[k]$ are the off-boresight and azimuth angle relative to the nadir vector and a reference azimuth vector in the satellite frame. In the fully unknown case, $\boldsymbol{x}_s$ includes all three

parameters:

$$\boldsymbol{x}_s[k] = [P_s[k]\,,\ \beta_s[k]\,,\ \boldsymbol{v}_s^{\mathsf{T}}[k]]^{\mathsf{T}} \tag{4.19}$$

Under hypothesis H_s, satellite $s \in \mathcal{S}_{\mathrm{e}}[k]$ is the source of signal $j \in \mathcal{J}$, in which case the measurement vector $\boldsymbol{z}_j[k]$ is distributed as

$$H_s: \quad \boldsymbol{z}_j[k] \sim \mathcal{N}(\bar{\boldsymbol{z}}_j[k, \boldsymbol{x}_s],\ R_j[k, \boldsymbol{x}_s]), \quad s \in \mathcal{S}_{\mathrm{e}}[k] \tag{4.20}$$

Here, $\bar{\boldsymbol{z}}_j[k, \boldsymbol{x}_s]$ is the vector of expected CNR measurements for signal $j \in \mathcal{J}$ emanating from satellite $s \in \mathcal{S}_{\mathrm{e}}[k]$, and $R_j[k, \boldsymbol{x}_s]$ is a square diagonal matrix with diagonal elements $\sigma_{ij}^2[k]$ for $i \in \mathcal{I}_j[k]$, where the angle $\theta_{\mathrm{R}}^{ij}[k]$, on which $\sigma_{ij(s)}^2[k]$ depends, is calculated assuming signal j comes from satellite s.

This association problem amounts to a multi-hypothesis composite general Gaussian problem [132]. Hypotheses may be compared based on their negative log likelihood function evaluated at the maximum-likelihood value for the parameter vector $\boldsymbol{x}_s$, which, for the model in (4.20), leads to a cost function of the form

$$J_j[k, \boldsymbol{x}_s^*] = (\boldsymbol{z}_j[k] - \bar{\boldsymbol{z}}_j[k, \boldsymbol{x}_s^*])^{\mathsf{T}} (R_j[k, \boldsymbol{x}_s^*])^{-1} (\boldsymbol{z}_j[k] - \bar{\boldsymbol{z}}_j[k, \boldsymbol{x}_s^*]), \quad s \in \mathcal{S}_{\mathrm{e}}[k] \tag{4.21}$$

where $\boldsymbol{x}_s^*$ is the value of $\boldsymbol{x}_s$ that minimizes $J_j[k, \boldsymbol{x}_s]$. Thus, for $a, s \in \mathcal{S}_{\mathrm{e}}[k]$, if $J_j[k, \boldsymbol{x}_a^*] < J_j[k, \boldsymbol{x}_s^*]$, then H_a is more likely than H_s. Since $\mathcal{S}_{\mathrm{e}}[k]$ is assumed to include the true source for signal j, then for the case of hypotheses with equally likely priors, the probability that H_s is true is given by [6]

$$P(s|\boldsymbol{z}_j[k]) = \frac{|R_j[k, \boldsymbol{x}_s^*]|^{-1/2} \exp\left(-\frac{1}{2} J_j[k, \boldsymbol{x}_s^*]\right)}{\sum_{a \in \mathcal{S}_{\mathrm{e}}[k]} |R_j[k, \boldsymbol{x}_a^*]|^{-1/2} \exp\left(-\frac{1}{2} J_j[k, \boldsymbol{x}_a^*]\right)}, \quad s \in \mathcal{S}_{\mathrm{e}}[k] \tag{4.22}$$

Based on this expression, one could choose the highest-probability satellite, or the top N most probable, or all those for which $P(s|\boldsymbol{z}_j[k])$ exceeds a threshold. Nearly equivalent to the latter approach, one may retain all $s \in \mathcal{S}_\mathrm{e}[k]$ for which $J_j[k, \boldsymbol{x}_s^*] \leq \nu_j[k]$ for a given threshold $\nu_j[k]$. A principled choice for $\nu_j[k]$ can be obtained by limiting the probability $P(J_j[k, \boldsymbol{x}_u^*] > \nu_j[k])$ to a specified size, where $u \in \mathcal{S}_\mathrm{e}$ is the true source satellite. It is straightforward to show that if $\boldsymbol{x}_u^*$ is an accurate estimate then $J_j[k, \boldsymbol{x}_u^*]$ is distributed as a chi-square random variable with $|\mathcal{I}_j[k]|$ degrees of freedom:

$$J_j[k, \boldsymbol{x}_u^*] \sim \chi^2_{|\mathcal{I}_j[k]|} \tag{4.23}$$

From this one can readily calculate $\nu_j[k]$ to ensure $P(J_j[k, \boldsymbol{x}_u^*] > \nu_j[k])$ is below a chosen value.

4.5.2.1 Simulation

1000 Monte Carlo trials were conducted at each time epoch depicted in Fig. 4.12. A trajectory for GPS PRN 15 from February 2020 was chosen for the true source satellite s. At each epoch, (1) a subset $\mathcal{S}_\mathrm{e}[k]$ of candidate satellites was generated by applying an elevation mask $\alpha_0 = 0$; (2) the subset $\mathcal{I}_j[k]$ was determined for signal j; (3) the noise-free CNR_{ij} at each reference station $i \in \mathcal{I}_j[k]$ was computed based on the true parameter vector $\boldsymbol{x}_s$ under the model (4.1); (4) a realization of the measurement $z_{ij}[k]$ was generated according to (4.8) assuming $\sigma_{ij}^2[k] = f(\theta_\mathrm{R}^{ij}[k])$ for each $i \in \mathcal{I}_j[k]$.

Costs $J_j[k, \boldsymbol{x}_s^*]$ were calculated for all $s \in \mathcal{S}_\mathrm{e}[k]$, with $\boldsymbol{x}_s$ estimated under five scenarios reflecting varying levels of *a priori* knowledge of the transmitter:

(S1) all parameters known; (S2) unknown transmit power P_s; (S3) unknown P_s and beam width β_s; (S4) unknown P_s and beam vector $\boldsymbol{v}_s$; and (S5) unknown P_s, β_s, and $\boldsymbol{v}_s$. The maximum likelihood (minimum cost) estimates β_s^*, and $\boldsymbol{v}_s^*$ were found via grid search, while P_s^* was taken as the mean of the elements of the difference $\boldsymbol{z}_j[k] - \bar{\boldsymbol{z}}_j[k, \tilde{\boldsymbol{x}}_s^*]$, with $\tilde{\boldsymbol{x}}_s^* = [P_s = 0\,,\, \beta_s^*\,,\, \boldsymbol{v}_s^{*\mathsf{T}}]^{\mathsf{T}}$.

4.5.2.2 Results and Discussion

Fig. 4.13 shows the number of satellites that, for at least 5% of the Monte Carlo trials, were considered viable, i.e., had costs $J_j[k, \boldsymbol{x}_s^*]$ below $\nu_j[k]$ when $P(J_j[k, \boldsymbol{x}_u^*] > \nu_j[k])$ was limited to under 10^{-3}. As shown, the analysis was performed across various epochs, transmitter scenarios, and receiver antenna types. The results reveal that the GLRT-based multi-hypothesis technique is far superior to multi-station elevation masking in narrowing the number of candidate satellites. When transmitter power, beamwidth, and beam vector are known *a priori*, only the true satellite and a small number of others are consistent with the data at any given epoch, as shown by the dark blue trace along the bottom of both panels in Fig. 4.13. Even when none of these parameters is known *a priori* and so they must be estimated from the network data (green traces in Fig. 4.13), the method offers useful discrimination. But it is clear that, at the geometries arising during many of the epochs, the CNR-based association method alone would not be sufficient for confident single-epoch satellite identification when there are unknown parameters.

The two panels in Fig. 4.13 also show that the receiver antenna gain

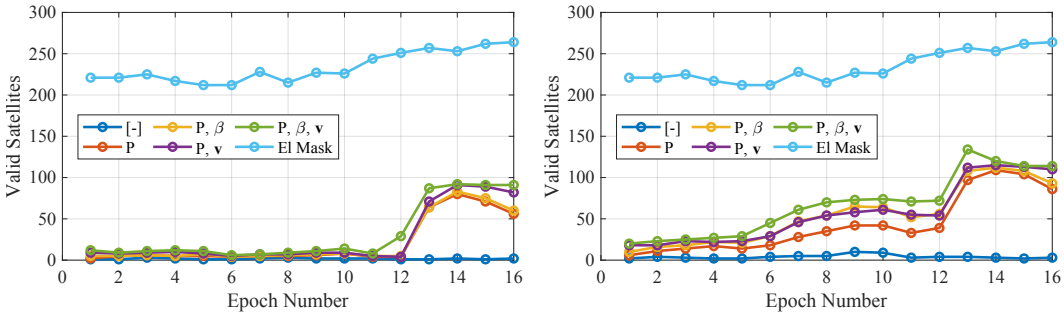


Figure 4.13: The number of satellites that, for at least 5% of the Monte Carlo trials, remained viable candidates after association testing with $P(J_j[k, \boldsymbol{x}_s^*] > \nu_j[k]) < 10^{-3}$. The legends indicate which parameters are assumed unknown. The number of satellites satisfying all elevation masks at each epoch, $|\mathcal{S}_e[k]|$, is shown in light blue. All reference stations modeled with Leica AR20 antennas (Left)or Leica AR25 antennas (right).

pattern substantially influences the test's discrimination power. Specifically, the Leica AR20 (left panel) exhibits a more pronounced variation in gain over the zenith angle compared to the Leica AR25 (right panel), thereby enhancing the sensitivity of measurements to changes in zenith angle and facilitating the exclusion of false associations.

The impact of transmitter satellite geometry on correct association is also evident. Note how the number of viable satellites increases when the true satellite approaches the equator during the last few epochs. Fig. 4.14 offers an explanation: Compared to epoch $k = 1$, which is far from the equator and saw relatively few false associations, at epoch $k = 16$ many more satellites are falsely associated because a whole band of geostationary satellites are also consistent with the data when any transmitter parameters are unknown *a priori*.

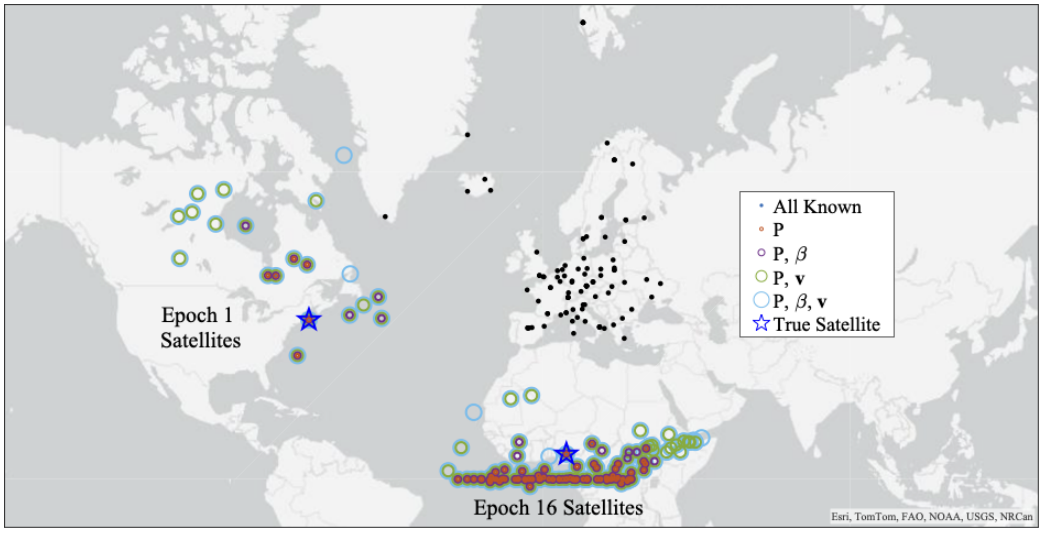


Figure 4.14: Satellites that remained viable candidates on more than 5% of Monte Carlo trials for epoch $k = 1$ and epoch $k = 16$ when the reference stations are modeled as using the Leica AR25 antenna. The legend gives symbols for the true satellite and for valid candidates under scenarios with an increasing number of unknown transmitter parameters.

4.5.3 GNSS Satellite Identification with Real CNR Data

The GNSS satellite identification technique developed here and demonstrated by simulation was also tested using real CNR measurements taken from reference stations in $\mathcal{I}$. Let $\mathcal{J}_\mathrm{g} \subset \mathcal{J}$ be the subset of GPS L1 C/A signals. Based on historical data from each site $i \in \mathcal{I}$ and on the known azimuthally symmetric antenna gain pattern $G_\mathrm{T}^j(\theta)$ and transmit power P_T^j for $j \in \mathcal{J}_\mathrm{g}$ from a select number of GPS satellites $\mathcal{S}_\mathrm{g} \subset \mathcal{S}_\mathrm{e}$, empirical estimates of the combined receiver antenna gain pattern and receiver noise floor

$$G_{\mathrm{R}0}^i(\theta) = G_\mathrm{R}^i(\theta) - N_0^i \tag{4.24}$$

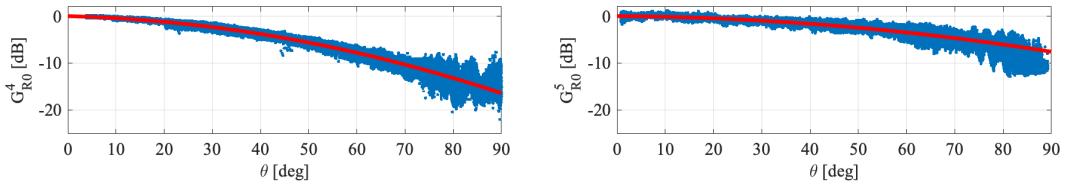


Figure 4.15: Normalized estimates of $G^i_{\mathrm{R0}}(\theta)$ for stations MEDI (Leica AR20, $i = 4$) and NICO (Leica AR25, $i = 5$). The blue dots represent collected values of $\mathrm{CNR}_{ij} - G^j_{\mathrm{T}} - L_{ij}$, while the red line indicates a weighted third-order polynomial fit, which becomes $G^i_{\mathrm{R0}}(\theta)$, all normalized to $G^i_{\mathrm{R0}}(0) = 0$. As expected, measurement noise increases significantly as the elevation angle decreases.

were obtained for all $i \in \mathcal{I}$ across zenith angles $\theta \in [0, 90°]$, assuming azimuthal symmetry. Accurate estimates of $G^j_{\mathrm{T}}(\theta)$ for $j \in \mathcal{J}_{\mathrm{g}}$ and $s \in \mathcal{S}_{\mathrm{g}}$ are available in public documents [91], and corresponding estimates for P^j_{T} were obtained by observation with a calibrated high-gain antenna. Normalized estimates of $G^i_{\mathrm{R0}}(\theta)$ for stations MEDI (Leica AR20) and NICO (Leica AR25) are shown in Fig. 4.15.

The GLRT-based multi-hypothesis association test was then applied based on measured CNR_{ij} values for $i \in \mathcal{I}$ and $j \in \mathcal{J}_{\mathrm{g}}$ for $s \in \mathcal{S}_{\mathrm{g}}$, according to (4.8) with $\sigma^2_{ij} = f_i(\theta^{ij}_{\mathrm{R}})$. The site-specific function f_i was determined by the zenith-angle-dependent spread in the empirical data underlying each station's antenna model (e.g., the blue dots in Fig. 4.15). The method was tested under full *a priori* knowledge of the transmit power, beamwidth, and beam vector of $s \in \mathcal{S}_{\mathrm{g}}$, and under a case where the transmit power P_s was unknown. The results of this empirical study were qualitatively consistent with those of the simulation study: the true $s \in \mathcal{S}_{\mathrm{g}}$ was often uniquely identifiable in

a single epoch when its transmitter parameters were all known *a priori*, but false associations increased in geometry-dependent ways when P_s had to be estimated from the data.

4.5.4 Implications for CINR-Based Interference Source Identification

Based on CNR and CINR measurements alone, interference source identification is more challenging than GNSS satellite identification in two key respects. First, the transmitter gain pattern may be wholly unknown, not just unknown in beamwidth β_s as in the simulation study. Second, and more importantly, the receiver antenna gain pattern $G_{\mathrm{R}}^{i}(\theta)$ and noise floor N_0^i must be known independently, rather than combined as in (4.24), because the model for CINR_{ij} given in (4.4) involves an isolated value of $\tilde{N}_0^i$. Yet it is not possible to independently estimate $G_{\mathrm{R}}^{i}(\theta)$ and N_0^i for $i \in \mathcal{I}$ strictly from historical CNR measurements because CNR is a linear combination of the two quantities, as evident in (4.1) and (4.2). Separate estimates of N_0^i can be obtained for $\mathcal{I}_{\mathrm{a}} \in \mathcal{I}$, the subset of reference stations for which an accurate antenna gain pattern $G_{\mathrm{R}}^{i}(\theta)$ is available, e.g., one provided by the manufacturer. But the authors of the current chapter found $\mathcal{I}_{\mathrm{a}}$ was limited to stations with Leica AR20 or AR25 antennas. For all other stations (all $i \in \mathcal{I} \setminus \mathcal{I}_{\mathrm{a}}$), either $G_{\mathrm{R}}^{i}(\theta)$ was not available or the available model did not match the shape of the empirically derived combined model $G_{\mathrm{R0}}^{i}(\theta)$. (Note that, whereas the IGS provides extremely accurate models of antenna phase center variation as a function of θ and ϕ [84], it provides no such models for antenna gain.) Importantly, the

geographic distribution of $\mathcal{I}_{\mathrm{a}}$ is not as extensive as that shown in Fig. 4.12, leading to weaker observability of the unknown transmitter parameters, and thus weaker discrimination of candidate satellites.

In view of these challenges, one can expect the current performance of CNR-based multi-hypothesis association for interference source identification to fall somewhere between the light blue and green curves shown in Fig. 4.13, tending towards the light blue curve when the unknown parameters multiply. It becomes clear that, to confidently identify a GNSS interference source, more powerful techniques are required. Nonetheless, CNR-based multi-hypothesis association could offer useful prior probabilities for other interference source identification techniques if accurate gain patterns $G_{\mathrm{R}}^{i}(\theta)$ were available for more stations in $\mathcal{I}$. Moreover, many stations provide CNR measurements only as integer-quantized values. Reporting these measurements at a higher resolution would significantly enhance measurement accuracy.

Assuming these desiderata for IGS stations will eventually be fulfilled, the mapping from the GNSS satellite identification problem to the interference source identification problem is briefly outlined. First, supposing that $G_{\mathrm{R}}^{i}(\theta)$ is known, and leveraging the known transmit gain pattern $G_{\mathrm{T}}^{j}(\theta)$ and transmit power P_{T}^{j} for $j \in \mathcal{J}_{\mathrm{g}}$, N_0^i is estimated from historical CNR data for all $i \in \mathcal{I}$ so that the models (4.4) and (4.9) can be applied. Then, assuming that at epoch k there is an interference event, let $\mathcal{I}_{\mathrm{d}}[k] \subset \mathcal{I}$ be the subset of stations that detect the event, and let the vector $\boldsymbol{z}[k]$ contain ordered station-specific detection statistics at epoch k, i.e., values from the set $\{\Lambda_i[k] \mid i \in \mathcal{I}_{\mathrm{d}}[k]\}$.

For every satellite $s \in S[k]$ and for an assumed parameter vector $\boldsymbol{x}_s$, one can develop a hypothesis model for $\boldsymbol{z}[k]$ like the one in (4.20) and find the subset of $S[k]$ consistent with $\boldsymbol{z}[k]$, just as with the GNSS satellite identification problem. Let the resulting probabilities, analogous to $P(s|\boldsymbol{z}_j[k])$ in (4.22), be denoted $P(s|\boldsymbol{z}[k])$. The following section shows how these probabilities can act as prior probabilities for a more powerful identification technique.

4.6 TDOA-Based Interference Source Identification

This section presents a satellite association framework based on TDOA measurements, together with a TDOA analysis based on recently captured raw in-phase and quadrature (IQ) samples from two spatially separated receivers in Europe during a continent-scale interference event in February 2026.

4.6.1 TDOA Measurement Models

Consider two receivers at fixed locations $\boldsymbol{g}_1$ and $\boldsymbol{g}_2$ in the ECEF reference frame. Let $\mathcal{K} = \{1, \ldots, K\}$ be the set of TDOA measurement epochs over an observation interval, $\mathcal{S}_\mathrm{e}$ be the set of candidate satellites that satisfy elevation mask vetting for all $k \in \mathcal{K}$, and $\boldsymbol{r}_s[k]$ be the ECEF position of satellite $s \in \mathcal{S}_\mathrm{e}$ at epoch k. For model simplicity, $\boldsymbol{r}_s[k]$ is assumed to remain constant over the signal's time-of-flight for each measurement. Denoting by $\rho_{is}[k] = \|\boldsymbol{g}_i - \boldsymbol{r}_s[k]\|$ the range between receiver $i \in \{1, 2\}$ and satellite $s \in \mathcal{S}_\mathrm{e}$, the modeled TDOA for s is the range difference between the two receivers: $\bar{y}_s[k] = \rho_{2s}[k] - \rho_{1s}[k]$. Assume that $u \in \mathcal{S}_\mathrm{e}$ is the true interference source.

Then the actual TDOA measurement is modeled as

$$y[k] = \bar{y}_u[k] + w[k], \quad k \in \mathcal{K} \tag{4.25}$$

where the AWGN noise term $w[k] \sim \mathcal{N}(0, \sigma^2)$ accounts for the aggregate effects of thermal noise and quantization error from nearest-neighbor sampling. Epoch-wise measurements are stacked as $\boldsymbol{y} = [y[1], \ldots, y[K]]^\mathsf{T}$, and modeled values for $s \in \mathcal{S}_\mathrm{e}$ as $\bar{\boldsymbol{y}}_s = [\bar{y}_s[1], \ldots, \bar{y}_s[K]]^\mathsf{T}$.

The location $\boldsymbol{r}_s[k]$ of each $s \in \mathcal{S}_\mathrm{e}$ is not known exactly but can be approximated with the public TLE data available for s. Let $\tilde{\boldsymbol{r}}_s[k]$ be the TLE-based orbit for s propagated to epoch k, and $\boldsymbol{e}_s$ be the error in $\tilde{\boldsymbol{r}}_s[k]$, with $\boldsymbol{e}_s \sim \mathcal{N}(\boldsymbol{0}, Q_s)$. Modeling $\boldsymbol{e}_s$ as constant is acceptable for observation intervals limited to a few seconds as in the current case. It is assumed that $\boldsymbol{e}_s$ and $w[k]$ are mutually independent for all $s \in \mathcal{S}_\mathrm{e}$ and $k \in \mathcal{K}$. The true satellite location is related to $\tilde{\boldsymbol{r}}_s[k]$ and $\boldsymbol{e}_s$ by

$$\boldsymbol{r}_s[k] = \tilde{\boldsymbol{r}}_s[k] + \boldsymbol{e}_s, \quad s \in \mathcal{S}_\mathrm{e}, \quad k \in \mathcal{K} \tag{4.26}$$

The framework established in [80, 97] can be leveraged to investigate the effects of $\boldsymbol{e}_s$ on the TDOA measurement. Let the range Jacobian of the ith receiver with respect to $\boldsymbol{r}_s[k]$ be given by

$$\Gamma_i[k] = \frac{\partial \rho_{is}[k]}{\partial \boldsymbol{r}_s[k]} = -\frac{(\boldsymbol{g}_i - \boldsymbol{r}_s[k])^\mathsf{T}}{\|\boldsymbol{g}_i - \boldsymbol{r}_s[k]\|}, \quad i \in \{1, 2\}; \quad k \in \mathcal{K} \tag{4.27}$$

This is approximated as constant over the short observation interval: $\Gamma_i[k] = \Gamma_i[1],\ i \in \{1, 2\}$. The modeled TDOA for $s \in \mathcal{S}_\mathrm{e}$ can be approximated by

retaining up to the linear terms in a Taylor series expansion of $\bar{y}_s[k]$ about $\tilde{\boldsymbol{r}}_s[k]$:

$$\begin{aligned}\bar{y}_s[k] &= \rho_{2s}[k] - \rho_{1s}[k] \\ &= \|(\tilde{\boldsymbol{r}}_s[k] + \boldsymbol{e}_s) - \boldsymbol{g}_2\| - \|(\tilde{\boldsymbol{r}}_s[k] + \boldsymbol{e}_s) - \boldsymbol{g}_1\| \\ &= \|\tilde{\boldsymbol{r}}_s[k] - \boldsymbol{g}_2\| - \|\tilde{\boldsymbol{r}}_s[k] - \boldsymbol{g}_1\| + (\Gamma_2 - \Gamma_1)\boldsymbol{e}_s \\ &= \tilde{y}_s[k] + \eta_s\end{aligned} \tag{4.28}$$

One thus observes that the random ephemeris error $\boldsymbol{e}_s$ manifests as a random TDOA measurement bias $\eta_s = (\Gamma_2 - \Gamma_1)\boldsymbol{e}_s$ over all $k \in \mathcal{K}$, with $\eta_s \sim \mathcal{N}(0, \sigma_s^2)$ and $\sigma_s^2 = (\Gamma_2 - \Gamma_1)Q_s(\Gamma_2 - \Gamma_1)^\mathsf{T}$. One may also lump into η_s any residual time synchronization error between the reference stations, which, like $\boldsymbol{e}_s$ can be considered constant over $\mathcal{K}$. Finally, let $\tilde{\boldsymbol{y}}_s = [\tilde{y}_s[1], \ldots, \tilde{y}_s[K]]^\mathsf{T}$ be the TLE-based modeled TDOA vector for $s \in \mathcal{S}_\mathrm{e}$.

4.6.2 TDOA-Based Association Framework

Based on the foregoing measurement models, an association framework can be developed with TLE-derived satellite positions. The objective is to assess whether a candidate satellite $s \in \mathcal{S}_\mathrm{e}$ is consistent with the measurement $\boldsymbol{y}$. For each s, the elements of $\tilde{\boldsymbol{y}}_s$ can be readily calculated as in (4.28). Let $\mathbf{1}$ be the all-ones vector of appropriate size, and let $\boldsymbol{w} = [w[1], \ldots, w[K]]^\mathsf{T} \sim \mathcal{N}(\mathbf{0}, R)$, with $R = \sigma^2 I$. Then the TDOA measurement residual with respect to the TLE-modeled satellite s is

$$\boldsymbol{\gamma}_s = \boldsymbol{y} - \tilde{\boldsymbol{y}}_s = \bar{\boldsymbol{y}}_u - \bar{\boldsymbol{y}}_s + \eta_s \mathbf{1} + \boldsymbol{w} \tag{4.29}$$

Note that if $s = u$, i.e., if s is the true satellite, then $\boldsymbol{\gamma}_s$ is zero mean. A cost can be assigned to each satellite $s \in \mathcal{S}_\mathrm{e}$ under a Bayesian framework that treats η_s as an unknown random parameter to be estimated. This is similar to how unknown transmitter parameters are treated in Section 4.5 except that η_s is modeled as a random variable with a known prior distribution: $\eta_s \sim \mathcal{N}(0, \sigma_s^2)$. The cost function follows from the negative log of the *a posteriori* probability $p(\eta_s|\boldsymbol{y})$:

$$J_s(\eta_s) = (\boldsymbol{\gamma}_s - \eta_s \mathbf{1})^\mathsf{T} R^{-1} (\boldsymbol{\gamma}_s - \eta_s \mathbf{1}) + \left(\frac{\eta_s}{\sigma_s}\right)^2 \tag{4.30}$$

For satellite s, the maximum *a posteriori* estimate of η_s is the value that minimizes the cost: $\eta_s^* = \mathrm{argmin}_{\eta_s} J_s(\eta_s)$. Under the true hypothesis H_u, the minimum cost is distributed as

$$H_u: \quad J_u(\eta_u^*) \sim \chi^2_K \tag{4.31}$$

This distribution arises because the total number of independent measurements is $K + 1$ (K from $\boldsymbol{y}$ and one from the prior constraint on η_u), but one degree of freedom is lost to estimate η_u [6].

Thus, as in Section 4.5.4, the satellite identification problem reduces to a multi-hypothesis association problem where comparisons between hypotheses are made on the basis of a cost function. In fact, the two techniques can be combined by taking the probability $P(s|\boldsymbol{z})$ from Section 4.5.4 as the prior probability for the TDOA-based identification problem. Then, since $\mathcal{S}_\mathrm{e}$ is assumed to include the true interference satellite, and approximating σ_s as

equal for all $s \in \mathcal{S}_\mathrm{e}$, the probability that H_s is true is

$$P(s|\boldsymbol{y}, \boldsymbol{z}) = \frac{\exp\bigl(-\frac{1}{2}J_s(\eta_s^*)\bigr)P(s|\boldsymbol{z})}{\sum_{a\in\mathcal{S}_\mathrm{e}} \exp\bigl(-\frac{1}{2}J_a(\eta_a^*)\bigr)P(a|\boldsymbol{z})}, \quad s \in \mathcal{S}_\mathrm{e} \tag{4.32}$$

Here, the vector $\boldsymbol{z}$ contains the ordered station-specific detection statistics for those stations $\mathcal{I}_\mathrm{d}$ that detect the interference event associated with the TDOA measurement, i.e., values from the set $\{\Lambda_i \mid i \in \mathcal{I}_\mathrm{d}\}$. As before, one can choose the highest-probability satellite, or the top N most probable, or all those for which $P(s|\boldsymbol{y}, \boldsymbol{z})$ exceeds a threshold.

4.6.3 A Wideband Raw-Signal Capture

On February 11, 2026, several wide-area transient GNSS interference events occurred across Europe. Fig. 4.16 shows the CNR of the tracked GPS L1 C/A and BeiDou B1I signals at the METG station in Finland for an interval on this day during which two events occurred. For each event, the GPS L1 C/A signals first dropped by 5 dB, then recovered, immediately following which the BeiDou B1I signals experienced two cycles of 5 dB degradation and recovery. The CNR drops of the BeiDou B1I signals persisted for ten seconds—more than twice as long as the drops observed on the GPS L1 C/A signals. The interference heat maps shown in Fig. 4.17 reveal a spatial pattern consistent with previously observed interference events.

The remainder of this section analyzes contemporaneous raw IQ samples captured by a receiver located in Amsterdam, Netherlands (R1), and one located in Trondheim, Norway (R2). R1 employed a complex sampling rate

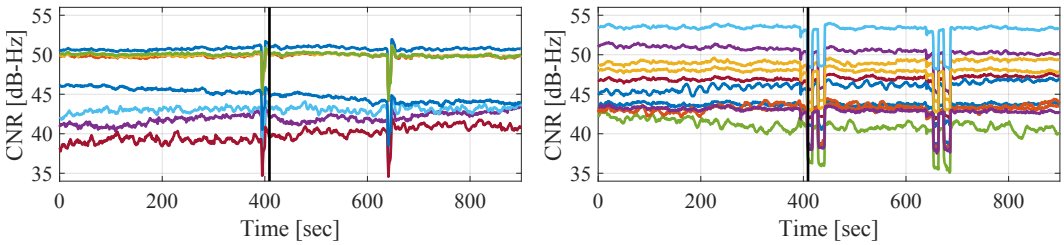


Figure 4.16: CNR time history of tracked GPS L1 C/A (left) and BeiDou B1I (right) signals at the IGS station METG in Finland. The GPS L1 C/A signals are first affected by each interference event, followed by the BeiDou B1I signals. The black line indicates the 2.3-second time-overlapped raw IQ capture interval.

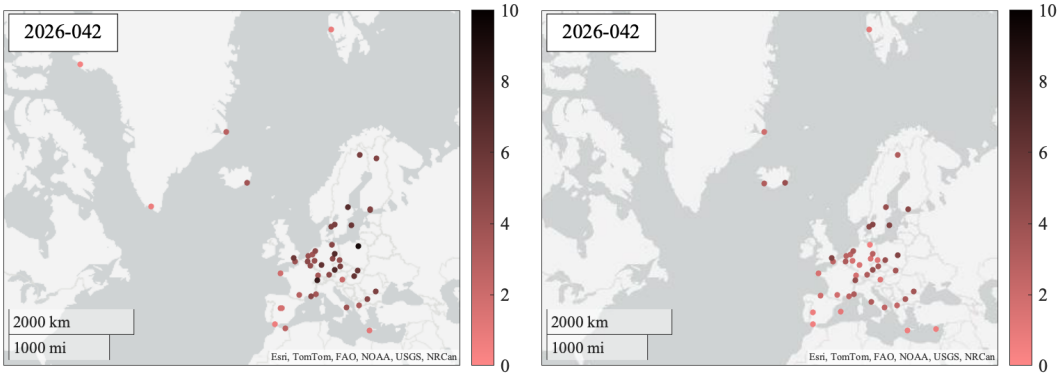


Figure 4.17: Interference heat maps for day 42 of year 2026 from tracked GPS L1 C/A (left) and BeiDou B1I (right) signals.

of 60 MHz, while R2 sampled at 75 MHz. Both captures were centered at 1585 MHz, utilized three-bit quantization, and were driven by high-stability oven-controlled crystal oscillators (OCXOs). These receivers are part of the Advanced RFI Detection, Analysis, and Alerting System, which focuses on the capture and collection of GNSS interference signals [98, 122]. The system was developed under the European Space Agency Navigation, Innovation, and Support Program (NAVISP).

R1 captured the interference signals responsible for the spatial degradation pattern shown in Fig. 4.17 and the power spectrum shown in Fig. 4.9. The interference waveform that caused CNR drops on the GPS L1 C/A signals was centered at 1577.5 MHz, which is consistent with the interference observed in Fig. 4.8. This signal exhibited cyclostationary properties, with repeating structure every 12 us. The interference waveform that caused the CNR drops on the BeiDou B1I signals was centered at 1558.5 MHz. It too exhibited cyclostationary properties, at first with repeating structure every 255.8 us, then later at 292.9 us.

On February 11, 2026 at 05:36:30 UTC, R1 and R2 recorded 2.3 seconds of time-overlapped raw IQ samples during the 1558.5-MHz-centered interference event. Several GPS and Galileo signals were contained within the captured bandwidth, which permitted processing and GPS time registration to better than 30 ns using the GRID software-defined GNSS receiver [102,109].

4.6.4 TDOA Measurement and Method Validation

The IQ samples from the two receivers were synchronized in both time and frequency and the samples from R1 were upsampled to 75-MHz to match those of R2. To obtain the maximum likelihood estimates of TDOA and frequency-difference of arrival (FDOA), the complex ambiguity function (CAF) was computed [124]. The TDOA and FDOA values that maximize the CAF magnitude correspond to the maximum likelihood estimates. The left panel of Fig. 4.18 shows the magnitude of the normalized cross correlation

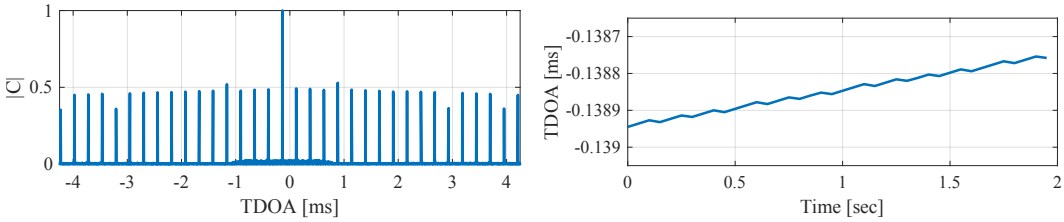


Figure 4.18: Left: Example normalized cross correlation at the optimal FDOA. Right: TDOA measurement time history.

C at the optimal FDOA estimate (155 Hz) for a 50 ms integration interval. Similar functions were repeatedly calculated every 50 ms over two seconds to obtain the 40 TDOA measurements shown as a time history in the right panel of Fig. 4.18. The TDOA measurements are magnitude-bounded by 4.25 ms, which corresponds to the delay of a hypothetical light pulse traveling between the two receivers. The first 0.3 seconds of the time-overlapped IQ samples were excluded because near-equal-magnitude cyclostationary peaks every 255.8 us caused ambiguity in the CAF. Apparently, the interference signal was purely periodic during this interval. After 0.3 seconds, significant repetition remained present, but the CAF peak became unambiguous, as shown in the left panel of Fig. 4.18.

To validate the TDOA measurement technique, TDOAs were measured for each available GPS L1 C/A signal captured in the wideband data and compared against the expected TDOA for the corresponding overhead GPS satellite. Expected TDOA values were modeled using the true satellite positions as provided by IGS final ephemerides. Besides differential time-of-flight, the model also accounted for tropospheric delay via the Saastamoinen model,

ionospheric delay via the Klobuchar model, and ECEF frame rotation during signal time of flight. The resulting measured and modeled TDOAs exhibited excellent agreement, remaining within 15 ns of each other during the 2.3-second overlapping interval of the wideband captures.

4.6.5 Satellite Association

The measurement vector $\boldsymbol{y}$ was formed from the time history shown in Fig. 4.18. Values for the TDOA measurement noise variance σ^2 and for the TLE-based position error variance Q_s were chosen as follows for all $s \in \mathcal{S}_\mathrm{e}$. Because the interference signal has a high interference-to-noise ratio and a bandwidth spanning several MHz, the TDOA estimation errors are dominated by nearest-sample quantization error, the variance of which, denoted σ_n^2, can be determined by calculating the variance of a uniform distribution with a width of one sample, or $\sigma_\mathrm{n}^2 = 0.76^2$ m^2 for 75-MHz sampling. Alternatively, a conservative choice of σ^2 can be determined by selecting the width of half a sample, $\sigma^2 = 2^2$ m^2 for 75-MHz sampling, which is an over-bound when compared to σ_n^2 and more amenable to the developed Gaussian framework. Thus, $\sigma^2 = 2^2$ m.

Position estimates derived from propagated TLEs typically exhibit kilometer-level inaccuracies, which can escalate to 10 km under volatile orbital conditions [133]. Error is typically concentrated in the along-track direction and increases with the age of ephemeris. A conservative approach was taken to over-bound the ephemeris position error. For all $s \in \mathcal{S}_\mathrm{e}$, the position covari-

ance in ECEF was set to $Q_s = 10000^2 I$. The TLEs of satellites identified as either "payload" or "unknown" within seven days of the interference event were retrieved from space-track.org. For each satellite, the TLE with the epoch nearest to the measurement time was used to calculate position. The set $\mathcal{S}_\mathrm{e}$ was composed of all satellites satisfying elevation angle masking with $\alpha_0 = 0$ for all stations that detected the event.

For each satellite $s \in \mathcal{S}_\mathrm{e}$, the cost $J_s(\eta_s^*)$ was computed. For a threshold ν chosen such that $P(J_u(\eta_u^*) > \nu) \leq 10^{-3}$, only one satellite satisfied $J_s(\eta_s^*) \leq \nu$: Cosmos 2546 (NORAD ID 45608). This is one of six Molniya-orbit satellites that compose the Russian Edinaya Kosmicheskaya Sistema (EKS), an early warning constellation. For Cosmos 2546, even assuming equally probable priors $P(s|\boldsymbol{z}) = 1/|\mathcal{S}_\mathrm{e}|$ for $s \in \mathcal{S}_\mathrm{e}$, the probability $P(s|\boldsymbol{y}, \boldsymbol{z})$ calculated according to (4.32) for Cosmos 2546 was numerically indistinguishable from unity. Cosmos 2546 also satisfied an $\alpha_0 = 35^\circ$ elevation mask with respect to every $i \in \mathcal{I}_\mathrm{d}$. This is consistent with its highly elliptical orbit, which positioned it far above the northern Atlantic Ocean during the event captured by the wideband data.

T/FDOA residuals for Cosmos 2546 are shown in Fig. 4.19. The estimated value of the TLE error was $\eta_s^* = 188$ m and the standard deviation of the TDOA residuals was 1.2 m, indicating a highly plausible match. The projected TDOA standard deviation due to TLE error was $\sigma_s = 370$ m for this satellite given the assumed Q_s. The FDOA residuals shown in the right panel of Fig. 4.19 also indicate tight consistency. FDOA measurements were

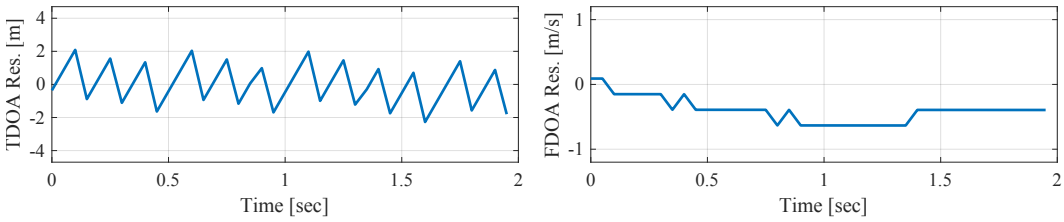


Figure 4.19: TDOA (left) and FDOA (right) residuals for Cosmos 2546 (NORAD ID 45608). The effects of nearest-sample quantization and Doppler search quantization are evident in the respective plots.

not used for association but could easily be incorporated into an augmented association framework.

4.6.6 Further Analysis and Discussion

When all space objects are considered, not only those in $\mathcal{S}_e$, one additional satellite satisfies $J_s(\eta_s^*) < \nu$—one from SpaceX's Starlink constellation. But this satellite is not a viable candidate: during the interval of the TDOA measurement it was located over the Pacific Ocean west of South America, yielding elevation angles far below zero for all stations detecting the event. One may conclude from these false associations that fusing TDOA-based association with elevation angle considerations or CNR-based priors is necessary for a confident unique association when data are limited to a single short TDOA time history.

The TLE for Cosmos 2546 closest to the TDOA measurement time is 40 hours offset, which motivates further examination of its expected accuracy. An analysis was performed based on the full historical archive of TLEs for Cosmos 2546 since its launch in 2020. Let $\mathcal{L} = \{1, \ldots, L\}$ be the index set for

this TLE archive, and let t_l represent the epoch for the lth TLE. Let $\boldsymbol{r}_k(t_l)$ denote the position of Cosmos 2546 based on the kth TLE but propagated to the epoch of the lth TLE. For all $l \in \mathcal{L}$ and all $k \in \mathcal{L} \setminus l$, one may assume that $\boldsymbol{r}_l(t_l)$ is more accurate than $\boldsymbol{r}_k(t_l)$. Let $\mathcal{L}_l \subset \mathcal{L}$ be the set of TLEs within a 60-hour window of t_l, and let $\mathcal{R}_l = \{\|r_k(t_l) - r_l(t_l)\| \mid k \in \mathcal{L}_l\}$ be the corresponding residual set. Statistical analysis of $\mathcal{R} = \bigcup_{l \in \mathcal{L}} \mathcal{R}_l$ (after outlier removal) revealed that the maximum per-dimension residual standard deviation was below 4 km and occurred in the along-track direction. This value represents an appropriate upper bound for Cosmos 2546's *a priori* position error and falls below 10 km, the determinant of Q_s assumed in the association framework. This methodology for gauging orbital accuracy mirrors established techniques for verifying the accuracy of the IGS's real-time service products [49, 57].

Strictly speaking, identification of Cosmos 2546 as a source for the transient wide-area interference holds only for the 1558.5 MHz interference during the time of wideband IQ capture. But due to the time proximity with the 1577.5 MHz interference, and to their similar time-domain and spatial patterns, interference in the two bands around the time of the IQ capture almost surely originated from the same satellite.

Two other assumptions should be acknowledged as underlying this chapter's results: (1) the TLEs for all interference sources are available on space-track.org, and (2) neither Cosmos 2546 nor any satellite near the quasi-hyperboloid surface defined by the TDOA measurement time history in Fig.

4.18 maneuvered so significantly between its nearest TLE epoch and the TDOA measurement as to alter the association findings. Both assumptions seem reasonable. In fact, it was found that the TLE error standard deviation in Q_s could be expanded from 10 km to 500 km before any additional satellites became consistent with the TDOA data.

Note that Cosmos 2546 was launched in May 2020 and so cannot be responsible for the interference events that occurred in 2019. Moreover, Cosmos 2546 was not over Europe during some interference events after May 2020. But during all events on the 75 days shown in Table 4.1 there was at least one EKS satellite above a 35° elevation angle with respect to every reference station that observed the interference. Thus, it is highly probable that the EKS constellation is collectively responsible for the wide-area transient GNSS interference events noted since 2019.

4.7 Conclusion

This chapter presented a comprehensive analysis of a novel GNSS interference phenomenon: wide-area transient interference from a space-based source causing up to 10-dB GNSS degradation across Europe since 2019 in the important L1 band. The interference's spatial, temporal, and spectral properties were detailed. A framework was developed to detect events using 1-Hz carrier-to-noise ratio observables from a network of 165 reference stations. A total of 75 days were identified since 2019 on which at least one wide-area interference event occurred. Three techniques were developed to identify candidate

source satellites: (1) a simple technique that finds all satellites satisfying a given elevation mask (e.g., 0°) for each station detecting the interference; (2) a more advanced technique based on each station's detection statistic rather than on a binary decision, and on estimation of unknown parameters such as the interference transmit power; and (3) a technique based on time-difference-of-arrival measurements made possible by capture of wideband (e.g., 60-MHz) raw samples during an interference event. By a combination of these techniques the satellite Cosmos 2546 (NORAD ID 45608) was identified with high confidence as one source of the interference. Further analysis pointed to the Russian Edinaya Kosmicheskaya Sistema, an early warning constellation to which Cosmos 2546 belongs, as collectively responsible for the wide-area transient interference causing GNSS degradation across Europe since 2019.

Chapter 5

Conclusions

This dissertation has established comprehensive frameworks for the detection and geolocation of both terrestrial and space-based GNSS interference. The developed theory has been implemented and validated through experimental and real-world data. This research enables safe and secure navigation for GNSS users.

This dissertation's first contribution 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.

This dissertation's second contribution explored two-step and direct geolocation of terrestrial emitters from LEO. The full algorithms for two-step and single-step direct geolocation were presented. It was demonstrated that the direct approach is a powerful geolocation technique for low SNR signals

during short captures with multiple emitters. This chapter also investigated emitter geolocation in crowded signal environments and explored geolocating cyclostationary signals. The relationship between the CAF domain and the position domain was explored and shown to be mappings of each other. Finally, this chapter demonstrated two-step and direct geolocation on raw intermediate frequency samples recorded from Spire Global's LEO constellation. Recent real-world GNSS interference signals captured by two time-synchronized LEO receivers over the Eastern Mediterranean were characterized and their emitters geolocated.

This dissertation's third contribution presented a comprehensive analysis of a novel GNSS interference phenomenon: wide-area transient interference from a space-based source causing up to 10-dB GNSS degradation across Europe since 2019 in the important L1 band. The interference's spatial, temporal, and spectral properties were detailed. A framework was developed to detect events using 1-Hz carrier-to-noise ratio observables from a network of 165 reference stations. A total of 75 days were identified since 2019 on which at least one wide-area interference event occurred. Three techniques were developed to identify candidate source satellites: (1) a simple technique that finds all satellites satisfying a given elevation mask (e.g., 0°) for each station detecting the interference; (2) a more advanced technique based on each station's detection statistic rather than on a binary decision, and on estimation of unknown parameters such as the interference transmit power; and (3) a technique based on time-difference-of-arrival measurements made possible by capture of

wideband (e.g., 60-MHz) raw samples during an interference event. By a combination of these techniques the satellite Cosmos 2546 (NORAD ID 45608) was identified with high confidence as one source of the interference. Further analysis pointed to the Russian Edinaya Kosmicheskaya Sistema, an early warning constellation to which Cosmos 2546 belongs, as collectively responsible for the wide-area transient interference causing GNSS degradation across Europe since 2019.

Together, these contributions strengthen GNSS security by providing the theory and experimental verification of GNSS interference detection and geolocation algorithms across several platforms. These developed capabilities can help protect the GNSS spectrum against harmful interference.

Bibliography

[1] Kibrom Ebuy Abraha, Anders Frisk, and Peter Wiklund. GNSS interference monitoring and detection based on the Swedish CORS network SWEPOS. *Journal of Geodetic Science*, 14(1):20220157, 2024.

[2] Jason Anderson, Sherman Lo, Andrew Neish, and Todd Walter. Authentication of satellite-based augmentation systems with over-the-air rekeying schemes. *NAVIGATION*, 70(3), 2023.

[3] Jason Anderson, Sherman Lo, and Todd Walter. Authentication security of combinatorial watermarking for GNSS signal authentication. *NAVIGATION*, 71(3), 2024.

[4] Jane Arraf. Israel fakes GPS locations to deter attacks, but it also throws off planes and ships. NPR, April 2024.

[5] Robert J Bamberger, Jay G Moore, Ravi P Goonasekeram, and David H Scheidt. Autonomous geolocation of RF emitters using small, unmanned platforms. *Johns Hopkins APL technical digest*, 32(3):636–646, 2013.

[6] Yaakov Bar-Shalom, X. Rong Li, and Thiagalingam Kirubarajan. *Estimation with Applications to Tracking and Navigation*. John Wiley and Sons, New York, 2001.

[7] Matt Berg. GPS 'spoofing' thickens the fog of war. Politico, October 2023.

[8] Marc 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.

[9] Jahshan Bhatti. *Sensor Deception Detection and Radio-Frequency Emitter Localization*. PhD thesis, The University of Texas at Austin, Aug. 2015.

[10] Jahshan Bhatti and Todd E Humphreys. Hostile control of ships via false GPS signals: Demonstration and detection. *NAVIGATION*, 64(1):51–66, 2017.

[11] Jahshan A. Bhatti, Todd E. Humphreys, and Brent M. Ledvina. Development and demonstration of a TDOA-based GNSS interference signal localization system. In *Proceedings of the IEEE/ION PLANS Meeting*, pages 1209–1220, April 2012.

[12] Samuel S Blackman and Robert Popoli. Kinematic state estimation: Filtering and prediction. In *Design and analysis of modern tracking systems*. Artech House, Massachusetts, 1999.

[13] Daniele Borio, Cillian O'Driscoll, and Joaquim Fortuny. GNSS jammers: Effects and countermeasures. In *2012 6th ESA Workshop on Satellite Navigation Technologies (Navitec 2012) & European Workshop on GNSS Signals and Signal Processing*, pages 1–7. IEEE, 2012.

[14] Robert G. Brown and Patrick Y.C. Hwang. *Introduction to Random Signals and Applied Kalman Filtering*. Wiley, 2012.

[15] Matt Burgess. The dangerous rise of GPS attacks. WIRED, April 2024.

[16] C4ADS. Above us only stars: Exposing GPS spoofing in Russia and Syria, April 2019. https://c4ads.org/reports.

[17] Sofia G Catalan, Brandon A Jones, and James S McCabe. Implementation of machine learning methods for crater-based navigation. In *Astrodynamics Specialist Conference*, 2021.

[18] Alessandro P. Cerruti, Pau M. Kintner, Dale E. Gary, Louis J. Lanzerotti, Eurico R. de Paula, and Hien B. Vo. Observed solar radio burst effects on GPS/Wide Area Augmentation System carrier-to-noise ratio. *Space Weather*, 4(S10006), Oct. 2006.

[19] Alessandro P Cerruti, Paul M Kintner, Dale E Gary, Anthony J Mannucci, Robert F Meyer, Patricia Doherty, and Anthea J Coster. Effect of intense December 2006 solar radio bursts on GPS receivers. *Space Weather*, 6(10), 2008.

[20] X Chen, YJ Morton, W-X Yu, and T-K Truong. GNSS spoofer localization with counterfeit clock bias observables on a mobile platform. *IEEE Sensors Journal*, 2023.

[21] Clara Chew, T Maximillian Roberts, and Steve Lowe. RFI mapped by spaceborne GNSS-R data. *NAVIGATION*, 70(4), 2023.

[22] Jacob S. Clements and Zachary L. Clements. GPU-accelerated direct geolocation of GNSS interference. *Masters Scholarly Article, the University of Maryland at College Park*, 2024.

[23] Zachary Clements, Patrick Ellis, and Todd E Humphreys. Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit. In *Proceedings of the IEEE/ION PLANS Meeting*, pages 458–469, Monterey, CA, 2023.

[24] Zachary Clements, Patrick Ellis, and Todd E Humphreys. Pinpointing GNSS interference from low Earth orbit. *Inside GNSS*, 18(5):42–55, 2023.

[25] Zachary Clements, Patrick Ellis, Mark L. Psiaki, and Todd E. Humphreys. Geolocation of terrestrial GNSS spoofing signals from low Earth orbit. In *Proceedings of the ION GNSS+ Meeting*, pages 3418–3431, Denver, CO, 2022.

[26] Zachary Clements, Iain Goodridge, Patrick Ellis, Matthew J. Murrian, and Todd E. Humphreys. Demonstration of single-satellite GNSS spoofer geolocation. In *Proceedings of the ION International Technical Meeting*, pages 361–373, Long Beach, CA, 2024.

[27] Zachary Clements, Peter A. Iannucci, Todd E. Humphreys, and Thomas Pany. Optimized bit-packing for bit-wise software-defined GNSS radio. In *Proceedings of the ION GNSS+ Meeting*, pages 3749–3771, St. Louis, MO, 2021.

[28] Zachary Clements, James E. Yoder, and Todd E. Humphreys. Carrier-phase and IMU based GNSS spoofing detection for ground vehicles. In *Proceedings of the ION International Technical Meeting*, pages 83–95, Long Beach, CA, 2022.

[29] Zachary Clements, James E. Yoder, and Todd E. Humphreys. GNSS spoofing detection: An approach for ground vehicles using carrier-phase and inertial measurement data. *GPS World*, 34(2):36–41, 2023.

[30] Zachary L. Clements, Patrick B. Ellis, Matthew J. Murrian, Mark L. Psiaki, and Todd E. Humphreys. Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit. *NAVIGATION*, 73(1), 2026.

[31] Zachary L. Clements and Todd E. Humphreys. Transient space-based GNSS interference: Observations and analysis. In *Proceedings of the ION GNSS+ Meeting*, pages 71–81, Baltimore, MD, 2025.

[32] Zachary L. Clements and Todd E. Humphreys. Orbital outlaw: GNSS interference from space. *Royal Institute of Navigation: Navigation News*, Jan/Feb:10–16, 2026.

[33] Michael Dacus, Zixi Liu, Sherman Lo, and Todd Walter. Improved RFI localization through aircraft position estimation during losses in ADS-B reception. In *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*, pages 947–957, 2022.

[34] John M Dow, Ruth E Neilan, and Chris Rizos. The international GNSS service in a changing landscape of global navigation satellite systems. *Journal of geodesy*, 83:191–198, 2009.

[35] Daniel Egea-Roca, José A López-Salcedo, Gonzalo Seco-Granados, and H Vincent Poor. Performance bounds for finite moving average tests in transient change detection. *IEEE Transactions on Signal Processing*, 66(6):1594–1606, 2018.

[36] Patrick Ellis and Farid 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)*, pages 1–5. IEEE, 2018.

[37] Patrick Ellis, Donald Van Rheeden, and Farid Dowla. Use of Doppler and Doppler rate for RF geolocation using a single LEO satellite. *IEEE Access*, 8:12907–12920, 2020.

[38] Patrick B Ellis and Farid Dowla. Single satellite emitter geolocation in the presence of oscillator and ephemeris errors. In *2020 IEEE Aerospace Conference*, pages 1–7. IEEE, 2020.

[39] ESA NAVISP-EL3-014. CGI GNSS event notification service final presentation, 2023. https://navisp.esa.int/project/details/116/show.

[40] European Space Agency. Lessons from the November 2025 solar storm,

2025. https://www.esa.int/Space_Safety/Space_weather/Lessons_from_the_November_2025_solar_storm.

[41] Bryn Fanger and Ryan P Russell. Thrust magnitude constraint handling for low-thrust trajectory optimization. *Journal of Spacecraft and Rockets*, 62(5):1506–1518, 2025.

[42] Bryn N Fanger and Ryan P Russell. Augmented lagrangian methods for handling terminal constraints in spacecraft trajectory optimization. *Astrodynamics*, 10(1):179–197, 2026.

[43] Michael Felux, Patric Fol, Benoit Figuet, Manuel Waltert, and Xavier Olive. Impacts of global navigation satellite system jamming on aviation. *NAVIGATION*, 71(3), 2024.

[44] Ignacio Fernandez-Hernandez, Jon Winkel, Cillian O'Driscoll, Simon Cancela, Rafael Terris-Gallego, Jose A López-Salcedo, Gonzalo Seco-Granados, Andrea Dalla Chiara, Carlo Sarto, Daniel Blonski, et al. Semi-assisted signal authentication for Galileo: Proof of concept and results. *IEEE Transactions on Aerospace and Electronic Systems*, 2023.

[45] Selam Gebrekidan. Electronic warfare confounds civilian pilots, far from any battlefield. The New York Times, Nov. 2023.

[46] Selam Gebrekidan. An Israeli air base is a source of GPS spoofing attacks, researchers say. The New York Times, July 2024.

[47] Selam Gebrekidan, K.K. Rebecca Lai, Pablo Robles, and Jeremy White. Why GPS is under attack. The New York Times, July 2024.

[48] Inside GNSS. The path to assured PNT: A framework for national leadership, September 2025.

[49] Jake Griffiths and Jim R Ray. On the precision and accuracy of IGS orbits. *Journal of Geodesy*, 83(3):277–287, 2009.

[50] Jason Gross and Todd 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.

[51] Jason N Gross, Cagri Kilic, and Todd E Humphreys. Maximum-likelihood power-distortion monitoring for GNSS-signal authentication. *IEEE Transactions on Aerospace and Electronic Systems*, 55(1):469–475, 2018.

[52] RIN Maritime GNSS Interference Working Group. The impacts of GNSS interference on maritime saftey report. Technical report, Royal Institute of Navigation, 2026.

[53] Blaise Kévin Guépié, Lionel Fillatre, and Igor Nikiforov. Sequential detection of transient changes. *Sequential Analysis*, 31(4):528–547, 2012.

[54] Sanjeev Gunawardena and Frank Van Graas. Multi-channel wideband GPS anomalous event monitor. In *Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2011)*, pages 1957–1968, 2011.

[55] Sanjeev Gunawardena, Zhen Zhu, Maarten Uijt de Haag, and Frank van Graas. Remote-controlled, continuously operating GPS anomalous event monitor. *NAVIGATION*, 56(2):97–113, 2009.

[56] Christoph Günther. A survey of spoofing and counter-measures. *NAVIGATION*, 61(3):159–177, 2014.

[57] Tomasz Hadas and Jaroslaw Bosy. IGS RTS precise orbits and clocks verification and quality degradation over time. *GPS solutions*, 19(1):93–105, 2015.

[58] DP Haworth, NG Smith, Riccardo Bardelli, and T Clement. Interference localization for EUTELSAT satellites–the first European transmitter location system. *International journal of satellite communications*, 15(4):155–183, 1997.

[59] Tucker Haydon, Andy Huang, and Todd E Humphreys. Analysis of the Trusted Inertial Terrain-Aided Navigation measurement function. *NAVIGATION*, 72(3), 2025.

[60] Tucker Haydon and Todd E. Humphreys. The effect of coarse Doppler resolution on the performance of a vertical synthetic aperture radar navigation system. In *Proceedings of the ION GNSS+ Meeting*, Baltimore, MD, 2025.

[61] K.C. Ho and Y.T. Chan. Geolocation of a known altitude object from TDOA and FDOA measurements. *IEEE Transactions on Aerospace and Electronic Systems*, 33(3):770–783, July 1997.

[62] Todd 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.

[63] Todd E Humphreys. Detection strategy for cryptographic GNSS anti-spoofing. *IEEE Transactions on Aerospace and Electronic Systems*, 49(2):1073–1090, 2013.

[64] Todd E. Humphreys. TEXBAT data sets 7 and 8, 2016. http://radionavlab.ae.utexas.edu/datastore/texbat/texbat_ds7_and_ds8.pdf.

[65] Todd E. Humphreys. Interference. In *Springer Handbook of Global Navigation Satellite Systems*, pages 469–503. Springer International Publishing, 2017.

[66] Todd E. Humphreys, Jahshan A. Bhatti, Daniel P. Shepard, and Kyle D. Wesson. The Texas Spoofing Test Battery: Toward a standard for evaluating GNSS signal authentication techniques. In *Proceedings of the ION GNSS Meeting*, 2012.

[67] Todd E. Humphreys, Peter A. Iannucci, Zacharias M. Komodromos, and Andrew M. Graff. Signal structure of the Starlink Ku-band downlink. *IEEE Transactions on Aerospace and Electronic Systems*, 59(5):6016–6030, Oct. 2023.

[68] Todd E Humphreys, Brent M Ledvina, Mark L Psiaki, Brady W. O'Hanlon,

and Paul M Kintner, Jr. Assessing the spoofing threat: Development of a portable GPS civilian spoofer. In *Proceedings of the ION GNSS Meeting*, Savannah, GA, 2008. Institute of Navigation.

[69] Patrick Y Hwang and Gary A McGraw. Receiver autonomous signal authentication (RASA) based on clock stability analysis. In *Proceedings of the IEEE/ION PLANS Meeting*, pages 270–281. IEEE, 2014.

[70] Peter Anthony Iannucci and Todd E. Humphreys. Fused low-Earth-orbit GNSS. *IEEE Transactions on Aerospace and Electronic Systems*, 60(4):3730–3749, Jun. 2022.

[71] Ali Jafarnia-Jahromi, Ali Broumandan, John Nielsen, and Gerard Lachapelle. GPS vulnerability to spoofing threats and review of antispoofing techniques. *International Journal of Naivgation and Observation*, pages 1–16, 2012.

[72] Ali Jafarnia-Jahromi, Saeed Daneshmand, Ali Broumandan, John Nielsen, and Gérard Lachapelle. PVT solution authentication based on monitoring the clock state for a moving GNSS receiver. In *European navigation conference (ENC)*, volume 11, 2013.

[73] John A. Volpe National Transportation Systems Center. Vulnerability assessment of the transportation infrastructure relying on the Global Positioning System, 2001.

[74] Gary Johnston, Anna Riddell, and Grant Hausler. The international GNSS service. In *Springer Handbook of Global Navigation Satellite*

Systems, pages 967–982. Springer International Publishing, 2017.

[75] Z. M. Kassas. *Analysis and Synthesis of Collaborative Opportunistic Navigation Systems*. PhD thesis, The University of Texas at Austin, May 2014.

[76] Andrew J. Kerns, Daniel P. Shepard, Jahshan A. Bhatti, and Todd E. Humphreys. Unmanned aircraft capture and control via GPS spoofing. *Journal of Field Robotics*, 31(4):617–636, 2014.

[77] Andrew J. Kerns, Kyle D. Wesson, and Todd E. Humphreys. A blueprint for civil GPS navigation message authentication. In *Proceedings of the IEEE/ION PLANS Meeting*, May 2014.

[78] Ali Khalajmehrabadi, Nikolaos Gatsis, David Akopian, and Ahmad F Taha. Real-time rejection and mitigation of time synchronization attacks on the Global Positioning System. *IEEE Transactions on Industrial Electronics*, 65(8):6425–6435, 2018.

[79] Zacharias M. Komodromos and Todd E. Humphreys. Signal parameter estimation and demodulation of the OneWeb Ku-Band downlink. *npj Wireless Technology*, 2(7):1–10, 2026.

[80] Zacharias M. Komodromos, Samuel C. Morgan, Zachary L. Clements, Wenkai Qin, W. Jeremy Morrison, and Todd E. Humphreys. Network-aided pseudorange-based LEO PNT from OneWeb. In *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.

[81] Jan Kouba. A guide to using international GNSS service (IGS) products, 2009.

[82] Argyris Kriezis, Yu-Hsuan Chen, Dennis Akos, Sherman Lo, and Todd Walter. GNSS RFI detection and impact characterization in various interference environments using low-cost receivers. In *Proceedings of the ION GNSS+ Meeting*, pages 3348–3360, 2024.

[83] Argyris Kriezis, Yu-Hsuan Chen, Dennis Akos, Sherman Lo, and Todd Walter. GNSS jamming and spoofing monitoring using low-cost COTS receivers. *NAVIGATION*, 2025. Submitted for review.

[84] Grzegorz Krzan, Karol Dawidowicz, and Pawel Wielgosz. Antenna phase center correction differences from robot and chamber calibrations: the case study LEIAR25. *GPS Solutions*, 24(2):44, 2020.

[85] Birendra Kujur, Samer Khanafseh, and Boris Pervan. Optimal INS monitor for GNSS spoofer tracking error detection. *NAVIGATION*, 71(1), 2024.

[86] Daniel M. LaChapelle, Lakshay Narula, and Todd E. Humphreys. Orbital war driving: Assessing transient GPS interference from LEO. In *Proceedings of the ION GNSS+ Meeting*, St. Louis, MO, 2021.

[87] Jinzhou Li, Le Yang, Fucheng Guo, and Wenli Jiang. Coherent summation of multiple short-time signals for direct positioning of a wideband source based on delay and Doppler. *Digital Signal Processing*, 48:58–70, 2016.

[88] Zixi Liu, Sherman Lo, Todd Walter, and Juan Blanch. Real-time detection and localization of GNSS interference source. In *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*, pages 3731–3742, 2022.

[89] E Pérez Marcos, Andriy Konovaltsev, Stefano Caizzone, Manuel Cuntz, Kazeem Yinusa, Wahid Elmarissi, and Michael Meurer. Interference and spoofing detection for GNSS maritime applications using direction of arrival and conformal antenna array. In *Proceedings of the 31st International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2018)*, pages 2907–2922, 2018.

[90] E. Pérez Marcos, S. Caizzone, A. Konovaltsev, M. Cuntz, W. Elmarissi, K. Yinusa, and M. Meurer. Interference awareness and characterization for GNSS maritime applications. In *2018 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, page 908–919. IEEE, April 2018.

[91] Willard A Marquis and Daniel L Reigh. The GPS block IIR and IIR-M broadcast L-band antenna panel: its pattern and performance. *NAVIGATION*, 62(4):329–347, 2015.

[92] Austin McKibben, Ryan McKnight, Brian C Peters, Zachary Arnett, and Sabrina Ugazio. Interference effects on a multi-GNSS receiver onboard a cubesat in LEO. In *Proceedings of the ION GNSS+ Meeting*, pages 1245–1258, Denver, CO, 2023.

[93] Tara Mina, Ashwin Kanhere, Akshay Shetty, and Grace Gao. GPS

spoofing-resilient filtering using self-contained sensors and chimera signal enhancement. *NAVIGATION*, 71(2), 2024.

[94] R.H. Mitch, R.C. Dougherty, M.L. Psiaki, S.P. Powell, B.W. O'Hanlon, J.A. Bhatti, and T.E. Humphreys. Signal characteristics of civil GPS jammers. In *Proceedings of the ION GNSS Meeting*, 2011.

[95] Ryan Mitch, Mark Psiaki, and Tunc Ertan. Chirp-style GNSS jamming signal tracking and geolocation. *NAVIGATION*, 63(1):15–37, 2016.

[96] Oliver Montenbruck, Peter Steigenberger, Lars Prange, Zhiguo Deng, Qile Zhao, Felix Perosanz, Ignacio Romero, Carey Noll, Andrea Stürze, Georg Weber, et al. The multi-GNSS experiment (MGEX) of the international GNSS service (IGS)–achievements, prospects and challenges. *Advances in space research*, 59(7):1671–1697, 2017.

[97] Samuel C. Morgan, Zacharias M. Komodromos, Wenkai Qin, Zachary L. Clements, Andrew M. Graff, W. Jeremy Morrison, and Todd E. Humphreys. A mock implementation of fused LEO GNSS. In *Proceedings of the IEEE/ION PLANS Meeting*, Salt Lake City, UT, 2025.

[98] Aiden Morrison, Nadezda Sokolova, Nicolai Gerrard, Anders Rødningsby, Christian Rost, and Laura Ruotsalainen. Radio-frequency interference considerations for utility of the Galileo E6 signal based on long-term monitoring by ARFIDAAS. *NAVIGATION*, 70(1), 2023.

[99] George V Moustakides. Multiple optimality properties of the shewhart test. *Sequential Analysis*, 33(3):318–344, 2014.

[100] Matthew J. Murrian, Lakshay Narula, Peter A. Iannucci, Scott Budzien, Brady W. O'Hanlon, Mark L. Psiaki, and Todd E. Humphreys. First results from three years of GNSS interference monitoring from low Earth orbit. *NAVIGATION*, 68(4):673–685, 2021.

[101] D. Musicki, R. Kaune, and W. Koch. Mobile emitter geolocation and tracking using TDOA and FDOA measurements. *IEEE Transactions on Signal Processing*, 58(3):1863–1874, March 2010.

[102] Hailey A. Nichols, Matthew J. Murrian, and Todd E. Humphreys. Software-defined GNSS is ready for launch. In *Proceedings of the ION GNSS+ Meeting*, Denver, CO, 2022.

[103] Carey E Noll. The crustal dynamics data information system: A resource to support scientific analysis using space geodesy. *Advances in Space Research*, 45(12):1421–1440, 2010.

[104] Dennis Odijk. Positioning model. In *Springer Handbook of Global Navigation Satellite Systems*, pages 605–638. Springer International Publishing, 2017.

[105] B. O'Hanlon, J. Bhatti, Todd 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, 2010. Institute of Navigation.

[106] B.W. O'Hanlon, M.L. Psiaki, J.A. Bhatti, and T.E. Humphreys. Real-time spoofing detection using correlation between two civil GPS receiver.

In *Proceedings of the ION GNSS Meeting*, Nashville, Tennessee, 2012. Institute of Navigation.

[107] Nickens Okello. Emitter geolocation with multiple UAVs. In *2006 9th International Conference on Information Fusion*, pages 1–8. IEEE, 2006.

[108] Okuary Osechas, Friederike Fohlmeister, Thomas Dautermann, and Michael Felux. Impact of GNSS-band radio interference on operational avionics. *NAVIGATION*, 69(2), 2022.

[109] Thomas Pany, Dennis Akos, Javier Arribas, M. Zahidul H. Bhuiyan, Pau Closas, Fabio Dovis, Ignacio Fernandez-Hernandez, Carles Fernández-Prades, Sanjeev Gunawardena, Todd Humphreys, Zaher M. Kassas, José A. López Salcedo, Mario Nicola, Mark L. Psiaki, Alexander Rügamer, Young-Jin Song, and Jong-Hoon Won. GNSS software defined radio: History, current developments, and standardization efforts. *NAVIGATION*, 71(1), 2024.

[110] Akshata Patil, R Eric Phelts, Yu-Hsuan Chen, Sherman Lo, and Todd Walter. Detecting space based interference on GNSS signals. In *Proceedings of the 36th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2023)*, pages 1232–1244, 2023.

[111] Akshata Patil, R Eric Phelts, Todd Walter, and Steffen Thoelert. Detecting and localizing space based interference on GNSS signals using

machine learning. In *Proceedings of the 2024 International Technical Meeting of The Institute of Navigation*, pages 532–545, 2024.

[112] H. Vincent Poor. *An Introduction to Signal Detection and Estimation, 2nd Edition.* Springer, 1994.

[113] Mark L Psiaki. Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids. *NAVIGATION*, 68(3):621–641, 2021.

[114] Mark L Psiaki and Todd E Humphreys. GNSS spoofing and detection. *Proceedings of the IEEE*, 104(6):1258–1270, 2016.

[115] Mark L. Psiaki and Todd E. Humphreys. *Position, Navigation, and Timing Technologies in the 21st Century: Integrated Satellite Navigation, Sensor Systems, and Civil Applications*, volume 1, chapter Civilian GNSS Spoofing, Detection, and Recovery, pages 655–680. Wiley-IEEE, 2020.

[116] Mark L. Psiaki, Brady W. O'Hanlon, Steven P. Powell, Jahshan A. Bhatti, Kyle D. Wesson, Todd E. Humphreys, and Andrew Schofield. GNSS spoofing detection using two-antenna differential carrier phase. In *Proceedings of the ION GNSS+ Meeting*, Tampa, FL, 2014. Institute of Navigation.

[117] Wenkai Qin, Andrew M. Graff, Zachary L. Clements, Zacharias M. Komodromos, and Todd E. Humphreys. Timing properties of the Starlink

Ku-band downlink. *IEEE Transactions on Aerospace and Electronic Systems*, pages 1–19, 2025.

[118] Katarina Radoš, Marta Brkić, and Dinko Begušić. Recent advances on jamming and spoofing detection in GNSS. *Sensors*, 24(13):4210, 2024.

[119] Alit Mendelson Reuven and Anthony J Weiss. Direct position determination of cyclostationary signals. *Signal Processing*, 89(12):2448–2464, 2009.

[120] Logan Scott. Anti-spoofing and authenticated signal architectures for civil navigation systems. In *Proceedings of the ION GNSS Meeting*, pages 1542–1552, 2003.

[121] A.Y. Sidi and A.J. Weiss. Delay and Doppler induced direct tracking by particle filter. *IEEE Transactions on Aerospace and Electronic Systems*, 50(1):559–572, January 2014.

[122] Nadezda Sokolova, Aiden Morrison, and Anja Diez. Characterization of the GNSS RFI threat to DFMC GBAS signal bands. *Sensors*, 22(22), 2022.

[123] John Stader and Sanjeev Gunawardena. Leveraging worldwide, publicly-available data to create an automated satnav interference detection system. In *Proceedings of the 2021 International Technical Meeting of The Institute of Navigation*, pages 69–83, 2021.

[124] S. Stein. Algorithms for ambiguity function processing. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 29(3):588–599, June

1981.

[125] S. Stein. Differential delay/Doppler ML estimation with unknown signals. *IEEE Transactions on Signal Processing*, 41(8):2717–2719, Aug. 1993.

[126] Andrew Tangel and Drew FitzGerald. Electronic warfare spooks airlines, pilots and air-safety officials. Wall Street Journal, Sept. 2024.

[127] 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*, 54(1):131–143, Feb. 2018.

[128] Robert Tenny and Todd E. Humphreys. Robust navigation for urban air mobility via tight coupling of GNSS with terrestrial radionavigation and inertial sensing. In *Proceedings of the ION GNSS+ Meeting*, Denver, CO, 2022.

[129] Robert Tenny and Todd E. Humphreys. Robust navigation for urban air mobility. *Inside GNSS*, 18(2):47–47, 2023.

[130] Robert M Tenny, Lisong C Sun, Alperen Duru, and Todd E Humphreys. Robust absolute headset tracking for extended reality. In *Proceedings of the IEEE/ION PLANS Meeting*, Monterey, CA, 2023.

[131] Tom Tirer and Anthony J Weiss. High resolution localization of narrowband radio emitters based on Doppler frequency shifts. *Signal Processing*, 141:288–298, 2017.

[132] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory*. Wiley, 2001.

[133] Jerome R Vetter. Fifty years of orbit determination. *Johns Hopkins APL technical digest*, 27(3):239, 2007.

[134] A.J. Weiss. Direct geolocation of wideband emitters based on delay and Doppler. *IEEE Transactions on Signal Processing*, 59(6):2513–2521, June 2011.

[135] 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*, 54(2):739–754, April 2018.

[136] Kyle Wesson. *Secure Navigation And Timing Without Local Storage Of Secret Keys*. PhD thesis, The University of Texas at Austin, May 2014.

[137] Kyle D. Wesson, Mark P. Rothlisberger, and Todd E Humphreys. Practical cryptographic civil GPS signal authentication. *NAVIGATION*, 59(3):177–193, 2012.

[138] Hanna Witzgall. Two-sensor tracking of maneuvering transmitters. In *2018 IEEE Aerospace Conference*, pages 1–7. IEEE, 2018.

[139] GPS Spoofing Workgroup. GPS spoofing: Final report of the GPS spoofing workgroup. Technical report, OPSGROUP, 2024.

[140] Johnathan York, Andrew Joplin, Michael Bratton, and David Munton. A detailed analysis of GPS live-sky signals without a dish. *NAVIGATION*, 61(4):311–322, 2014.