The Dissertation Committee for Jahshan Ahmed Bhatti
certifies that this is the approved version of the following dissertation:

Sensor Deception Detection and Radio-Frequency Emitter Localization

Committee:

Todd Humphreys, Supervisor

Srinivas Bettadpur

Behçet Açıkmeşe

Maruthi Akella

Aaron Kerkhoff

Sensor Deception Detection and Radio-Frequency Emitter Localization

by

Jahshan Ahmed Bhatti, B.S.As.E.; M.S.E.

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

December 2015

To not knowing what comes next.

Acknowledgments

First, I would like to acknowledge the various funding sources I have received throughout my graduate school program:

- National Science Foundation CAREER Award under Grant No. 1454474.
- The Data-supported Transportation Operations and Planning Center (D-STOP), a Tier 1 USDOT University Transportation Center.
- U.S. Navy STTR "EMLOC System for Emitter Detection and Localization" through Coherent Navigation, contract CN-STTR-12-001.
- Coherent Navigation under sponsored research agreement UTA10-000514.
- Lockheed Martin through the sponsored research agreement titled "FOTON Sensor Development and Risk Reduction Effort."
- U.S. Air Force STTR "Connected Autonomous Space Environment Sensors" through ASTRA, LLC, under sponsored research agreement UTA09-000852.

In particular, I thank Dr. Brent Ledvina for providing me the opportunity to intern at Coherent Navigation, an experience which helped shape my present career goals.

Second, I thank all current and former members of my committee, Drs. Todd Humphreys, Srinivas Bettadpur, Behçet Açıkmeşe, Maruthi Akella, Aaron Kerkhoff, Bob Schutz, and Glenn Lightsey, for their time, patience, and guidance. In addition, I thank Dr. Andy Norris for his comments on maritime navigation and security in practice.

Third, I thank all those who helped support the various experiments conducted for this dissertation: my colleagues at Cornell University, Brady O'Hanlon, Dr. Ryan Mitch, Steven Powell, and Dr. Mark Psiaki, for collaborating, and taking lead, on various projects; the Department of Homeland Security and the 746th Test Squadron, for organizing and executing the civilian GPS jamming test exercise at White Sands Missile Range; Master Andrew Schofield and his crew, for inviting me aboard the White Rose of Drachs to conduct GPS spoofing experiments in the Mediterranean, a truly unforgettable experience; Ken Pesyna, for providing an easy-to-use carrier-phase differential GPS implementation; Mark Maughmer II, for expertly piloting a very overweight quadcopter; and all members of the Radionavigation Laboratory, for their camaraderie over the last six years.

Finally, I thank the various people who have had a positive influence on me: Dr. Henri Kjellberg, for all the endless and entertaining conversation, especially when I would get frustrated with research, and for introducing me to cycling—our group rides on Shoal Creek were truly blissful; Drs. Kyle Wesson and Zak Kassas, for being role models and providing actionable advice; all the cheerful and friendly people on the fourth floor of Woolrich Labs and beyond, for providing much needed social interaction; Kathleen Garcia and Laura de la

Garza, for encouraging me to be a better person; Michael Szmuk, Chirag Patel, and Daniel Dueri, for enduring all the ridiculous aspects of undergraduate and graduate school together; Ludwig Barragán, Lauren Waelder, and Judy Sifonté, for aiding my journey into adulthood ever since freshman year of college; my brother, Aqyan, for keeping the apartment clean; and my mother, Khaleda, for always raising the bar.

Sensor Deception Detection and Radio-Frequency Emitter Localization

Publication No. _____

Jahshan Ahmed Bhatti, Ph.D.
The University of Texas at Austin, 2015

Supervisor: Todd Humphreys

The Global Positioning System (GPS) is an invisible utility that has had enormous impact in areas such as navigation, telecommunications, and power grids. However, malicious so-called "field" attacks such as jamming and spoofing threaten to disrupt and damage an infrastructure that has become so dependent on an always available and trustworthy GPS. This dissertation provides solutions that, if deployed as part of a layered defense, can significantly mitigate the effects of these emerging threats.

The first type of attack considered in this dissertation is GPS spoofing. An attacker's ability to covertly control a maritime surface vessel by broadcasting counterfeit civil GPS signals is analyzed and demonstrated. It is shown that, despite access to a variety of high-quality navigation and surveillance sensors, modern maritime navigation depends crucially on satellite-based navigation. A simple innovations-based detection framework for GPS deception

is developed, and given real-world environmental and attack parameters, the probability of hazardously misleading information or integrity risk is minimized within the framework. A covert attack is designed to have a high integrity risk and is possible because attacker-induced deviations in the vessel's dynamics can be disguised as the effects of slowly-changing ocean currents and wind. A field experiment confirms the analysis by demonstrating covert control of a 65-m yacht in the Mediterranean Sea.

The second type of attack considered in this dissertation is GPS jamming. A system for passively locating radio-frequency emitters is developed and demonstrated. The system was originally motivated by the proliferation of GPS jammers, but has broad applicability to any emitter of unknown waveform. A model for the cross-correlation of the emitter signal captured by spatially distributed receivers with an independent local oscillator and an efficient digital cross-correlation implementation is presented. Algorithms based on grid search and the particle filter are developed to estimate the emitter state directly from the cross-correlation, avoiding the inefficiency of an intermediate time and frequency difference of arrival estimate. The system is proven in several field experiments with the emitter on stationary or vehicular platforms and with one experiment using a receiver on an airborne platform.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

The next few decades will see pervasive autonomous control systems become critical to the world economy—from autonomous cars and aircraft to smart homes, smart cities, and vast energy, communication, and financial networks controlled at multiple scales. Protecting these systems from malicious attacks is a matter of urgent societal interest. The study of secure control has made important advances over the past few years [2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16], but these constitute not solutions so much as problem framing and an emerging consensus that traditional fault detection and mitigation fails when confronted with a deliberate attacker: outlaws are different from outliers; fraud is different from faults [3, 2, 10, 15]. Moreover, the majority of this early literature focuses on standard cyber attacks—those that entail infiltration of communications networks or computer systems within which sensor measurements y and control commands u are conveyed or calculated.

This dissertation focuses on an emergent category of cyber-physical attack that has seen little scrutiny in the secure control literature. Like cyber attacks, these attacks are hard to detect and can be executed from a distance, but unlike cyber attacks, they are effective even against control systems whose software, data, and communications networks are secure, and so can be con-

sidered a more menacing long-term threat. These are field attacks: attacks on the physical fields—electromagnetic, magnetic, acoustic, etc.—measured by system sensors. As specialized sensor attacks, field attacks seek to compromise a system's perception of reality non-invasively—from without, not from within. This work emphasizes field attacks against navigation and timing sensors, as these are of special importance to the rise of autonomous vehicles and the smart grid.

1.1 Sensor Deception Problem

A particularly effective and mature field attack that exploits the insecurity of civil Global Positioning System (GPS) signals is known as GPS "spoofing" [17]. The University of Texas Radionavigation Laboratory has developed an in-house GPS spoofing testbed that has been used to investigate the effects of this attack on GPS receivers embedded in a diverse set of semi-autonomous control systems. For example, GPS spoofing has been demonstrated against: (i) an autonomous helicopter, which appeared to be as if caught in a tractor beam [18, 19], (ii) a phasor measurement unit used in smart grids for microsecond-accurate timing [20], and (iii) an $80M superyacht, which was driven several kilometers off course without triggering alarms [21]. Emblematic of the current literature's limitations, no published secure estimation technique would be capable of thwarting the unmanned aerial vehicle (UAV) attack described in [19] or the yacht attack described in this dissertation, given the vehicles' respective sensor suites.

1.2 GPS Jamming Problem

Despite its marvelous success over the last three decades, the Global Positioning System has an Achilles' heel: its weak signals are an easy target for jamming. GPS jamming is a blunt denial-of-service type of field attack when compared to its more sophisticated, expensive, and targeted cousin, GPS spoofing. The National Space-Based Positioning, Navigation, and Timing Advisory Board in a recent white paper has concluded that the "United States is now critically dependent on GPS" [22]. The paper notes an alarming increase in the incidence rate of deliberate and unintentional GPS interference, which in some cases renders GPS inoperable for critical infrastructure operations. The white paper also notes the increasing availability of small and cheap GPS jammers known as personal privacy devices (PPDs). Although the advertised jamming coverage radius for these devices is small, typically 10 to 20 meters, their actual range may extend to tens of kilometers [23].

In one recent case of interest, a test version of the GPS ground-based augmentation system (GBAS) at Newark International Airport suffered from periodic interference due to a PPD aboard a truck traveling on a nearby highway [24, 25]. The authorities took four months to track down the jammer. Continued monitoring in the Newark airport area after this incident indicates that during rush hours, there occur 4 to 5 interference events per hour, presumably due to PPDs [26]. GPS-synchronized cellular communications networks also report an increasing rate of periodic GPS outages, most likely due to passing PPDs. Although these networks are designed to fall back to a hold-over mode that is capable of maintaining adequate synchronization for several days,

such interference is nonetheless an annoyance for network operators.

Despite a recent effort by the Federal Communications Commission to discourage sale, purchase, and use of PPDs [27], there is reason to believe that they will only become more widespread in the future. The miniaturization and proliferation of GPS trackers will likely lead to an increased use of PPDs, despite their being illegal, as people seek to protect their privacy from invasive tracking [28]. To aid in enforcing laws against PPDs and jamming devices, there is a need for a persistent system capable of detecting and locating sources of jamming.

The work by Scott (J911) [29], Brown (JLOC) [30], and Chronos Technology (GAARDIAN) [31] focus on building cheap, low-network-throughput jamming-to-noise ratio sensors based on monitoring GPS carrier-to-noise ratio and automatic gain control (AGC) values, making them suited only for triggering and coarse localization. The work by Akos [32, 33] considers a network of sensor nodes using a low-cost Global Navigation Satellite System (GNSS) front end with AGC monitoring capability. Single-emitter interference localization is implemented using AGC values coupled with power-law path loss models for strong sources and cross-correlation-based time difference of arrival (TDOA) estimation coupled with hyperbolic positioning for weak sources.

However, this dissertation will focus on direct geolocation techniques, first explored by Weiss and Sidi [34, 35], where "direct" refers to estimating the emitter state directly from the cross-correlation of the received signals, without making an intermediate time and frequency difference of arrival (T/FDOA)

estimate. Note that the methods developed in this dissertation can not only find GPS jammers and spoofers, but any kind of radio-frequency (RF) emitter transmitting an unknown waveform.

1.3 Dissertation Contributions

This dissertation makes two primary contributions:

(i) A GNSS deception detection technique implemented at the state estimation level that minimizes the integrity risk (the probability of undetected hazardous conditions) for a pre-determined false alarm probability. The technique is a novel adaptation and optimization of the standard innovations test for model correctness. As such, it is intuitive and simple to implement, as shown in a demonstration of a GPS spoofing attack on a yacht. The technique is developed and implemented in this dissertation for maritime surface vessels, but is generalizable to any estimation and control system that depends on GNSS such as for timing and aviation.

(ii) A particle filter for direct geolocation of radio-frequency emitters with unknown deterministic waveforms. The technique, which was developed, implemented, and tested concurrent with and independently of [35], exploits the estimation efficiency of a direct geolocation approach and the reduction of the emitter state search space when compared to naive grid search. Although this contribution is similar to [35], this dissertation extends beyond [35] in three important ways. First, it considers practical implementation issues of a deployed system such as time synchronization

of receivers and long coherent integration for non-constant T/FDOA. Second, it goes beyond the simulations in [35] to conduct and report on three field experiments. Third, it provides a comparative study of the performance and complexity of various estimation algorithms, including but not limited to those presented in [34, 35].

1.4 Published Works

The author contributed to the following publications.

Journal Publications

1. K. B. Deshpande, G. S. Bust, C. R. Clauer, H. Kim, J. E. Macon, T. E. Humphreys, J. A. Bhatti, S. B. Musko, G. Crowley, and A. T. Weatherwax, "Initial GPS scintillation results from CASES receiver at South Pole, Antarctica," *Radio Science*, vol. 47, no. 5, 2012

2. C. R. Clauer, H. Kim, K. Deshpande, Z. Xu, D. Weimer, S. Musko, G. Crowley, C. Fish, R. Nealy, T. E. Humphreys, J. A. Bhatti, and A. J. Ridley, "Autonomous adaptive low-power instrument platform (AALPIP) for remote high latitude geospace data collection," *Geoscientific Instrumentation, Methods and Data Systems*, vol. 3, pp. 211–227, 2014

   The papers analyze scintillation events captured by instruments deployed at South Pole, including the CASES software-defined GPS receiver. The author helped develop the GPS receiver and scintillation triggering algo-

rithms, aided in integration of the receiver with AAL-PIP and mitigation of interference with Iridum modem, and provided data analysis support.

3. M. Psiaki, B. O'Hanlon, J. Bhatti, D. Shepard, and T. Humphreys, "GPS spoofing detection via dual-receiver correlation of military signals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2250–2267, 2013

4. B. W. O'Hanlon, M. L. Psiaki, T. E. Humphreys, J. A. Bhatti, and D. P. Shepard, "Real-time GPS spoofing detection via correlation of encrypted signals," *Navigation, Journal of the Institute of Navigation*, vol. 60, no. 4, pp. 267–278, 2013

   The papers develop GPS spoofing detection algorithm via correlation of encrypted signals and demonstrate a real-time implementation of the algorithm. The author helped develop the GPS spoofer and conduct live tests in the lab to test the algorithm.

5. A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Unmanned aircraft capture and control via GPS spoofing," *Journal of Field Robotics*, vol. 31, no. 4, pp. 617–636, 2014

   The paper analyzes the effects of GPS spoofing on an unmanned aerial vehicle (UAV) and demonstrate capture and partial control in a live test. The author helped develop the GPS spoofer and conduct live tests. In addition, the author developed a model for closed-loop control of a UAV via GPS spoofing to predict the performance of innovations-based detection.

6. E. G. Lightsey, T. E. Humphreys, J. A. Bhatti, A. J. Joplin, B. W. O'Hanlon, and S. P. Powell, "Demonstration of a space capable miniature dual frequency GNSS receiver," *Navigation, Journal of the Institute of Navigation*, vol. 61, no. 1, pp. 53–64, 2014

   The paper demonstrates capabilities of the FOTON software-defined GPS receiver for space-based ionospheric sounding. The author helped develop the GPS receiver, provided a description of the software architecture, and aided in analysis of data from a sounding rocket demonstration flight.

7. J. Bhatti and T. Humphreys, "Hostile control of surface vessels via counterfeit GPS signals: Demonstration and detection," *Navigation, Journal of the Institute of Navigation*, 2015, (In preparation.)

   The paper develops a GPS spoofing detection framework for maritime surface vessels. The paper provides theoretical performance of the detection framework for various system parameters. Finally, the paper gives results for an over-the-air GPS spoofing test exercise against a 65-meter yacht.

8. J. Bhatti, B. Ledvina, and T. Humphreys, "Analysis and experimental results of direct geolocation techniques," *Navigation, Journal of the Institute of Navigation*, 2015, (In preparation.)

   The paper develops a direct-geolocation radio-frequency emitter localization system. The paper considers the performance of estimation algorithms based on grid search, Kalman filter, and particle filter against

two experimental scenarios, one with an airborne sensor and the other with a moving emitter.

Selected Conference Publications

1. T. E. Humphreys, J. Bhatti, T. Pany, B. Ledvina, and B. O'Hanlon, "Exploiting multicore technology in software-defined GNSS receivers," in *Proceedings of the ION GNSS Meeting*. Savannah, GA: Institute of Navigation, 2009, pp. 326–338

   The paper examines the performance gain achieved by various parallelization strategies in a software-defined GNSS receiver running on a shared-memory four-core desktop computer. The author developed and implemented the parallelization strategies via OpenMP and conducted tests to determine load balancing and overall performance gain.

2. J. A. Bhatti, T. E. Humphreys, and B. M. Ledvina, "Development and demonstration of a TDOA-based GNSS interference signal localization system," in *Proceedings of the IEEE/ION PLANS Meeting*, April 2012, pp. 1209–1220

   The paper develops a TDOA-based multiple-emitter localization system using GPS-synchronized sensor nodes and subspace-based estimation methods. The paper gives results for a test exercise in which the system with three sensors locates at least one emitter within 20 meters among a field of two emitters with significant multipath.

Magazine Articles

1. T. E. Humphreys, J. Bhatti, and B. M. Ledvina, "The GPS Assimilator: Upgrading receivers via benign spoofing," *Inside GNSS*, vol. 5, no. 4, pp. 50–58, June 2010

2. R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. E. Humphreys, "Know your enemy: Signal characteristics of civil GPS jammers," *GPS World*, Jan. 2012

3. D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Drone hack: Spoofing attack demonstration on a civilian unmanned aerial vehicle," *GPS World*, Aug. 2012

1.5 Dissertation Organization

Chapter 2 develops a GPS spoofing detection framework for maritime surface vessels. The components of modern integrated bridge system are described and their dependency on GPS is investigated. The effect of a GPS spoofing attack on an unprotected vessel is illustrated by analogy to the description of the Royal Majesty grounding in 1995 [47]. A model for the closed-loop ship and spoofing dynamics is presented. The performance of the detection framework, characterized in terms of the integrity risk, or equivalently the probability of hazardously misleading information, is optimized against a set of spoofing attack profiles. Lastly, the detection framework is applied to simulated and experimental data, and the performance is compared to the theoretical integrity risk.

Chapter 3 begins with the necessary background on the cross-correlation model that allows direct geolocation. The model takes into account the time-delay and Doppler effect of the geometric propagation of the transmitted emitter signal and the receiver clock offset. A zero-order hold approximation of the cross-correlation model is derived that allows an efficient digital implementation. A tightly-coupled RF receiver architecture that can simultaneously capture the emitter signal and a reference signal is presented. The reference signal, typically GPS, in a tightly-coupled receiver allows coherent cross-correlation between spatially distributed receivers with independent local oscillators. Algorithms for direct emitter localization based on grid search, the Kalman filter, and the particle filter are developed. The algorithms are tested against three different types of emitter dynamics model with field experiments conducted in amateur radio and GPS bands.

Chapter 4 concludes this dissertation with a summary of contributions and suggestions for future work.

Chapter 2

Hostile Control of Surface Vessels via Counterfeit GPS Signals: Demonstration and Detection

Surface vessels, from fishing boats to container ships to deep-water oil rigs, depend crucially on Global Positioning System (GPS) signals for navigation, station keeping, and surveillance [48, 49, 50, 51]. GPS, its ground and satellite-based augmentation systems, and other Global Navigation Satellite Systems (GNSS) are used as the primary position-fixing system, especially in open waters. They are an important maritime navigation aid even for vessels actively piloted by human operators, except in familiar littoral waters such as port entry and within natural or man-made channels where conventional "optical" navigation is used. Moreover, as surface craft become more autonomous, the trend is toward increased reliance on GNSS: current autopilot systems, dynamic-positioning systems, and fully unmanned surface vehicles are designed under the assumption that GNSS signals are usually available and trustworthy [49, 52, 53, 50]. Even autonomous underwater vehicles typically depend indirectly, or periodically, on GNSS [54].

Given the fragility of GNSS signals under conditions of signal blockage or jamming, and given that the signals do not penetrate underwater, there

is interest in developing GNSS-independent maritime navigation and control systems [49, 55]. Terrain-relative navigation has been successfully employed in autonomous submersibles [55], and could serve as a backup to GNSS for surface vessels. This technique has historically required high-resolution (e.g., m-level) underwater terrain maps, which are available for only a tiny fraction of the seafloor, but recent results indicate that coarser (e.g., 20-m-resolution) ship-based bathymetry maps may be adequate for 10-meter-level positioning, provided sufficient terrain variability [56]. Nonetheless, for the present, terrain-relative navigation does not even appear to be an active research topic for civil surface maritime transportation. What is more, the only widespread radionavigation backup to GNSS, Loran-C, was abandoned by the U.S. Coast Guard in 2010 [57], and there are no official U.S. plans for a successor, despite continued lobbying for deployment of its upgrade, eLoran, which is available in other parts of the world [58]. Consequently, one can expect most maritime navigation systems to rely primarily on GNSS for position-fixing for years to come.

By standard practice marine craft are equipped with redundant GNSS units so that one serves as backup if the other experiences a fault. And for extremely critical applications, an entirely GNSS-free positioning system may be available, such as the acoustic positioning system required as a backup to GNSS on dynamically-positioned deepwater drilling vessels [50]. But these fail-safe systems are designed to handle obvious faults or GNSS outages caused by signal blockage or ionospheric effects. They are likely to fail when confronted with a sophisticated and deliberate attacker: outlaws are different from out-

liers; fraud is different from faults.

A GNSS deception attack, in which counterfeit GNSS signals are generated for the purpose of manipulating a target receiver's reported position, velocity, or time, is a potentially dangerous tool in the hands of a deliberate attacker. While there have been no confirmed reports of such attacks performed with malice, convincing demonstrations have been conducted both in the laboratory and in the field with low-cost equipment against a wide variety of GPS receivers [17, 59, 19]. The key to the success of these so-called GPS spoofing attacks is that, whereas the military GPS waveforms are by design unpredictable and therefore resistant to spoofing, civil GPS waveforms—and those of other civil GNSS—are unencrypted, unauthenticated, and openly specified in publicly-available documents [60, 61]. Also, although not entirely constrained by the GNSS signal specifications, the navigation data messages modulating these civil waveforms are highly predictable. The combination of known signal structure and navigation data predictability makes civil GNSS signals an easy target for spoofing attacks.

The departure point for development of a spoofing detection framework is the impressive corpus of fault detection and isolation (FDI) literature, the result of more than four decades of effort. Sensor deception can be thought of as a special type of sensor fault in which a strategic attacker has some level of control over the fault behavior and applies this control with malicious intent. Several classes of methods for sensor FDI in stochastic linear dynamic systems are surveyed in [62, 63, 64, 65]. Although many sophisticated approaches have been developed in this mature field, most fault-detection methods focus on

minimizing time-to-detect without regard to integrity risk, as noted by Joerger [66]. Integrity risk is the appropriate figure of merit for dynamic systems with clearly specified alert limits such as aviation and maritime navigation and time transfer. For these systems, state estimation errors that remain within the alert limits cause no performance degradation or heightened safety risk, but undetected errors exceeding the alert limit can have severe consequences.

The first attempt to address sensor deception by minimizing integrity risk appears to be [67], where a model-based spoofing detection method was developed for an aircraft's GPS-aided inertial navigation system. However, the analysis considered a batch detection test whose batch interval is aligned with the attack interval, a coincidence that cannot be expected in practice. The current work adopts a sequential detection approach, which is more appropriate for attacks of unknown start time and duration. But as opposed to sequential detection techniques designed to minimize time-to-detect for fixed probabilities of false alarm and detection, such as the sequential probability ratio test [68], the current work adopts a fixed time-to-detect approach and follows [66] and [67] in seeking to minimize integrity risk. More precisely, this work minimizes mean integrity risk, or integrity risk averaged over all possible attack start times.

The heart of a detection technique is the so-called detection statistic, a function of the sensor measurements that gets compared to a threshold [69]. This work adopts an innovations-based detection statistic whose performance is insensitive to the particular time history of false differential position and velocity induced by the attacker.

A key feature of the current work's detection framework is that it optimizes the measurement sampling interval; the standard innovations-based detection approach makes no attempt at such optimization [70, 62]. The optimization seeks to minimize worst-case integrity risk over a set of reasonable attack profiles. Measurement sampling interval optimization was previously considered in [71], but that work minimized time-to-detect, whereas the current work's criterion is integrity risk.

This chapter makes three contributions. First, it details the pathways and effects of GNSS deception on maritime navigation and surveillance. Whereas maritime transportation's vulnerability to GNSS jamming has been previously established [49], this work offers the first detailed analysis of the effects of GNSS deception on a surface vessel. Second, it develops an innovations-based spoofing detection framework and optimizes the worst-case mean integrity risk within this framework given a set of reasonable attack profiles. Third, it presents the results of an unprecedented field experiment demonstrating hostile control of a 65-m yacht in the Mediterranean Sea.

2.1 GNSS Dependencies of a Modern Integrated Bridge System

This section details the pathways and effects of GNSS deception on maritime navigation and surveillance. Besides providing a deeper understanding of the vulnerability of maritime vessels to GNSS spoofing, this overview will identify a subset of ship sensors that can conveniently and effectively be applied to the problem of spoofing detection. Although the focus here and

throughout the rest of the chapter is on manned surface vessels, the conclusions apply with slight modification to unmanned surface vessels.

2.1.1 Compass

The magnetic compass and gyrocompass (a gyroscope designed to be north-seeking by taking advantage of Earth's rotation) depend only weakly on GNSS. A magnetic compass requires knowledge of latitude and longitude to correct for magnetic variation [72]. A gyrocompass requires knowledge of the latitude and speed in the north/south direction to correct for "northing" error [72]. However, outside of the polar regions, position errors on the order of tens of kilometers and velocity errors on the order of meters per second will only cause pointing errors on the order of a degree. Therefore, this work will neither exploit nor model the weak coupling between GNSS and traditional ship compasses.

However, a satellite compass [73], which provides both the position and three-axis attitude of the ship, is fully reliant on GNSS. A common satellite compass comprises two GPS receivers separated by a 0.2–10 meter baseline coupled with miniature accelerometers, gyros, and a magnetometer. The low cost, size, weight, and power consumption of satellite compasses, and the fact that they never require calibration, make these devices an increasingly popular compass option for surface vessels.

2.1.2 Collision Avoidance

The Automatic Radar Plotting Aid (ARPA) is the primary tool used for collision avoidance by the navigator, along with the view from the bridge windows. The ARPA processes and displays the raw radar data in a polar azimuth-range plot, tracks targets, and computes time and distance of closest approach for each target [74]. Without the additional information that sensors like compass, speed log, and GNSS provide, the ARPA can still perform collision-avoidance functions but can only display target information oriented along the ship's heading, the so-called heads-up mode, with relative motion. With compass information, the ARPA can present the radar data oriented along the ship's velocity vector, the so-called course-up mode, which prevents smearing of the returns during course-change manuevers. Similarly, the ARPA can present the radar data in a so-called true motion mode, where the motion is either sea-stabilized by compass and speed log or ground-stabilized by GNSS. Additionally, GNSS information allows the ARPA to compute latitude and longitude for the tracked targets. Nevertheless, convenience features such as ground stabilization and target localization that depend on GNSS signals play a relatively minor role in collision avoidance with other moving targets. Finally, an interesting effect of a GNSS deception attack with ground stabilization enabled on the ARPA makes radar echos from land masses appear to move when they should be stationary.

The Automatic Identification System (AIS) allows ships to communicate their position, heading, and speed in a self-organizing radio network to aid in collision avoidance [72]. A ship's AIS transceiver typically relies on a

GNSS-based positioning source, although it can revert to a pre-determined backup sources during a manifest GNSS failure. Under a GNSS deception attack, a ship may transmit misleading AIS reports and incorrectly compute the point of closest approach to surrounding ships, raising the collision risk. An ARPA can typically overlay the AIS over the radar return and a modern ARPA with integrated AIS can automatically correlate AIS and radar positions into a single target.

2.1.3 Dead Reckoning

Dead reckoning (DR) is the process of propagating a known position based solely on a ship's course and speed, derived from compass and speed log measurements. An estimated position (EP) corrects a dead-reckoned position by applying approximate knowledge of the effects of environmental disturbances such as leeway (drift due to wind), and tidal and ocean currents. Typically, the effects of environmental disturbances are lumped together into a velocity error vector, whose angle and magnitude are referred to as set and drift, respectively. The set and drift can be estimated by comparing a dead-reckoned position to a position fix derived from either a GNSS receiver (typically), observations of celestial bodies, or radar and visual bearings. On paper charts, DR would be reset with a position fix at least every hour, or as often as every three minutes, depending on the accuracy required for navigating the surrounding waters [72]. Electronic chart systems, discussed in the next section, all have the ability to automate DR, making it easier to detect GNSS faults or deception.

2.1.4 Electronic Chart Display and Information System

The Electronic Chart Display and Information System (ECDIS) consolidates the measurements available from various ship sensors and integrates systems such as ARPA, AIS, and DR as shown in Fig. 2.1 to provide complete situational awareness to the ship's crew [72]. ECDIS is the primary tool for route planning and tertiary to the ARPA and AIS for collision avoidance, as mandated by legislation and made explicit in maritime training. Most ECDIS allow overlaying ARPA and AIS information on the charts and planned route for convenience. The overlay may be useful in detecting discrepancies that would arise due to GNSS deception of the own-ship position, e.g., failure of radar returns to match coastal features and buoys on the charts, or the AIS-reported position of nearby vessels. Maritime training emphasizes the need to look for and investigate discrepancies as they normally indicate an equipment problem. But, these discrepancies may simply confuse a crew unaware of GNSS deception, although mariner training manuals have begun to identify GNSS deception as a potential issue for crew to monitor [75]. In any case, when the distance to shore exceeds the range of radar (20 km for low-frequency radar, less for X-band) and when there are few ships nearby, GNSS deception attacks are not likely to be detected solely with radar. Most electronic chart systems such as the Totem ECDIS allow configuring the reset interval of the built-in DR and raising an alarm if the difference between the position fix and DR exceeds a threshold [76]. Section 2.3 will develop an analytically rigorous foundation for this approach by relating the detection threshold to the probability of hazardously misleading information (HMI) for a given false alarm

rate and fix interval.

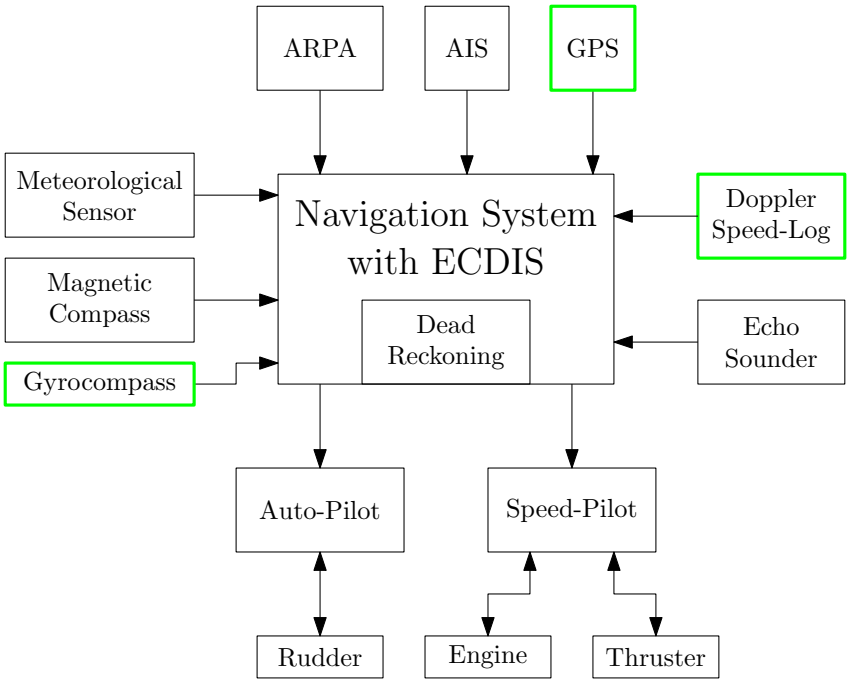

Figure 2.1: Block diagram showing relationship between sensors, actuators, and the ECDIS on a modern integrated bridge system. In this work, only the highlighted sensors, which are used for dead-reckoning, are used for GNSS spoofing detection.

2.1.5 Autopilot System

Virtually all large ships have a course autopilot, which maintains a prescribed heading through rudder actuation in response to compass feedback. Some ships will additionally have a speed autopilot, which maintains a prescribed speed through water by varying the engine thrust in response to feedback from the Doppler speed log sensor. Neither of these rudimentary autopilot systems depends on GNSS directly. However, the course autopilot is typically driven by a higher-level track-keeping system that requires GNSS

feedback. This work focuses on conventional PID-based control systems because they are commonly implemented in practice and typically perform just as well as adaptive model-based control systems under nominal sea and ship conditions [1].

2.1.6 GNSS-Independent Sensors

Sensors which do not have any dependency on GNSS include inertial, acoustic, visual, and meteorological sensors. An inertial sensor found on most ships is a gyroscope-based rate-of-turn (ROT) sensor, which is independent of the compass and GNSS, for derivative course control feedback. The modern speed log uses acoustic Doppler measurements from particles in the water column to compute three-axis speed through water. Other acoustic sensors include conventional downward-looking sonar, also known as an echo sounder, for sea depth measurements and round-trip acoustic ranging to transponders embedded in the sea floor for dynamic positioning [50]. Meteorological sensors provide measurements of temperature, wind, and pressure that can help predict, for example, the effect of leeway and surface currents [72]. Visual bearing measurements of known reference points such as terrestrial landmarks or celestial bodies can be used for positioning. However, landmarks such as buoys can be misidentified, as in the case study in Sec. 2.1.8, or are not available in open waters. Celestial navigation requires knowing the time, either from GNSS or a free-running quartz crystal clock, to look up the position of celestial bodies from an almanac [72]. A jump in ship time by 5 seconds (e.g., due to leap second spoofing) would cause a longitude error of 0.02 degrees. Nevertheless,

errors less than ten seconds from either a drifting or GNSS-deceived clock are comparable with other errors in celestial navigation.

These GNSS-independent sensors feed into alternative position sources that can be used to cross-check GNSS in a modern integrated bridge system. However, a subtle-enough spoofing attack can be consistent with dead reckoning or celestial navigation and thus escape detection. Also, acoustic positioning is only useful for vessels operating in the small neighborhood of the transponders. Although this work's focus is on GNSS deception, it is worth mentioning that radar and acoustic sensors systems on modern civil surface vessels are also vulnerable to deception and jamming. Thus, although these systems are assumed herein to be trustworthy and potentially useful for detecting GNSS deception, a more thorough security analysis would need to consider a coordinated, self-consistent attack on GNSS, radar, and acoustic sensors.

2.1.7 Summary of GNSS Deception Vulnerabilities

The ship's crew can cross-check GNSS with (1) compass, (2) speed log, (3) ship dynamics model, (4) radar and AIS from other ship, (5) radar and charts, (6) echo sounder, (7) meteorological sensors. But, as will be shown later on, even an optimal combination of (1)–(3), which amounts to sophisticated DR, would not be sufficient to reliably detect a subtle attack before the ship's positioning error exceeds a reasonable hazardous condition threshold. If (4) and (5) are properly and fully exploited, then the security situation improves significantly. But alignment of charted objects such as the shoreline and buoys with radar returns is often quite poor even under normal condi-

tions because of (i) shoreline changes with tide, (ii) inadequate resolution of charts, and (iii) positioning, bearing, and radar-ranging errors. Consequently many ships' crews either do not attempt radar overlay or would not consider it a trustworthy cross-check for own-ship positioning errors. Also, comparing radar with AIS from other ships is not trustworthy because (i) AIS data can be manipulated, (ii) AIS-repeated location data ultimately depends on GNSS, and (iii) ships' crews are accustomed to discrepancies in AIS data and so many have come to mistrust it.

For avoidance of collisions with radar-reflective objects, the ARPA remains trustworthy and its collision avoidance function does not depend on GNSS. But cross-track ship excursions outside the planned corridor are nevertheless dangerous precisely because some threatening objects (e.g., underwater hazards) are not visible to radar and will not be detected by downward-looking sonar. Moreover, along-track errors in a ship's position can also be hazardous because such errors can confuse the interpretation of radar returns or cause a ship to over- or under-shoot the location of a planned maneuver.

2.1.8 Illustrative Example: The Grounding of the Royal Majesty

To appreciate the possible effects of a GNSS deception attack on a surface vessel, it is instructive consider the grounding of the 174-meter cruise liner Royal Majesty [47, 77]. Shortly after the ship departed Bermuda for Boston in June of 1995, the cable connecting its GPS antenna to the unit on the bridge became detached, forcing the GPS unit to transition to a dead-reckoning mode in which the ship's location was extrapolated from the last known good loca-

tion based solely on gyro compass and water speed measurements. The crew and autopilot, unaware of the transition to DR mode, accepted the position indicated on the radar display's map as truthful even as the ship accumulated a 31 km cross-track navigational error. As the ship approached Nantucket, the crew misidentified one buoy and ignored the absence of another. The ship's GPS-based navigation system had performed so utterly reliably in the past that the crew's trust in the ship's displayed position was not shaken even as a lookout sighted blue and white water ahead. Minutes later, the ship ran aground on shoals invisible to its radar system.

In the aftermath of the Royal Majesty grounding, integrated bridge systems were modified to more clearly indicate loss of GNSS signals, and redundant GNSS units became standard. In addition, the incident is used as an important lesson on the dangers of over-reliance on GNSS in maritime training colleges since the crew of the Royal Majesty clearly acted in a manner inconsistent with proper training, a contributing cause to the incident. Nevertheless, the risk of a repeat of the Royal Majesty grounding, or a similar incident, caused by deliberate, strategic GNSS deception remains because there would be no apparent loss of GNSS, the DR would appear to remain consistent with GNSS, and because primary and backup GNSS units would be equivalently affected. The possibility of an improperly-trained or fatigued crew encountering a GNSS failure that does not trigger alarms can never be completely eliminated, and proper quantification of this risk under a GNSS deception attack is needed.

Having offered an overview of the possible effects of GNSS deception on

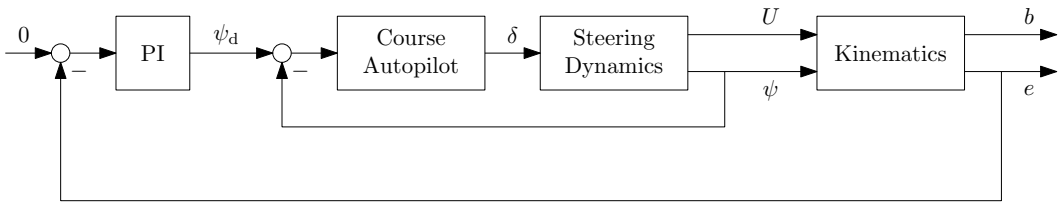

Figure 2.2: Conventional track-keeping system based on an existing course autopilot system [1, p. 293]. Here, ψ_d is the desired heading angle, δ is the rudder angle, U is the ship speed through water, ψ is the heading angle, b is the along-track position, and e is the cross-track position.

surface vessels, this chapter now turns to developing a framework for analysis of GNSS spoofing detection based on comparison of GNSS data with a modified version of DR. This detection strategy is appealing because of its broad applicability: all sizable surface vessels can perform at least rudimentary DR, and the DR technique works both far from shore and in littoral waters. The next two sections introduce the dynamics model and the detection framework.

2.2 Ship and Spoofing Model

Consider a simplified ship dynamics model with a conventional track-keeping guidance system as presented in [1]. A conventional track-keeping system attempts to zero the ship's cross-track position using a proportional-integral (PI) controller wrapped around a course autopilot, as shown in Fig. 2.2.

2.2.1 Ship Dynamics

The ship dynamics model presented here, although quite simplified compared to a more expressive six degree-of-freedom model, captures the low-frquency ship motion relevant for control and spoofing. The ship's steering

dynamics is described by a 1st order Nomoto model [1],

$$T\dot{r} + r = K\delta + r_{\mathrm{b}},$$

where T is the ship's time constant (s), K is the rudder gain ($^{1}/_{\mathrm{s}}$), δ is the rudder angle (rad), r is the ship's turn rate ($^{\mathrm{rad}}/_{\mathrm{s}}$), and r_{b} is a slowly-varying parameter that models environmental disturbances ($^{\mathrm{rad}}/_{\mathrm{s}}$). The rudder angle δ and angular rate $\dot{\delta}$ are physically constrained by saturation conditions $|\delta| < \delta_{\max}$ and $|\dot{\delta}| < \dot{\delta}_{\max}$, respectively, but the controller is designed such that the rudder angle dynamics remain linear under typical conditions. The ship's kinematics are given by [1]

$$\begin{aligned}\dot{\psi} &= r \\ \dot{x} &= U\cos\psi + d_x \\ \dot{y} &= U\sin\psi + d_y,\end{aligned}$$

where U is the ship's speed through water ($^{\mathrm{m}}/_{\mathrm{s}}$), d_x and d_y model errors due to drift caused by slowly-varying environmental disturbances such as ocean currents and wind ($^{\mathrm{m}}/_{\mathrm{s}}$), x and y are the ship's northing and easting (m), respectively, and ψ is the ship's heading (rad). Zero heading is defined to be due north with increasing heading clockwise. The environmental disturbance parameters are modeled as Gauss-Markov processes,

$$\begin{aligned}\dot{d}_x &= -\frac{1}{T_d}d_x + v_x \\ \dot{d}_y &= -\frac{1}{T_d}d_y + v_y,\end{aligned}$$

where T_d is the disturbance time constant and v_x and v_y are additive white Gaussian noise (AWGN) sources with intensity σ_d^2 ($^{\mathrm{m}^2}/_{\mathrm{s}^3}$).

2.2.2 Ship Control Laws

Only conventional controllers are considered in the sequel because they perform just as well as adaptive and non-linear model-based controllers under nominal sea and ship conditions [1]. A conventional course autopilot controls the ship's heading ψ to a desired approximately-constant heading ψ_{d} using a proportional-integral-derivative (PID) control law. In modeling the course control law that follows, and the track-keeping control law presented thereafter, the measurements are assumed to be noiseless and continuous since the low-bandwidth controllers and ship dynamics act to suppress the effects of real-world discretization and measurement noise at the output of each closed-loop control system. The measurements $\psi(t)$ and $r(t)$ from the compass and ROT sensor, respectively, control the rudder angle $\delta(t)$ according to

$$\delta(t) = K_{\mathrm{i}} \int_0^t \left[\psi_{\mathrm{d}} - \psi(\tau)\right] d\tau + K_{\mathrm{p}} \left[\psi_{\mathrm{d}} - \psi(t)\right] - K_{\mathrm{d}} r(t),$$

where K_{i} is the integral gain, K_{p} is the proportional gain, and K_{d} is the derivative gain. Following conventional PID control design of second-order systems [1, p. 261], these gain parameters are derived from a chosen natural frequency ω_{n} and relative damping ratio ξ of the closed-loop system; the latter is typically chosen in the interval $0.8 \leq \xi \leq 1.0$. The closed-loop bandwidth ω_{b}, defined as

$$\omega_{\mathrm{b}} \triangleq \omega_{\mathrm{n}} \sqrt{1 - 2\xi^2 + \sqrt{4\xi^4 - 4\xi^2 + 2}},$$

is chosen such that

$$\frac{1}{T} < \omega_{\mathrm{b}} < \omega_{\delta},$$

where $\omega_\delta \triangleq \frac{\dot{\delta}_{\max}}{\delta_{\max}}$ is the rudder servo bandwidth. Finally, the PID gains are related to ω_n and ξ by

$$\begin{aligned} K_\mathrm{p} &= \frac{T}{K}\omega_\mathrm{n}^2 \\ K_\mathrm{d} &= \frac{1}{K}\left[2T\xi\omega_\mathrm{n} - 1\right] \\ K_\mathrm{i} &= \frac{T}{K}\frac{\omega_\mathrm{n}^3}{10}. \end{aligned}$$

An outer control loop for track-keeping is typically wrapped around the course autopilot. In some cases, a human operator in the loop may take the role of track-keeping controller. Whether mechanical or human, the controller can be modeled as a PI controller. The track, or rhumb line, can be approximated in the local Cartesian coordinates by a ray, which is parametrized by an angle ψ_0 (rad) and start position x_0 and y_0 (m). The along-track and cross-track position, b and e, respectively, are given by

$$\begin{aligned} b &= (x - x_0)\cos\psi_0 + (y - y_0)\sin\psi_0 \\ e &= (y - y_0)\cos\psi_0 - (x - x_0)\sin\psi_0. \end{aligned}$$

The relationship between the global and track coordinates is illustrated graphically in Fig. 2.3. Because GNSS is the most accurate positioning source, nearly always available, and assumed to be reliable when available, it is typically the primary positioning source [49]. The GNSS receiver's cross-track position measurement, which is taken to be equivalent to $e(t)$, is fed back with a PI control law given by

$$\psi_\mathrm{d}(t) = \psi_0 - K_\mathrm{i}' \int_0^t e(\tau)d\tau - K_\mathrm{p}' e(t),$$

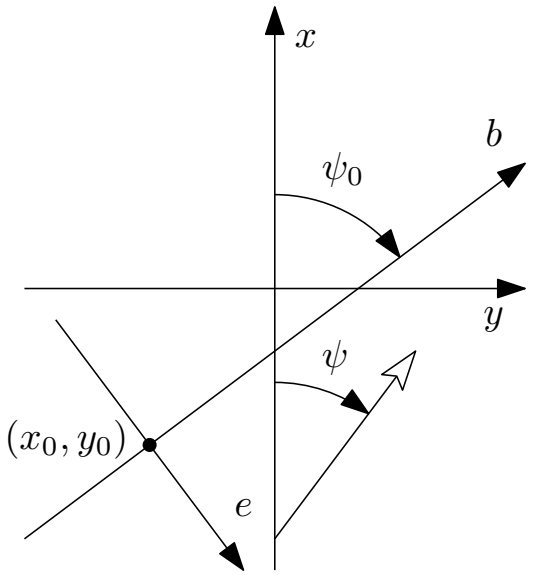

Figure 2.3: Coordinate systems for ship global position (x, y) and track position (b, e). The track coordinate system's origin and rotation with respect to the global coordinate system is given by (x_0, y_0) and ψ_0, respectively. The ship's orientation with respect to the global coordinate system is given by heading angle ψ.

where K'_{i} is the integral gain and K'_{p} is the proportional gain. The gains are chosen so that the inner course control loop and the outer track-keeping loop have significant time scale separation, with the inner loop faster, a typical practice for marine and aerial vehicle cascaded controller design [1, 78]. Thus, from the perspective of the outer loop, one can assume $\psi \approx \psi_{\mathrm{d}}$, and the full closed-loop cross-track dynamics can be approximated by a first-order system with bandwidth $\omega'_{\mathrm{b}} = UK'_{\mathrm{p}} \ll \omega_{\mathrm{b}}$. Note that the along-track position is not controlled by a feedback law but instead proceeds open-loop according to an approximate crew-selected velocity setpoint.

2.2.3 Spoofer Control Law

In a spoofing attack, the ship's GNSS receiver will report the position commanded by the spoofer. To remain covert, the spoofer will typically command positions that are gentle deviations, conveniently represented in along-track and cross-track coordinates, from the ship's true position. Cross-track deviations will prompt a response from the ship's track-keeping controller whereas along-track deviations will elicit no response unless the ship's track changes. Along-track spoofing can be an effective strategy from the point of view of the attacker, but this chapter will focus on cross-track spoofing because it is equally effective yet requires less knowledge of the ship's route.

In a cross-track spoofing attack, the spoofer generates a GNSS signal whose implied coordinates are the attacker's estimate of the ship's actual along-track position $\hat{b}_{\mathrm{a}}(t)$ and a spoofed cross-track position $e_{\mathrm{s}}(t)$. The latter can be written as the difference of two parts, the attacker's estimate of the ship's true cross-track position $\hat{e}_{\mathrm{a}}(t)$ and a spoofer-induced cross-track modulation $e_{\mathrm{m}}(t)$ so that $e_{\mathrm{s}}(t) = \hat{e}_{\mathrm{a}}(t) - e_{\mathrm{m}}(t)$. The modulation $e_{\mathrm{m}}(t)$ is also called the attack profile, as it represents the spoofer-intended depature from the cross-track position. Note that to form $e_{\mathrm{a}}(t)$ the attacker must continuously estimate both the ship's position and its rhumb line. This assumption is not particularly demanding: other-ship position estimation via radar is both accurate and routine, and surface vessels typically follow a route consisting of waypoints connected by readily-estimable lines of constant bearing.

The attacker's goal is to force the ship to track a spoofer-commanded

cross-track position, denoted $\bar{e}$, as quickly as possible without being detected. He evades detection by generating a subtle $e_{\mathrm{m}}(t)$ with limited velocity and acceleration magnitudes:

$$|\dot{e}_{\mathrm{m}}(t)| \leq v_{\max}, \quad |\ddot{e}_{\mathrm{m}}(t)| \leq u_{\max} \tag{2.1}$$

Solving the following minimum-time optimal-control problem yields the attack profile $e_{\mathrm{m}}(t)$ that achieves the spoofer's goal. Here, t_{f} is the final time and the control input $u(t)$ enters through the second derivative of $e_{\mathrm{m}}(t)$ as part of the dynamical constraint.

$$\begin{aligned} \min_{u(t)} \quad & t_{\mathrm{f}} \\ \text{s. t.} \quad & \ddot{e}_{\mathrm{m}}(t) = u(t) \\ & e_{\mathrm{m}}(0) = 0, \ \dot{e}_{\mathrm{m}}(0) = 0 \\ & e_{\mathrm{m}}(t_{\mathrm{f}}) = \bar{e}, \ \dot{e}_{\mathrm{m}}(t_{\mathrm{f}}) = 0 \\ & |\dot{e}_{\mathrm{m}}(t)| \leq v_{\max} \\ & |\ddot{e}_{\mathrm{m}}(t)| \leq u_{\max}. \end{aligned} \tag{2.2}$$

For $\bar{e} \to \infty$ and $t_c \triangleq v_{\max}/u_{\max}$, the solution is given by

$$e_{\mathrm{m}}(t) = \begin{cases} \frac{1}{2} u_{\max} t^2 & 0 < t \leq t_c \\ \frac{1}{2} u_{\max} t_c^2 + v_{\max}(t - t_c) & t_c < t \end{cases}$$

An attack profile generated as a solution to (2.2) is easily disguised as the effect of ocean currents. But it may not be optimal from the point of view of the attacker; i.e., it may not be the most hazardous undetectable profile. The optimal profile in this sense actually depends on the defender's particular detection test. Other strategies for generating $e_{\mathrm{m}}(t)$ that are more directly

related to plausible detection tests are considered in Appendix A.4. Nonetheless, the strategy outlined in (2.2) has the virtue of being intuitive and readily implementable yet generates $e_{\mathrm{m}}(t)$ profiles similar to those produced by the more complex strategies.

The maxima $v_{\max}$ and $u_{\max}$ are assumed to be sufficiently small that $e_{\mathrm{m}}(t)$ is slow compared to the time constant of the ship's track-keeping control law, ensuring the attacker can dictate the ship's true cross-track position $e(t)$ with only modest errors—errors due to the spoofer's imperfect estimate of the ship's true position and the rhumb line, and to the ship's own estimation and control errors. Under this assumption, the spoofer need not adapt $e_{\mathrm{m}}(t)$ to the ship's response but may simply generate $e_{\mathrm{m}}(t)$ open loop. A closed-loop spoofing controller is also possible, but its attacks are more difficult to maintain covert, as explained in [19].

2.3 Detection Framework

The detection framework developed in this chapter attempts to minimize the mean integrity risk $\bar{I}_R$, defined subsequently, for a given continuity risk $C_R \triangleq 1/M_F$, where M_F is the mean time between false alarms. This framework borrows concepts from GNSS integrity monitoring in aviation applications [66] and the fault detection literature [62], which are applied here to the "fraud detection" problem. Typically, the integrity and continuity risk are specified in terms of the probability of hazardously misleading information (HMI) per approach and the false-alarm rate, respectively.

2.3.1 Overview

The schematic in Fig. 2.4 offers a graphical overview of the detection problem. Time $t = 0$ denotes the beginning of an approach, or part of a journey, such as the final approach to a harbor. At each time $t_k = kT_s,\ k = 0, 1, ...,$ a detection test is performed to decide between two hypotheses—a null hypothesis H_0 indicating nominal operating conditions, and an alternative hypothesis H_1 indicating a spoofing attack is underway. At the beginning of the approach, the null hypothesis is true; at some time $t_0 \geq 0$, a transition to the alternative hypothesis occurs. After t_0, the attack continues until either hazardous conditions occur or the attack is detected. In this framework, the constant time between tests T_s is a key parameter: it is taken as the free parameter for the integrity optimization problem.

The detection strategy envisioned here is decoupled from the ship's track-keeping controller, which is assumed to ingest GNSS measurements at its usual rate—typically much faster than $1/T_s$—without regard for the the periodic detection tests ocurring in parallel. A joint control-and-detection framework is possible and would have slightly superior performance compared to the proposed framework, but a disjoint framework is simpler and has the benefit of being applicable to existing ships without re-certification of their integrated bridge systems.

So long as the detection statistic $q(k)$ remains below a threshold λ, the detector assumes the null hypothesis; otherwise, it assumes the alternative hypothesis and continuity is broken as the crew attempts to neutralize the

potential spoofing threat. The threshold λ is chosen to satisfy

$$P\left(q(k) > \lambda | H_0\right) = T_s / M_F = C_R T_s$$

to maintain the prescribed false-alarm rate. Note that the probability distribution of $q(k)$ under H_0 is independent of k, so λ need not depend on k.

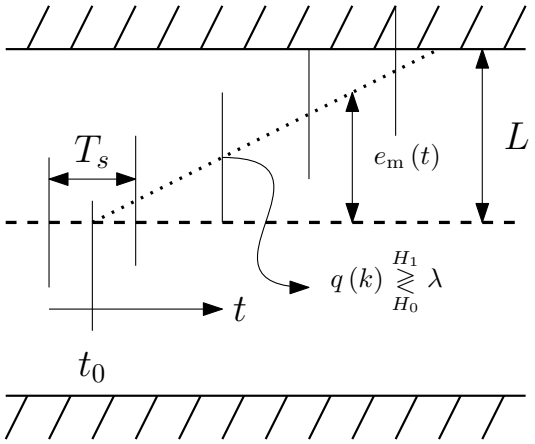

Figure 2.4: A graphical overview of the detection problem: A spoofing attack with cross-track profile $e_{\mathrm{m}}(t)$ begins at time t_0, which is unknown to the defender. The attacker attemps to drive the ship to exceed the alert limit L, beyond which lie potential hazards, without detection. At every time $t_k = kT_s,\ k = 0, 1, ...$, a GNSS measurement is taken and used to form the detection statistic $q(k)$. The time instants t_k are unknown to the attacker, though the measurement period T_s may be known. If $q(k)$ exceeds the threshold λ, the alternative hypothesis H_1 (spoofing attack) is declared; otherwise, the null hypothesis H_0 is assumed.

2.3.2 Integrity Risk

Leading up to a definition of mean integrity risk $\bar{I}_R$, it will be useful to define what is meant by hazardous conditions and by a so-called local HMI event. Let the total system error of a certain state element of interest be

denoted $\epsilon(t)$. The total system error is the departure of the true state element from the controller's desired value of that state element, and includes both estimation and control errors. Hazardous conditions are said to occur when $|\epsilon| > L$ for an alert limit L. Although the ship may not be in immediate danger if $|\epsilon| > L$, control decisions based on such divergent estimates are highly risky. In this chapter, the state of interest is the cross-track position $e(t)$, and a typical value for L may be 1 km. To account for worst-case control error, L must be substantially smaller than the distance that the ship's route clears charted hazards.

Assuming GNSS measurements are continuously available, as in the ship's control model, and that control errors remain small, then under H_1, $\epsilon(t) \approx e_{\mathrm{m}}(t)$. This deterministic approximation is a key simplifying assumption: it prevents the total system error from being correlated with the detection statistic. Lack of correlation greatly simplifies the expression for the mean integrity risk, as will be shown subsequently.

A local HMI event $E(t)$ for $t > t_0$ is defined as hazardous conditions under a spoofing attack that has not been detected. Mathematically, $E(t)$ is expressed as

$$E(t) \triangleq (|e_{\mathrm{m}}(t)| > L) \wedge \left(\bigwedge_{k \in S_t} q(k) < \lambda \right),$$

where $S_t \triangleq \{k | t_0 < kT_s < t\}$. The boolean event G indicates whether a local HMI event has occurred at any time $t > t_0$ during an approach:

$$G \triangleq \bigvee_{t > t_0} E(t).$$

Let the first time hazardous conditions occur under a spoofing attack be denoted t_L and let $S_L \triangleq \{k|t_0 < kT_s \leq t_L\}$. Then G can be reformulated as

$$G = \bigwedge_{k \in S_L} q(k) < \lambda.$$

Integrity risk is defined for a particular start time t_0 as $I_R(t_0) \triangleq P(G|t_0)$, assuming conservatively that the probability of a spoofing attack is unity. Assuming all spoofing start times are equally likely, mean integrity risk is then defined as

$$\bar{I}_R \triangleq \int_0^1 P(G|t_0 = \beta T_s) d\beta.$$

2.3.3 Detection Statistic

Detector performance depends strongly on the detection statistic $q(k)$. If the attack profile $e_{\mathrm{m}}(t)$ were precisely known to the defender *a priori*, then a detection statistic could be optimally tailored to the known profile. The statistic would amount to processing estimator innovations through a filter matched to the known profile Appendix A.3. If some attack profile parameters remained unknown, such as t_0 and $v_{\max}$, then the generalized likelihood ratio approach would be reasonable [62]. However, the stronger the defender's assumptions are about the attack profile, the more vulnerable he becomes to an attacker who violates those assumptions.

One recognizes a zero-sum game in the simultaneous incentives the defender has to optimize $q(k)$ for the defender's choice of $e_{\mathrm{m}}(t)$ and the attacker has to optimize $e_{\mathrm{m}}(t)$ for the defender's choice of $q(k)$. If an equilibrium pair $\{q^\star(k), e^\star_{\mathrm{m}}(t)\}$ were found to exist for this game, such that neither attacker

nor defender would benefit by unilateral departure from the equilibrium, then $q^\star(k)$ could be taken as an equilibrium-optimal detection statistic [79, 80]. However, the author was unable to discover such an equilibrium; its existence remains an open question. Instead, a normalized-innovations-squared (NIS) statisic [70, 62, 81] is adopted here. This statistic is not optimal in the sense of $q^*(k)$, but it is robust in that it makes no assumptions about the attack trajectory; rather, it penalizes all departures from the assumed model.

The innovation sequence $\nu(k)$ on which $q(k)$ is based is generated by a Kalman filter ingesting a GNSS measurement every T_s seconds. A simplified model for the Kalman filter is developed below in preparation for determining the probability distribution of $q(k)$. First, consider the continuous-time ship dynamics model

$$\dot{\eta}(t) = A\eta(t) + Bu(t) + \Gamma\tilde{v}(t),$$

where

$$\begin{aligned}
&\eta = \begin{bmatrix} x & y & d_x & d_y \end{bmatrix}^{\mathrm{T}} \text{ is the state vector,} \\
&A = \begin{bmatrix} 0 & I_2 \\ 0 & -\frac{1}{T_d} I_2 \end{bmatrix}, B = \begin{bmatrix} I_2 \\ 0 \end{bmatrix}, \Gamma = \begin{bmatrix} 0 \\ I_2 \end{bmatrix}, \\
&u = U \begin{bmatrix} \sin\psi & \cos\psi \end{bmatrix}^{\mathrm{T}} \text{ is the control, and} \\
&\tilde{v} = \begin{bmatrix} v_x & v_y \end{bmatrix}^{\mathrm{T}} \text{ is AWGN with intensity } Q_c = \sigma_d^2 I_2,
\end{aligned}$$

with I_n the n-by-n identity matrix and the other matrices appropriately dimensioned. The control $u(t)$ is derived from the ship's compass and speed log measurements. The potentially-spoofed GNSS measurements are sampled from

$$z(k) = H\eta(kT_s) - z_{\mathrm{m}}(kT_s) + w(k),$$

where $w(k)$ is a discrete AWGN sequence with covariance $R = \sigma_p^2 I_2$, $H = \begin{bmatrix} I_2 & 0 \end{bmatrix}$, and $z_{\mathrm{m}}(t)$ is the deterministic spoofer-induced two-dimensional position modulation for which, by definition, $z_{\mathrm{m}}(t) = 0$ for $t < t_0$.

The *a priori* and *a posteriori* estimation errors of the Kalman filter are defined as $\bar{\epsilon}(k) \triangleq \eta(k) - \bar{\eta}(k)$ and $\hat{\epsilon}(k) \triangleq \eta(k) - \hat{\eta}(k)$, respectively, while the innovation $\nu(k)$ is defined as

$$\nu(k) \triangleq H\bar{\epsilon}(k) - z_{\mathrm{m}}(kT_s) + w(k).$$

The recursion equations for the estimation error's means and covariances are given by

$$\begin{aligned}
\mathbb{E}\left[\bar{\epsilon}(k)\right] &= F\mathbb{E}\left[\hat{\epsilon}(k-1)\right] \\
\bar{P}(k) &\triangleq \mathbb{E}\left[\bar{\epsilon}(k)\bar{\epsilon}^T(k)\right] = FP(k-1)F^{\mathrm{T}} + Q \\
\mathbb{E}\left[\hat{\epsilon}(k)\right] &= (I - K(k)H)\mathbb{E}\left[\bar{\epsilon}(k)\right] - K(k)z_{\mathrm{m}}(kT_s) \\
P(k) &\triangleq \mathbb{E}\left[\hat{\epsilon}(k)\hat{\epsilon}^T(k)\right] = (I - K(k)H)\bar{P}(k),
\end{aligned}$$

where

$$\begin{aligned}
F &= e^{AT_s}, \\
Q &= \int_0^{T_s} e^{A\tau}\Gamma Q_c \Gamma^{\mathrm{T}} e^{A^{\mathrm{T}}\tau} d\tau, \\
S(k) &= H\bar{P}(k)H^{\mathrm{T}} + R, \\
K(k) &= \bar{P}(k)H^{\mathrm{T}}S^{-1}(k), \text{ and} \\
\mathbb{E}\left[\bar{\epsilon}(0)\right] &= 0.
\end{aligned}$$

In the sequel, it is assumed that the estimation error covariances have reached their steady-state values, which can be found by solving a discrete-time algebraic Riccati equation, and so the index k is dropped from P, $\bar{P}$, S, and K.

Note that during a spoofing attack, a nonzero $z_{\mathrm{m}}(kT_s)$ biases the estimation error and innovation.

The NIS detection statistic, defined as, $q(k) = \nu^{\mathrm{T}}(k)S^{-1}\nu(k)$, is distributed under H_0 as χ^2 with two degrees of freedom, and under H_1 as non-central χ^2 with two degrees of freedom and non-centrality parameter $\delta(k) = \bar{\nu}^{\mathrm{T}}(k)S^{-1}\bar{\nu}(k)$, where

$$\bar{\nu}(k) = H\mathbb{E}\left[\bar{\epsilon}(k)\right] - z_{\mathrm{m}}(kT_s)$$

is the mean of innovation at index k. Since the innovation sequence is white, each detection test is independent.

2.3.4 Optimization

A natural question arises in sequential detection: How often should the detection test be executed? If the measurement sampling interval T_s, which is also the detection test interval, is too small, then no innovation $\nu(k)$ will appear particularly surprising under H_1, and no cumulative effect will accrue as the Kalman filter accepts and adjusts to the biased innovations. As T_s is made longer, innovations under H_1 become more obviously biased, but if T_s is too long, the attacker may begin an attack and reach hazardous conditions between detection tests.

A distinguishing feature of the current framework is that it optimizes T_s to minimize $\bar{I}_R$ over a range of possible $v_{\max}$. Fig. 2.5 shows how $\bar{I}_R$ varies as a function of T_s and $v_{\max}$ for the parameters indicated. The optimal T_s is a minimax solution which minimizes the maximum $\bar{I}_R$ over the range of $v_{\max}$

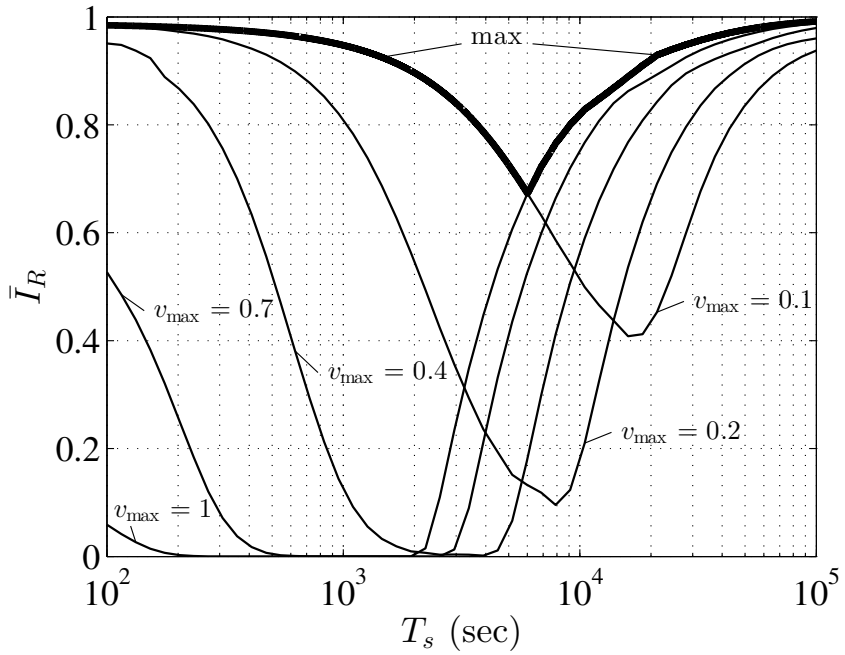

Figure 2.5: Mean integrity risk $\bar{I}_R$ vs. sampling time T_s for various choices of $v_{\max}$. The optimal sampling time $T_s^\star$ that minimizes the worst-case mean integrity risk is approximately 100 minutes, yielding $\bar{I}_R^\star \approx 0.6727$. Note that the worst-case attack is given by either $v_{\max} = 0.1$ or $1\,\mathrm{m/s}$. Other parameters are $u_{\max} = 0.03\,\mathrm{m/s^2}$, $M_F = 1\,\mathrm{month}$, $\bar{e} \gg L = 3\,\mathrm{km}$, $\sigma_p = 6\,\mathrm{m}$, $T_d = 200\,\mathrm{s}$, and $\sigma_d = 0.02\,\mathrm{m/s^{1.5}}$.

considered, in this case $0.1\,\mathrm{m/s}$ to $1\,\mathrm{m/s}$. More formally, a robust optimizer for T_s would be

$$\min_{T_s} \max_{\substack{v_{\max}\in\mathbb{V} \\ u_{\max}\in\mathbb{U}}} \bar{I}_R, \tag{2.3}$$

where $\mathbb{V}$ and $\mathbb{U}$ are bounded sets containing reasonable values for the attack parameters. $v_{\max}$ is bounded from below under the assumption that the attacker wishes to cause hazardous conditions before the end of a typical approach. If the average duration of an approach is $\bar{T}_{\mathrm{app}}$, then $v_{\max} \geq L/\bar{T}_{\mathrm{app}}$, reasonably assuming a linear relationship between the approach's alert limit and aver-

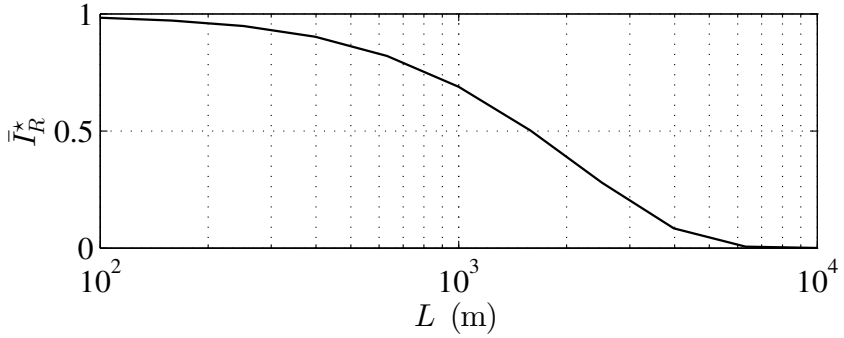

Figure 2.6: Minimax integrity risk $\bar{I}_R^\star$ vs. the hazardous condition threshold L. For $L \leq 400\,\text{m}$, the worst-case attack will likely cause HMI since $\bar{I}_R^\star > 0.9$. On the other hand, $L \geq 7\,\text{km}$ maintains an integrity risk near zero for any reasonable attack. Other parameters are set to the values indicated in Fig. 2.5.

age duration. Induced velocities greater than $1\,\text{m/s}$ would lead to physically impossible set and drift values that are not captured by the Gauss-Markov disturbance model and break the small control error assumption; this constraint places an upper bound on $v_{\max}$. Lastly, the impact of the acceleration regime of the attack is diminished for large T_s since the regime only occurs for a short period of time in the beginning of the attack. Therefore, the integrity risk optimization is not particularly sensitive to the choice of $u_{\max}$, which is fixed to a value of $0.03\,\text{m/s}^2$ for the rest of the analysis.

A closed form solution to (2.3) does not appear possible, but the optimal T_s can be found numerically based on the definition of $\bar{I}_R$ and on the known distributions for $q(k)$ under H_0 and H_1. Minimax results for two example scenarios are shown in Figs. 2.6 and 2.7.

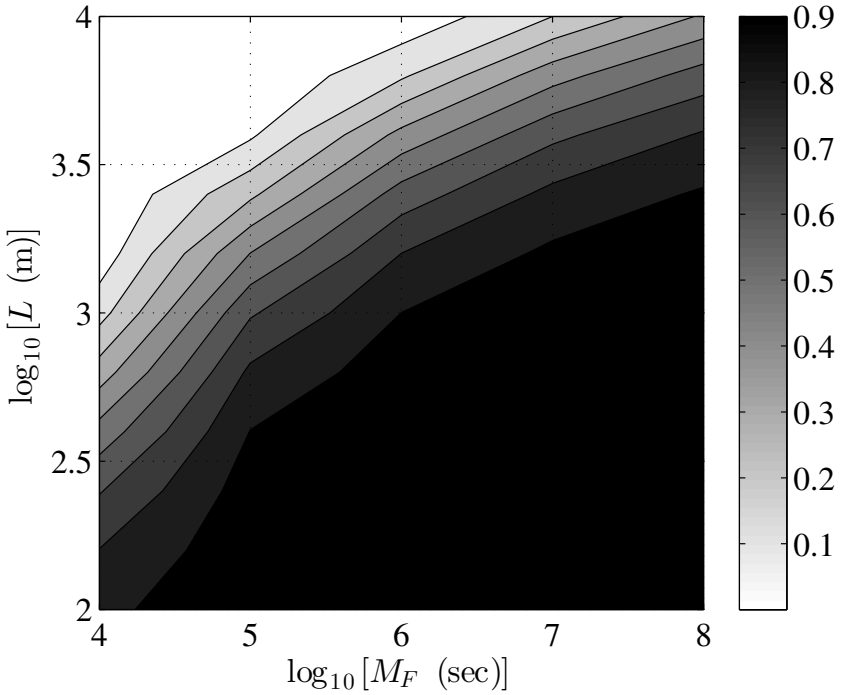

Figure 2.7: Minimax integrity risk $\bar{I}_R^\star$ vs. L and M_F. Depending on the alert limit and continuity risk requirements of the approach, the detection framework will maintain an integrity risk that can be either quite high (black region), in which covert attacks are possible, or quite low (white region). Other parameters are set to the values indicated in Fig. 2.5.

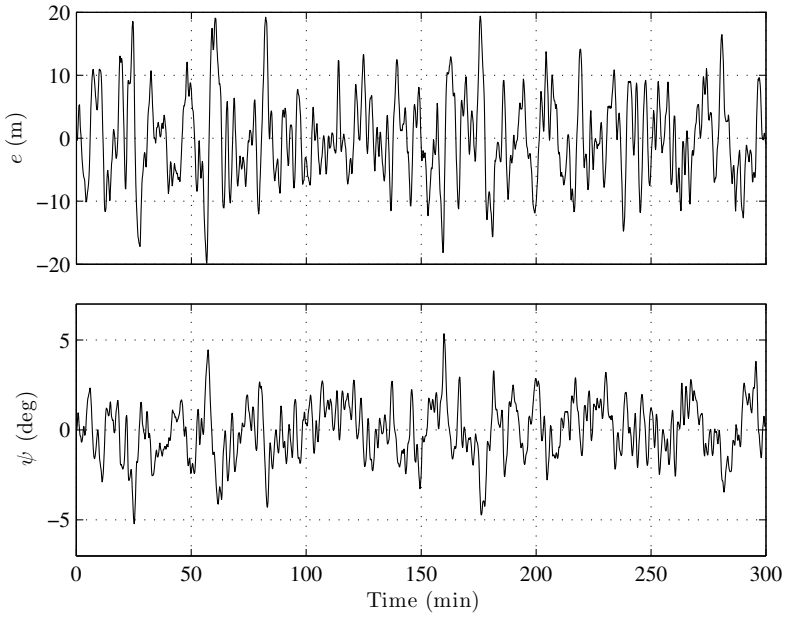

(a) Case I, no spoofing

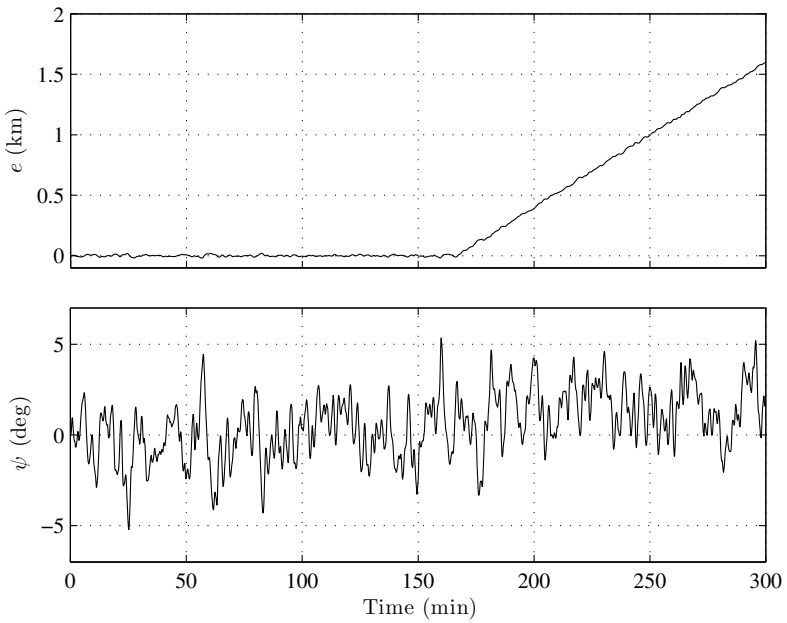

(b) Case II, $v_{\max} = 0.2$ m/s

Figure 2.8: Trajectory resulting from simulation of ship dynamics under nominal conditions and a spoofing attack. Model parameters are given by $T = 39.94\,\mathrm{s}$, $K = 0.211\,\mathrm{s}^{-1}$, $U = 8.23\,\mathrm{m/s}$, $K_\mathrm{p} = 1.4415$, $K_\mathrm{i} = 0.0126$, $K_\mathrm{d} = 21.6904$, $K'_\mathrm{p} = 0.0028$, $K'_\mathrm{i} = 1.8949 \times 10^{-5}$. Other parameters are set to the values indicated in Fig. 2.5.

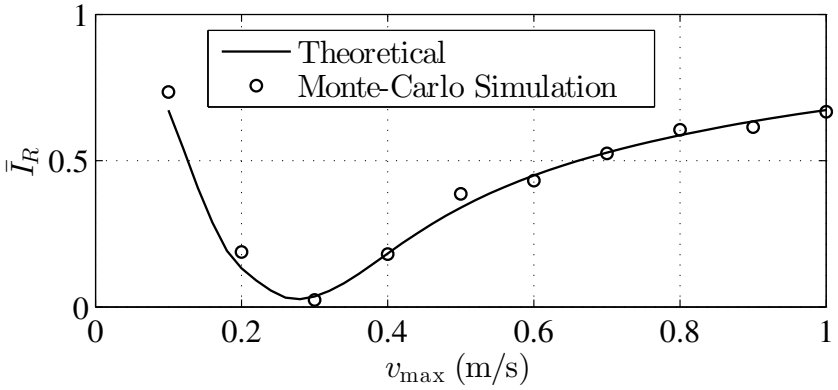

Figure 2.9: Theoretical vs. simulated integrity risk for different values of $v_{\max}$. Other parameters are set to the values indicated in Fig. 2.5.

2.4 Simulation

The spoofer control law and integrity risk calculations were verified with Monte-Carlo simulations. The simulations take into account the Nomoto ship model, closed-loop ship controller, and open-loop spoofer controller developed in Sec. 2.2, with $\bar{e} \gg L$. A couple of representative ship trajectories are shown in Fig. 2.8. The simulation-based integrity risk is determined by counting the number of HMI events over 100 sampling phases per simulation and 20 simulations per attack profile. As shown in Fig. 2.9, the simulation-based integrity risk for different values of $v_{\max}$ agrees quite well with the values predicted by the theory developed in Sec. 2.3.

2.5 Experiment

Covert control of a marine vessel by GPS spoofing was demonstrated in the Mediterranean sea in the summer of 2013. The author was invited to conduct the unprecedented experiment aboard the White Rose of Drachs, a

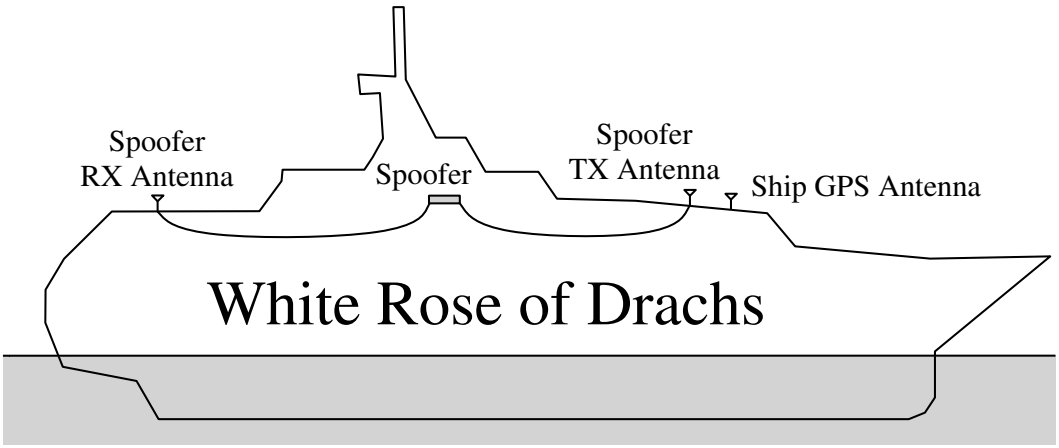

Figure 2.10: Sketch of the spoofer setup on the White Rose of Drachs.

65-meter superyacht. The experimental setup centered on the receiver-spoofer developed at the University of Texas at Austin [19]. The spoofer receives the authentic signals from an antenna placed in the stern. The spoofer transmits the false navigation signals towards the bow, where the ship's GPS antennas are located as shown in Fig. 2.10.

Once a safe route is established, the captain attempts to maintain the ship's reported position along a series of rhumb lines within some prescribed corridor. Control actions at sea are required to maintain course due to disturbances such as wind and ocean currents, which are typically not measured directly. Instead, the disturbance sources are lumped together, and measured indirectly through the GPS. Therefore, a spoofing attack can induce false disturbances, causing the captain to believe the ship is on course, when in reality, the ship is slowly drifting off course. In the aforementioned experiment, a spoofer-induced velocity was introduced in the cross-track direction—at first $0.5\,\mathrm{m/s}$, then increased to $2\,\mathrm{m/s}$ at $200\,\mathrm{m}$, and finally reset to zero at $700\,\mathrm{m}$.

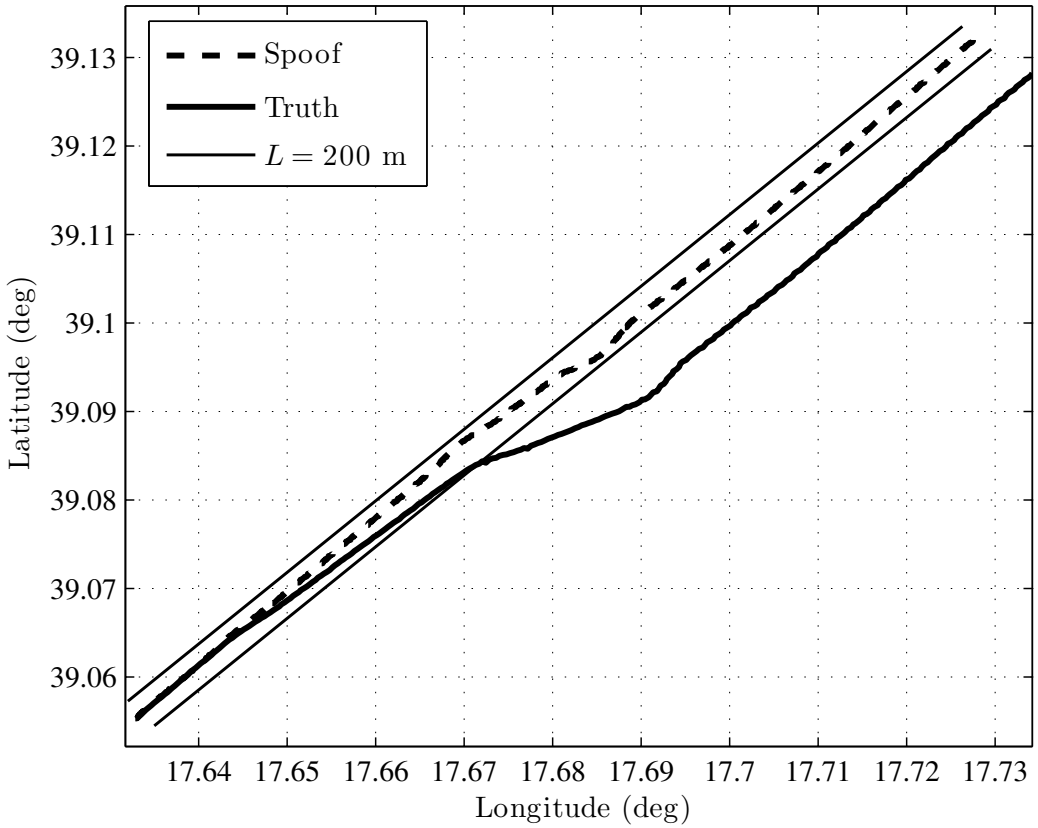

Figure 2.11: Comparison of the ship's reported position and the ship's actual position during a spoofing attack. The thin solid lines indicate ±200 m cross-track deviation.

The spoofer-induced acceleration in the first velocity change was $0.03\,{}^{\mathrm{m}}/_{\mathrm{s}^2}$, while for all other changes the acceleration was $0.1\,{}^{\mathrm{m}}/_{\mathrm{s}^2}$. Note that the maximum spoofer-induced velocity and acceleration exceed the limits assumed in Sec. 2.3 in order to reduce the duration of the experiment. As the captain performed typical correction maneuvers to maintain the spoofed trajectory within a ±200 m corridor, the actual ship's position deviated along a parallel track as shown in Fig. 2.11 and 2.12.

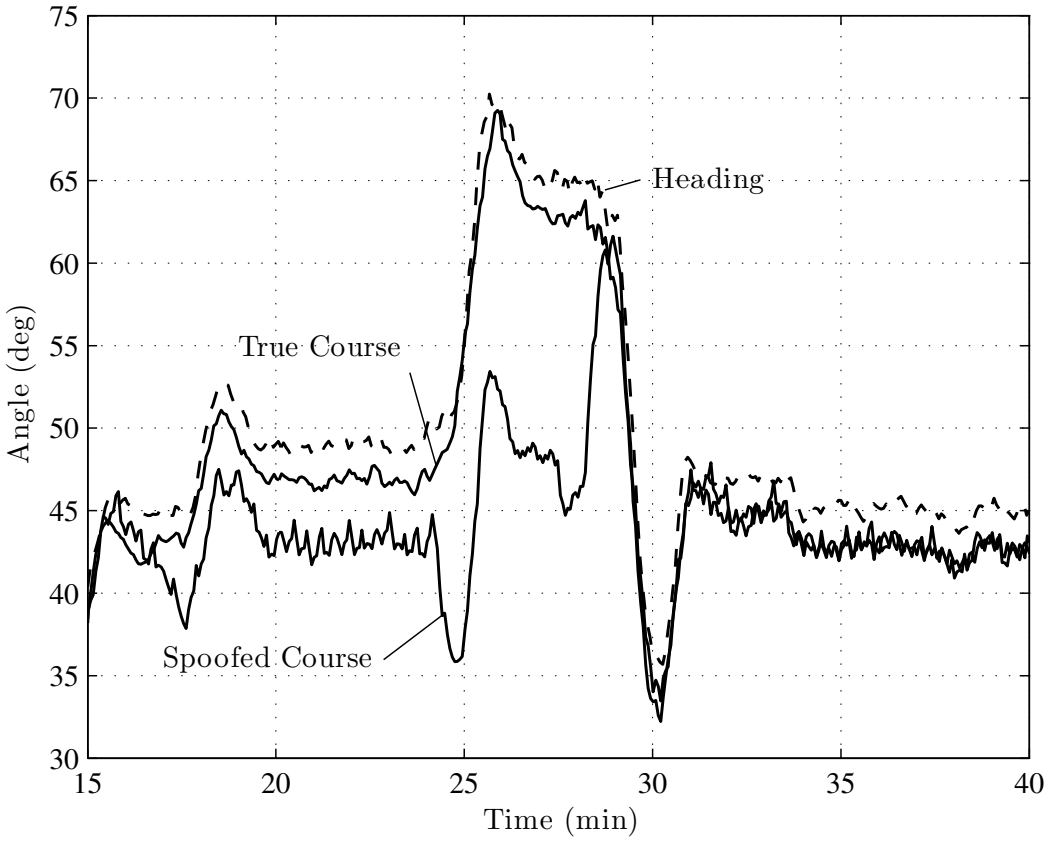

Figure 2.12: Comparison of the ship's heading, spoofed course, and true course during a spoofing attack. Course is defined as the direction of the ship's velocity over ground vector with respect to North.

The ship's reported position and heading were logged to a file during the spoofing attack. Unfortunately, the ship's Doppler log was not functional, but the ship's engine throttle control was set to Full Ahead, so the ship's speed through water U is assumed to be a nominal 15 knots. The logged measurements are fed post-facto into the innovations-based spoofing detection framework developed in Sec. 2.3. In order to determine the optimal sampling time $T_s^\star$ for the experiment, many of the same parameter values indicated in Fig. 2.5 were used, except $0.5\,\mathrm{m/s} \leq v_{\max} \leq 2\,\mathrm{m/s}$ and $L = 200\,\mathrm{m}$. Even though the ship was traveling in open waters, the narrow corridor was chosen to reduce the time scale of the experiment from hours to minutes, and could potentially represent approaches to harbors with many surrounding hazards. The resulting minimax optimization yields $T_s^\star \approx 250\,\mathrm{s}$ and integrity risk $\bar{I}_R^\star = 0.8956$ for the worst-case attacks. The first phase of the actual attack, while $|z_{\mathrm{m}}| \leq 200\,\mathrm{m}$, is a worst-case attack and remains covert with respect to the detection framework. The second phase of the attack is significantly less covert assuming $L = 700\,\mathrm{m}$, $u_{\max} = 0.1\,\mathrm{m/s^2}$, and $v_{\max} = 2\,\mathrm{m/s}$, which yields a theoretical integrity risk of $\bar{I}_R = 0.0067$, although the actual integrity risk is different due to the change in the spoofer-induced velocity in the middle of the attack. The NIS values generated by the detection framework for the experimental data with five different sampling phases are shown in Fig. 2.13. Recall that the integrity risk computed previously is the marginal risk over uniformly distributed sampling phases. A realization for a particular sampling phase leads to HMI if the associated NIS values never cross the detection threshold λ after the spoofing attack begins and before the attack leads to hazardous

conditions. If the NIS value falls into the shaded regions shown in Fig. 2.13, then the defender has succesfully detected an attack (i.e. declared H_1) before hazardous conditions occur.

2.6 Strategies for Mitigating Surface Vessel Vulnerability to GNSS Deception

A number of promising methods are currently being developed to defend against civil GNSS deception attacks. These can be categorized as (1) receiver-autonomous signal-processing-oriented techniques, which require no antenna motion or specialized antenna hardware [82, 83, 84, 85]; (2) receiver-autonomous antenna-oriented techniques, which require antenna motion or specialized antenna hardware [86, 87, 88]; (3) cryptographic techniques that require signal specification modifications to overlay unpredictable but verifiable modulations on existing or future civil GNSS signals [89, 90]; and (4) techniques that exploit the existing encrypted military signals to offer civil GPS signal authentication for networked GPS receivers [91, 92, 38, 93]. Among these methods, the dual-antenna technique described in [94, 87] seems an especially promising option for maritime protection because (1) it could be implemented in the near term, and (2) its chief drawbacks relative to the other techniques—larger size and higher cost—are not so critical for marine vessels as they are for handheld devices and small unmanned aerial vehicles, for example. Nonetheless, it will take years before this or other techniques mature and are implemented on a wide scale. Meanwhile, there are no off-the-shelf defenses against GNSS spoofing.

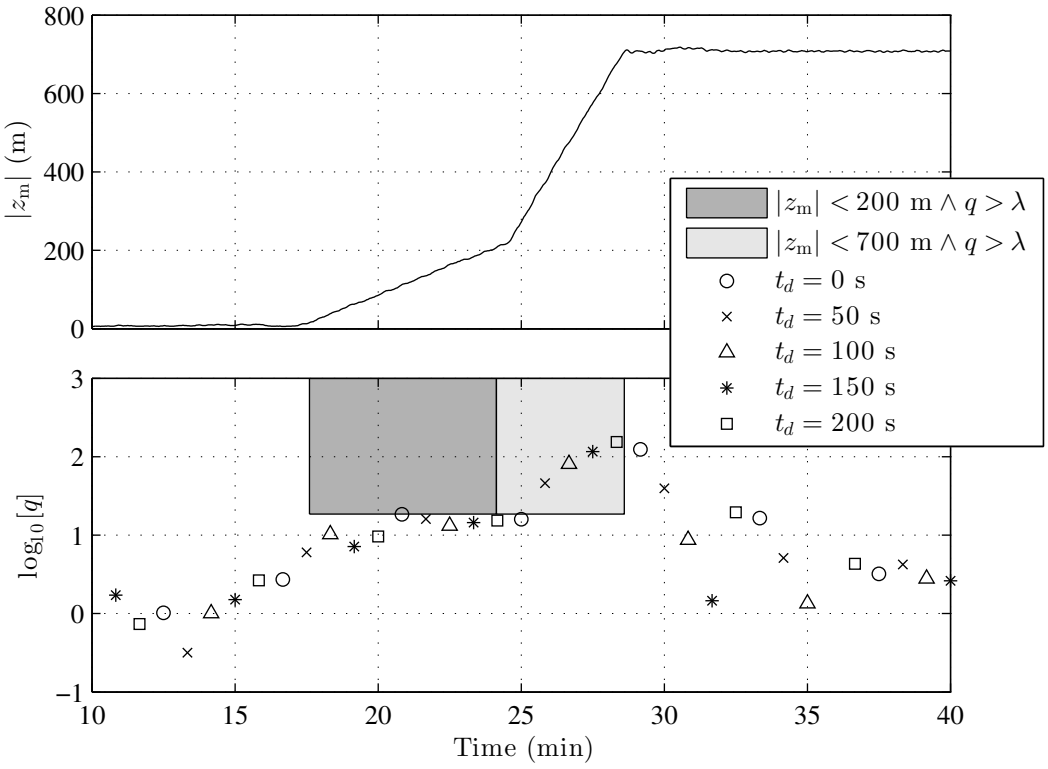

Figure 2.13: NIS values generated by the detection framework with a sampling time $T_s = 250\,\mathrm{s}$ for the experimental data collected on the White Rose of Drachs during a spoofing attack. NIS time history for five different sampling phases are shown. The shaded regions indicate areas where the NIS must fall in order to detect the attack before hazardous conditions occur, preventing an HMI event. The darker and lighter regions correspond to the first and second phase of the attack, respectively. The lower edge of the regions corresponds to the detection threshold λ.

Chapter 3

Emitter Localization

Passive radio-frequency (RF) emitter localization has many military and civilian applications, and in particular, finding GPS jammers as noted in the dissertation's introduction. Many methods for passive RF emitter localization exist that use various types of measurements derived from the unknown emitter signal such as received signal strength (RSS), angle of arrival (AOA), and time and frequency difference of arrival (T/FDOA) [95, 32, 96, 97, 98, 99, 100]. The measurement type has implications on the number of sensors required for observability, emitter state estimability, and sensor complexity. For example, AOA measurements from a single moving platform can provide good estimability, but may require complex sensors such as a rotating directional antenna or an antenna array. On the other hand, T/FDOA measurements obtained by cross-correlation requires at least two spatially-separated platforms with a high-throughput network link, but only a single omnidirectional antenna per platform is required. RSS measurements typically require the least complex sensors and platforms, but tend to contain less information about the emitter state due to unknown nuisance parameters such as transmitted power, path loss exponent, and multipath, leading to poor estimability.

The present work will focus on emitter localization based on cross-

correlation due its low antenna and receiver complexity requirements. In a traditional cross-correlation approach, T/FDOA measurements are obtained from processing the cross-correlated complex ambiguity function (CAF), using the classic algorithms in [101, 102]. Then, the measurements are ingested by a top-level geolocation algorithm such as the algebraic hyperbolic-positioning estimator in [103, 104], the Hough Transform-based estimator in [105], or the T/FDOA nonlinear filter in [96]. However, this two-step estimation approach is suboptimal, especially for low signal-to-noise ratio, because the approach ignores the constraint that all measurements must be consistent with a single emitter position and velocity [34]. In direct geolocation, the CAF is maximized directly by parameterizing the delay time histories in terms of the emitter state space and using known information about the receiver position and clock offset time histories. This chapter will consider direct geolocation methods and their implementation for various experimental scenarios.

This chapter makes three contributions, which distinguishes itself from the work in [34, 35]. First, it considers practical implementation issues of a deployed system such as time synchronization of receivers and long coherent integration for non-constant T/FDOA. In addition, three different types of emitter dynamics models—static, constant velocity, and constant velocity with path constraint—are considered. Second, it goes beyond the simulations in [35] to conduct and report on three field experiments. Third, it provides a comparative study of the performance and complexity of various estimation algorithms based on grid search, the Kalman filter, and the particle filter.

3.1 Received Signal Model

Consider the following model for the signal transmitted by an emitter:

$$s(t) = A_s(t)\cos(2\pi f_c t + \phi_s(t)). \tag{3.1}$$

Here, $A_s(t)$ is the instantaneous amplitude, f_c is the center frequency, and $\phi_s(t)$ is the transmitted beat carrier phase. For convenience, consider the complex envelope $\tilde{s}(t) = A_s(t)\exp(j\phi_s(t))$ and analytic representation $\hat{s}(t) = \tilde{s}(t)\exp(2\pi f_c t)$ of the transmitted signal $s(t)$. Note that analytic signals are a valid approximation when the complex envelope is slowly varying with respect to the center frequency (i.e. bandpass signals) [106]. This "narrowband" approximation is valid in typical scenarios because the bandwidth of the baseband signal $\tilde{s}(t)$ is small with respect to the carrier frequency f_c. Assume that the radio propagation channel induces a non-dispersive delay $\tau_\rho(t)$, an attenuation $A(\bar{\rho})$ that is a function of the average range $\bar{\rho}$ over the time-of-flight interval, and additive white Gaussian noise $n'(t)$. Then the received signal $r'(t)$ at the receiver can be modeled as

$$r'(t) = A(\bar{\rho})\, s(t - \tau_\rho(t)) + n'(t), \tag{3.2}$$

or with an analytic representation as

$$\hat{r}'(t) = A(\bar{\rho})\, \hat{s}(t - \tau_\rho(t)) + \hat{n}'(t),$$

where $\hat{n}'(t)$, the analytic representation of $n'(t)$, is a complex white Gaussian noise process with single-sided power spectral density N_0 in W/Hz. Other propagation effects like multipath and shadowing are not considered in this model.

For electromagnetic waves traveling in a vacuum, the propagation delay $\tau_\rho(t)$ satisfies the implicit relationship

$$c\tau_\rho(t) = \sqrt{(r_e(t-\tau_\rho) - r_s(t))^{\mathrm{T}}(r_e(t-\tau_\rho) - r_s(t))}, \tag{3.3}$$

where c is the speed of light, $r_s(t)$ is the receiver position vector, and $r_e(t)$ is the emitter position vector [107]. For short propagation distances commonly encountered in terrestrial applications, (3.3) can be approximated as

$$c\tau_\rho(t) = \rho(t) = \sqrt{r(t)^{\mathrm{T}} r(t)},$$

where $r(t) = r_e(t) - r_s(t)$ is the relative position vector and $\rho(t)$ is the instantaneous range. The range rate is given by $\dot{\rho}(t) = r(t)^{\mathrm{T}} \dot{r}(t) / \rho(t)$.

Let the relationship between the time t_r at the receiver and true time t be given by

$$t = t_r - \tau_r(t_r), \tag{3.4}$$

where $\tau_r(t_r)$ is the receiver's clock offset from true time. Suppose that a mixing signal with nominal center frequency f_c is generated with the receiver's clock. The mixing signal's phase $\phi_r(t_r)$ is related to t_r by

$$\phi_r(t_r) = 2\pi f_c t_r + \phi_{r,0},$$

where $\phi_{r,0}$ is the initial phase of the oscillator. Let the mixing operation be modeled such that the resulting baseband signal $\tilde{r}'(t)$ is given by

$$\begin{aligned}\tilde{r}'(t) &= \hat{r}'(t)\exp(-j\phi_r(t_r)) \\ &= A(\bar{\rho})\,\tilde{s}(t-\tau_\rho(t))\exp(-j\phi'(t,t_r)) + \tilde{n}'(t),\end{aligned} \tag{3.5}$$

where

$$\phi'(t, t_r) = 2\pi (t_r - t + \tau_\rho(t)) f_c + \phi_{r,0}$$

and $\tilde{n}'(t) = \hat{n}'(t) \exp(-j\phi_r(t_r))$ is a zero-mean baseband complex Gaussian process. The receiver clock model in (3.4) is used in (3.5) to express $\tilde{r}'(t)$ in the receiver's time base, denoted $\tilde{r}(t_r)$. The noise-free baseband received signal $\tilde{s}_r(t_r)$ is given by

$$\tilde{s}_r(t_r) = A(\bar{\rho})\, \tilde{s}(t_r - \tau_m(t_r)) \exp(-j\phi_m(t_r)), \tag{3.6}$$

with the apparent delay $\tau_m(t_r)$ defined as

$$\tau_m(t_r) = \tau_r(t_r) + \tau_\rho(t_r - \tau_r(t_r)) \tag{3.7}$$

and the received beat carrier phase $\phi_m(t_r)$ given by

$$\phi_m(t_r) = 2\pi f_c \tau_m(t_r) + \phi_{r,0}.$$

The full expression for the baseband received signal $\tilde{r}(t_r)$ is given by

$$\tilde{r}(t_r) = \tilde{s}_r(t_r) + \tilde{n}(t_r), \tag{3.8}$$

where $\tilde{n}(t_r) = \tilde{n}'(t_r - \tau_r(t_r))$ is still a zero-mean baseband complex Gaussian process.

Assuming a nominal sampling rate T_s, the digital representation of the signal $\tilde{r}(t_r)$ is given by $\tilde{r}[k] = \tilde{r}(kT_s)$. The noise $\tilde{n}(t_r)$ is generated at each receiver based on the noise power density N_0 in W/Hz over the two-sided noise-equivalent bandwidth B_n in Hz. Therefore, the noise power σ_n^2 in Watts is given by

$$\sigma_n^2 = N_0 B_n.$$

The complex noise time series $\tilde{n}[k]$ is a scaled and filtered version of a sequence of random samples whose real and imaginary components are independent and normally distributed. The noise samples are scaled so that

$$\mathbb{E}\left[\tilde{n}[k]\tilde{n}^{\star}[k]\right] = \sigma_n^2.$$

The emitter has an average transmitted power density P_s in W/Hz over the single-sided noise-equivalent bandwidth. The spreading loss $L\left(\bar{\rho}\right)$ is given by

$$L\left(\bar{\rho}\right) = \frac{\lambda_c^2}{4\pi^2\bar{\rho}^2},$$

where $\lambda_c = c/f_c$ is the nominal wavelength of the signal. Isotropic transmit and receive antennas and no cable loss are assumed. The received signal power σ_s^2 in Watts is given by

$$\sigma_s^2 = L\left(\bar{\rho}\right) P_s B_n.$$

The signal component of the received signal $\tilde{s}_r[k] = \tilde{s}_r\left(kT_s\right)$ is scaled so that

$$\mathbb{E}\left[\tilde{s}_r[k]\tilde{s}_r^{\star}[k]\right] = \sigma_s^2,$$

which constrains $A\left(\bar{\rho}\right)$ in (3.2) appropriately.

3.2 Generalized Cross-Correlation Function (GCCF)

Consider the generalized cross-correlation function (GCCF) for a pair of complex baseband signals $\tilde{z}_1\left(t\right)$ and $\tilde{z}_2\left(t\right)$:

$$S\left(\tilde{z}_1\left(t\right), \tilde{z}_2\left(t\right), \tau_1\left(t\right), \tau_2\left(t\right)\right) \triangleq \int_0^T \tilde{z}_1\left(t+\tau_1\left(t\right)\right) \tilde{z}_2^{\star}\left(t+\tau_2\left(t\right)\right) e^{j2\pi f_c[\tau_1(t)-\tau_2(t)]} dt, \quad (3.9)$$

where $\tau_i(t)$ is the delay time history for received signal $i \in \{1,2\}$ and T is the length of the integration interval. Note that the more familiar complex ambiguity function (CAF) adapted from the radar literature [106],

$$S'\left(\tilde{z}_1(t), \tilde{z}_2(t), \tau_0, f_D\right) \triangleq \int_0^T \tilde{z}_1(t)\, \tilde{z}_2^{\star}(t+\tau_0)\, e^{-j2\pi f_D t} dt, \tag{3.10}$$

where τ_0 is a constant delay and f_D is the Doppler frequency, can be expressed approximately in terms of the GCCF by

$$\tau_1(t) = 0, \quad \tau_2(t) = \bar{\tau}(t) = \tau_0 + \frac{f_D}{f_c} t. \tag{3.11}$$

The impact of assuming constant delay τ_0 in the CAF is negligible if the Doppler frequency is small i.e. $f_D \ll \frac{f_c}{TB_n}$. The GCCF can capture arbitrary geometric and clock motion, not just linear approximations usually only valid over short intervals as in the CAF. However, the discrete-time CAF can be efficiently evaluated for a large number of Doppler frequency with the Fast Fourier Transform (FFT) and can be used to quickly visualize the presence and path of strong emitters in the delay-Doppler domain.

Consider a pair of signals $\tilde{r}_1(t)$ and $\tilde{r}_2(t)$ from receivers 1 and 2, respectively, satisfying the single-emitter propagation model in (3.8). Note that although signals are recorded in different time bases, both signal's functional representation are indexed with the same integration variable t. Then, the GCCF for the received signals is given by

$$\begin{aligned} S\left(\tilde{r}_1(t), \tilde{r}_2(t), \tau_1(t), \tau_2(t)\right) = \alpha_1 \alpha_2^{\star} S\left(\tilde{s}(t), \tilde{s}(t), \tau_1'(t), \tau_2'(t)\right) \\ + N\left(\tilde{r}_1(t), \tilde{r}_2(t), \tau_1(t), \tau_2(t)\right), \end{aligned} \tag{3.12}$$

where the complex attenuation factor α_i is defined as

$$\alpha_i = A\left(\bar{\rho}_i\right) e^{-j\phi_{r_i,0}},$$

$S\left(\tilde{s}\left(t\right), \tilde{s}\left(t\right), \tau_1'\left(t\right), \tau_2'\left(t\right)\right)$ is the generalized auto-correlation function (GACF) of $\tilde{s}\left(t\right)$,

$$\tau_i'\left(t\right) = \tau_i\left(t\right) - \tau_{m,i}\left(t + \tau_i\left(t\right)\right), \tag{3.13}$$

and $N\left(\tilde{r}_1\left(t\right), \tilde{r}_2\left(t\right), \tau_1\left(t\right), \tau_2\left(t\right)\right)$ is a noise function, which includes all correlation terms involving the noise signal $\tilde{n}\left(t\right)$. Note that the GACF is maximized when

$$\forall i \;\; \forall t \;\; \tau_i'\left(t\right) = 0,$$

although if $\tilde{s}\left(t\right)$ has any periodic features, the GACF can reach its maximum value for additional delay time histories. (3.13) is an implicit equation like (3.3) and can be solved iteratively by the recursion

$$\tau_i^{n+1}\left(t\right) = \tau_{m,i}\left(t + \tau_i^n\left(t\right)\right),$$

where $\tau_i^0\left(t\right) = 0$. In a typical scenario,

$$\tau_i\left(t\right) = \lim_{n\to\infty} \tau_i^n\left(t\right)$$

is valid and in practice, the solution converges within reasonable tolerance for small n—usually three iterations. This recursive iteration technique is also used to solve the time-of-flight equation in GPS receivers [108].

The (not necessarily unique) delay time histories that maximize the GCCF in (3.12) are denoted by $\hat{\tau}_1\left(t\right)$ and $\hat{\tau}_2\left(t\right)$, which are typically restricted to some subset of the real function space. For example, if $\hat{\tau}_1\left(t\right) = 0$ and

$\hat{\tau}_2(t)$ is restricted to linear functions as in (3.11), then the GCCF reduces approximately to the CAF, which is given by

$$\begin{aligned} S'(\tilde{r}_1(t), \tilde{r}_2(t), \tau_0, f_D) &\approx S(\tilde{r}_1(t), \tilde{r}_2(t), 0, \bar{\tau}(t)) \\ &= \alpha_1 \alpha_2^{\star} S(\tilde{s}(t), \tilde{s}(t), -\tau_{m,1}(t), \bar{\tau}(t) - \tau_{m,2}(t + \bar{\tau}(t))) \\ &\quad + N(\tilde{r}_1(t), \tilde{r}_2(t), 0, \bar{\tau}(t)). \end{aligned} \tag{3.14}$$

The GACF term of (3.14) is maximized when

$$\forall t \;\; F(t) = \bar{\tau}(t) - \tau_{m,2}(t + \bar{\tau}(t)) + \tau_{m,1}(t) = 0 \tag{3.15}$$

is satisfied. However, given the two-dimensional delay-Doppler parameterization of $\bar{\tau}(t)$, (3.15) can only be approximately satisfied, where the best approximation is the solution to

$$\min_{\tau_0, f_D} \|F(t)\| \geq 0$$

with appropriate choice of function norm. For band-limited Gaussian white noise $\tilde{n}$, the CAF is equivalent to the log-likelihood of $\{\tau_0, f_D\}$, up to an additive constant and scaling [109]. Therefore, the delay and Doppler that maximize the magnitude of the CAF (including noise terms), denoted respectively as $\hat{\tau}_0$ and $\hat{f}_D$, are the corresponding maximum likelihood estimates (MLE) of the time and frequency difference of arrival (T/FDOA) between a pair of receivers for a single emitter. For multiple emitters, the ML approach would be to jointly estimate all unknown emitter T/FDOA while excising known in-band emitter sources such as GPS signals. A sub-optimal algorithm for extracting multiple T/FDOA estimates would be to find all the local maxima of

the CAF above a threshold. However, the auto-ambiguity terms generated by each emitter waveform may interfere with each other if their peaks are close in the delay-Doppler domain, leading to biases in the T/FDOA estimates. These biases can be mitigated by super-resolution methods such as MUSIC, although the number of emitters and reflectors must be known *a priori* for best performance [44, 110].

In order to compute the GCCF digitally, a zero-order hold (ZOH) approximation for $\tau_i(t)$ and $\tilde{z}_i(t)$ is used with hold time T_s. The ZOH continuous-time approximations can be written in terms of the discrete-time representations $\tau_i[k]$ and $\tilde{z}_i[k]$ as

$$\begin{aligned}\tau_i(t) &\approx \tau_i[\text{floor}(t/T_s)] \text{ and} \\ \tilde{z}_i(t) &\approx \frac{1}{\sqrt{T_s}}\tilde{z}_i[\text{floor}(t/T_s)], \text{ respectively.}\end{aligned}$$

The hold time T_s is related to the integration interval by an integer N_s such that $T = N_s T_s$. Then, the GCCF can be approximated as

$$S(\tilde{z}_1(t), \tilde{z}_2(t), \tau_1(t), \tau_2(t)) \approx \sum_{k=0}^{N_s-1} C(k)\, e^{-j2\pi f_c \Delta\tau[k]},$$

where $\Delta\tau[k] = \tau_2[k] - \tau_1[k]$ and the sub-cross-correlation term $C(k)$ is given by

$$C(k) = \int_{kT_s}^{(k+1)T_s} \tilde{z}_1(t + \tau_1[k])\, \tilde{z}_2^{\star}(t + \tau_2[k])\, dt. \tag{3.16}$$

A crude approximation of the integral in (3.16) is to hold the sample at the midpoint of the integral, i.e.

$$C(k) \approx \tilde{z}_1[k_{m,1}]\, \tilde{z}_2^{\star}[k_{m,2}],$$

where

$$k_{m,i} = \text{floor}\left(k + \frac{1}{2} + \frac{\tau_i[k]}{T_s}\right).$$

The GCCF implementation in this work uses the above midpoint approximation, although a more accurate interpolation scheme is considered in [39], which develops a real-time implementation of GPS spoofing detection via cross-correlation of encrypted signals.

3.3 Tightly-Coupled Radio-Frequency Frontend

"Tightly-coupled" refers to an RF receiver architecture in which emitter signals and reference signals are down-converted with the same oscillator and sampled in such a way that a nanosecond-accurate correspondence can be made between the two sampled signal streams (coherent signal conditioning and sampling). Fig. 3.1 shows one straightforward tightly-coupled receiver architecture. Tight coupling between the emitter and reference data enables the data streams from two separate receivers to be synchronized to within nanoseconds and for clock variations over the cross-correlation interval to be estimated and compensated at the carrier-phase level. The tightly-coupled receiver architecture draws from the success of ongoing work in opportunistic navigation at the University of Texas at Austin [111, 112, 113]. Experience with GNSS signals, terrestrial signals of opportunity such as cellular CDMA, and Iridium signals suggests that an emitter localization system could exploit any instance of these three signal types as a reference.

The simplest approach to a tightly-coupled receiver architecture is to

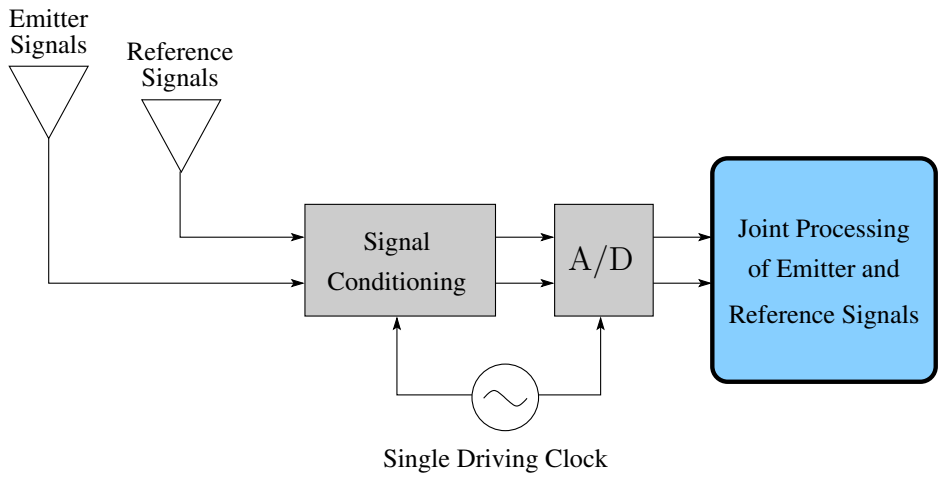

Figure 3.1: Basic tightly-coupled receiver architecture.

use GNSS signals as the reference signals. This approach allows one to exploit the well-known, clean, and stable signal characteristics of GNSS signals. An embeddable real-time software-defined GNSS receiver called GRID has been jointly developed at the University of Texas at Austin and Cornell University, with significant contributions to the code made by the author from 2009 to 2014 [43]. Therefore, software-defined GNSS signal processing can be done within the receiver, which may have limited computational resources, to minimize network throughput requirements. A typical GNSS navigation solution will provide estimates of the receiver position $\hat{r}'_s(t_r)$ and clock offset $\hat{\tau}_r(t_r)$ time histories. This information, coupled with a prediction of the emitter position time history $\bar{r}_e(t) = \bar{r}_e(t_r - \hat{\tau}_r(t_r))$, allows predicting the apparent delay $\bar{\tau}_m(t_r)$, given by

$$\bar{\tau}_m(t_r) = \hat{\tau}_r(t_r) + \frac{1}{c}\sqrt{(\bar{r}_e(t_r - \hat{\tau}_r(t_r)) - \hat{r}'_s(t_r))^{\mathrm{T}}(\bar{r}_e(t_r - \hat{\tau}_r(t_r)) - \hat{r}'_s(t_r))}.$$

3.4 Limits of Coherent Integration

The first source of coherent integration loss is from the ZOH approximation of the GCCF. In the sub-cross-correlation, the first-order delay rate error $\Delta\dot{\tau}$ (i.e. Doppler frequency error) causes power loss according to

$$L_{\text{ZOH}} = \text{sinc}^2\left(\Delta\dot{\tau} f_c T_a\right).$$

In order to avoid the first null of the sinc, the sub-accumulation time T_a should be chosen such that $\Delta\dot{\tau}_{\max} f_c T_a \ll 1$. For highway velocities (35 m/s) and 20 cm wavelengths, $T_a = 1$ ms is reasonable. In addition, errors in the assumed dynamics model of the emitter (which will typically be a low-order polynomial) will limit the overall coherent integration time T.

The second source of loss is from the noisy estimates of the receiver position and clock offset. Consider the coherent sum

$$w = \sum_{n=0}^{N_e-1} \exp\left(j\phi\left[n\right]\right),$$

where each phase noise sample is normally distributed $\phi\left[n\right] \sim \mathcal{N}\left(0, \sigma_\phi^2\right)$ and independent. Then, the coherent integration loss due to white Gaussian phase noise [114] is given by

$$L_{\text{WGN}} = \frac{1 + \exp\left(-\sigma_\phi^2\right)\left(N_e - 1\right)}{N_e}.$$

The result can be approximately extended to colored noise by letting the number of samples be given by

$$N_e = \frac{T}{\tau_d},$$

where τ_d is the decorrelation time of the noise. In this formulation, N_e represents the effective number of white noise samples. Note that although phase biases do not cause coherent integration loss and subsequent loss of estimation precision through lower SNR, the estimation accuracy of the emitter state is still affected.

The error covariance of a GNSS navigation solution

$$x = \begin{bmatrix} r_s \\ b_r \end{bmatrix},$$

where $b_r = c\tau_r$, is typically expressed in the form $P_x = \sigma^2_{\text{URE}} \left(G^{\text{T}} G\right)^{-1}$, where σ_{URE} is the user range error and G is the geometry matrix for the observations used in the solution [108]. Dropping the time arguments, the apparent delay in units of length is given by

$$b_m = f\left(x, r_e\right) = b_r + \sqrt{\left(r_e - r_s\right)^{\text{T}} \left(r_e - r_s\right)}$$

where $b_m = c\tau_m$. A first-order approximation of propagating the error covariance through the nonlinear function $f\left(\cdot\right)$ [81] is given by

$$P_{b_m} = H_x P_x H_x^{\text{T}},$$

where

$$H_x = \frac{\partial f}{\partial x}\left(x, r_e\right) = \begin{bmatrix} \frac{r_s - r_e}{\sqrt{\left(r_e - r_s\right)^{\text{T}} \left(r_e - r_s\right)}} & 1 \end{bmatrix}.$$

In a typical scenario where the emitter and receivers are constrained to the local horizontal plane, it can be shown that the error covariance of the apparent delay can be expressed as

$$P_{b_m} = \left(H^2_{\text{DOP}} + T^2_{\text{DOP}}\right) \sigma^2_{\text{URE}}$$

where the dilution of precision (DOP) terms are related to the elements of $\left(G^{\mathrm{T}}G\right)^{-1}$ as shown in [108].

The user range error depends on the type of observations used in the GNSS navigation solution. For example, code-phase measurements yield $\sigma_{\mathrm{URE,UDP}} \approx 6\,\mathrm{m}$, which includes multipath, atmospheric, and satellite clock and ephemeris errors, whereas single-differenced code-phase measurements yield $\sigma_{\mathrm{URE,SDP}} \approx 1\,\mathrm{m}$ and double-differenced (DD) carrier-phase measurements yield $\sigma_{\mathrm{URE,DDC}} \approx 1\,\mathrm{cm}$ [108]. The differencing operation nearly cancels common error sources from atmosphere propagation modeling and satellite clock and ephemeris parameters. The technique of using DD carrier-phase measurements with successful integer ambiguity resolution to compute a navigation solution is called carrier-phase differential GPS (CDGPS). With CDGPS, the coherent integration loss is bounded by $\exp\left(-\sigma_\phi^2\right)$ for large N_e where $\sigma_\phi \approx 4\pi\sigma_{\mathrm{URE,DDC}}/\lambda_c$, assuming a typical DOP = 4 for CDGPS. For emitter wavelengths of 20 cm, the loss is bounded by $-1.7\,\mathrm{dB}$. With code-phase measurements, the phase noise is much greater than a cycle, and so the coherent integration loss is unbounded. In fact, non-coherent integration yields less loss for $N_e > 1$, so the coherent integration time T is limited by the decorrelation time τ_d. Typically, a code-phase-based navigation solution is aided by carrier-phase and/or IMU measurements, which can significantly increase the decorrelation time. For example, a carrier-aided delay-locked loop can have bandwidths as low as 0.01 Hz (an effective decorrelation time of 100 seconds), and for $T \ll \tau_d$, the loss is nearly equivalent to CDGPS [108]. Similarly, a complementary Kalman filter ingesting both code-phase and IMU measure-

ments will have an effective decorrelation time on the order of 1–100 seconds depending on the accuracy of the IMU [81].

3.5 Single-Emitter Localization Algorithms

3.5.1 Emitter Dynamics Model

The emitter position time history throughout the integration interval T is parameterized by a low-dimensional state space η. The following models are considered: nearly static (NS), nearly constant velocity (NCV), and nearly constant velocity with path constraint (NCVP). The descriptions for each model, such as the state variables, time history within the integration interval, and update recursions, are given in Table 3.1. Note that for NCVP, the distance along the path s can be transformed to Cartesian coordinates by functions $T_x(s)$ and $T_y(s)$. The process noise $v(k)$ represents continuous AWGN velocity (NS) or acceleration (NCV and NCVP) integrated over the interval T, where the covariance matrices Q_n are given by

$$
\begin{aligned}
Q_0 &= q_0 I_{2\times 2} T \\
Q_1 &= q_1 \begin{bmatrix} Q_{cv} & 0 \\ 0 & Q_{cv} \end{bmatrix} \\
Q_2 &= q_2 Q_{cv}
\end{aligned}
$$

and

$$
Q_{cv} = \begin{bmatrix} \frac{1}{3}T^3 & \frac{1}{2}T^2 \\ \frac{1}{2}T^2 & T \end{bmatrix}.
$$

Type	State Space	Time History	Update
NS	$\eta = \begin{bmatrix} x_0 \\ y_0 \end{bmatrix}$	$x(t) = x_0$ $y(t) = y_0$	$\eta(k+1) = F_0\eta(k) + v(k)$ $v(k) \sim \mathcal{N}(0, Q_0)$
NCV	$\eta = \begin{bmatrix} x_0 \\ \dot{x} \\ y_0 \\ \dot{y} \end{bmatrix}$	$x(t) = x_0 + \dot{x}t$ $y(t) = y_0 + \dot{y}t$	$\eta(k+1) = F_1\eta(k) + v(k)$ $v(k) \sim \mathcal{N}(0, Q_1)$
NCVP	$\eta = \begin{bmatrix} s_0 \\ \dot{s} \end{bmatrix}$	$x(t) = T_x(s_0 + \dot{s}t)$ $y(t) = T_y(s_0 + \dot{s}t)$	$\eta(k+1) = F_2\eta(k) + v(k)$ $v(k) \sim \mathcal{N}(0, Q_2)$

Table 3.1: Description of three different types of emitter dynamics models.

For the update equations, the state transition matrices F_n are defined as

$$F_0 = I_{2\times 2}$$
$$F_1 = \begin{bmatrix} F_2 & 0 \\ 0 & F_2 \end{bmatrix}$$
$$F_2 = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}.$$

Lastly, for the rest of this work, the emitter position is constrained to the surface of the Earth and is expressed in a local two-dimensional Cartesian coordinate system.

3.5.2 Likelihood Function

For the following direct geolocation algorithms, given N_r receivers, the likelihood function $L(\eta|z)$ must be defined up to a scale factor, where $z = \{r_i | i = 1, \ldots, N_r\}$ and r_i is a vector of received complex samples from the ith receiver over some integration interval. The simplified signal model for the ith receiver is given by

$$r_i = \alpha_i H_i(\eta)\, s + n_i,$$

where s and $n_i \sim \mathcal{CN}(0, \sigma_n^2 I)$ are the complex baseband emitter signal and noise vectors, respectively, α_i is the complex path attenuation, and $H_i(\eta)$ is a complex matrix that time and phase shifts the signal vector according to the received signal model in (3.8) and the emitter state. Therefore, given the complex-normal distribution of the noise n_i, the likelihood function is given by

$$L''(\eta, s, \alpha|z) \propto \exp\left(-\frac{1}{\sigma_n^2}\sum_{i=1}^{N_r}\|r_i - \alpha_i H_i(\eta)\, s\|^2\right), \tag{3.17}$$

where $\alpha = \{\alpha_i | i = 1, ..., N_r\}$. For the case of unknown deterministic emitter signal, α and s are nuisance parameters. A reasonable approach is to replace them with their maximum likelihood estimates as in [35], i.e.

$$L(\eta|z) = \max_{s,\alpha} L''(\eta, s, \alpha|z).$$

For the path attenuation scalars, it can be easily shown that

$$\hat{\alpha}_i = \frac{s^{\mathrm{H}} H_i^{\mathrm{H}}(\eta)\, r_i}{\|H_i(\eta)\, s\|^2}. \tag{3.18}$$

Substituting (3.18) into (3.17) yields

$$\begin{aligned}
L'(\eta, s|z) &= L''(\eta, s, \hat{\alpha}|z) \\
&\propto \exp\left(-\frac{1}{\sigma_n^2}\sum_{i=1}^{N_r}\left\|r_i - \frac{s^{\mathrm{H}} H_i^{\mathrm{H}}(\eta)\, r_i}{\|H_i(\eta)\, s\|^2} H_i(\eta)\, s\right\|^2\right) \\
&= \exp\left(-\frac{1}{\sigma_n^2}\sum_{i=1}^{N_r} r_i^{\mathrm{H}} r_i - \frac{\left|r_i^{\mathrm{H}} H_i(\eta)\, s\right|^2}{\|H_i(\eta)\, s\|^2}\right) \\
&\propto \exp\left(\frac{1}{\sigma_n^2}\sum_{i=1}^{N_r}\frac{\left|r_i^{\mathrm{H}} H_i(\eta)\, s\right|^2}{\|H_i(\eta)\, s\|^2}\right) \\
&\approx \exp\left(\frac{1}{\sigma_n^2}\frac{s^{\mathrm{H}} D(z, \eta)\, s}{\|s\|^2}\right),
\end{aligned} \tag{3.19}$$

where

$$D(z,\eta) = V(z,\eta)\, V^{\mathrm{H}}(z,\eta), \text{ and}$$
$$V(z,\eta) = \left[H_1^{\mathrm{H}}(\eta)\, r_1, \ldots, H_L^{\mathrm{H}}(\eta)\, r_L\right].$$

Note that $H_i(\eta)$ is nearly square with dimension N_s and $\|H_i(\eta)\, s\| \approx \|s\|$ since the non-zero elements of $H_i(\eta)$ have unity magnitude. Finally, substituting the best estimate of s into (3.19), which is proportional to the eigenvector associated with the maximum eigenvalue of $D(z,\eta)$, an $N_s \times N_s$ matrix, yields

$$\begin{aligned} L(\eta|z) &= L'(n, \hat{s}|z) \\ &\propto \exp\left(\frac{1}{\sigma_n^2}\lambda_{\max}(D(z,\eta))\right). \end{aligned} \tag{3.20}$$

For large N_s, computing $\lambda_{\max}(D(z,\eta))$ can be quite expensive. However, note that $\lambda_{\max}(D(z,\eta)) = \lambda_{\max}\left(\bar{D}(z,\eta)\right)$, where $\bar{D}(z,\eta) = V^{\mathrm{H}}(z,\eta)\, V(z,\eta)$ is an $N_r \times N_r$ matrix. Typically $N_r \ll N_s$, so computational savings can be achieved using $\bar{D}(z,\eta)$. In fact, the elements of $\bar{D}(z,\eta)$ are simply evaluations of the GCCF for a pair of receivers, i.e. for received signals $\tilde{r}_i(t)$ and associated delay time histories $\tau_i(t,\eta)$, then

$$\begin{aligned} \bar{D}_{i,j}(z,\eta) &= S\left(\tilde{r}_i(t), \tilde{r}_j(t), \tau_i(t,\eta), \tau_j(t,\eta)\right) \\ &\approx r_i^{\mathrm{H}} H_i(\eta)\, H_j^{\mathrm{H}}(\eta)\, r_j, \end{aligned}$$

where $S(\cdot)$ is approximated by a zero-order hold representation in a digital implementation. Note that for two receivers, the cost function is nearly equivalent to the magnitude of the GCCF when the path attenuation at each receiver is approximately equal, i.e. $|\alpha_1| \approx |\alpha_2|$, as shown in [115].

In practice, using the likelihood function $L(\eta|z)$ in (3.20) can lead to problems when the received signal power is high and the effective noise floor is dominated by the signal's sidelobes in the ambiguity domain rather than the receiver noise samples. Although the ideal ambiguity function is the so-called "thumbtack" function, actual signals under a finite integration interval will have some power off the peak in the so-called "pedestal" as described in the radar literature [106]. Therefore, a normalized likelihood function $\hat{L}(\eta|z)$ is considered such that

$$\hat{L}(z|\eta) = \exp\left(\gamma \lambda_{\max}\left(\hat{D}(z,\eta)\right)\right),$$

where γ is some scaling factor and

$$\hat{D}_{i,j}(z,\eta) = \frac{D_{i,j}(z,\eta)}{\sqrt{D_{i,i}(z,\eta)\, D_{j,j}(z,\eta)}}.$$

Note that the diagonal elements of $\hat{D}(z,\eta)$ are exactly one, which imposes the condition

$$1 \leq \lambda_{\max}\left(\hat{D}(z,\eta)\right) \leq N_r.$$

Although no theoretical guidance is offered here, experimental results in the sequel show that the range $1 \leq \gamma \leq 3$ yields satisfactory results.

3.5.3 Grid Search

A straightforward algorithm to determine the maximum likelihood (ML) estimate of the emitter state, denoted as GS, is to grid up the state space and search for the grid point that yields the maximum value of the likelihood function for each integration interval as in [34]. The choice of likelihood function,

either $L(\eta|z)$ or $\hat{L}(\eta|z)$, may yield different ML state estimates, although the difference is negligible in most cases. However, the values of σ_n and γ do not affect the ML estimate since the relative values of the likelihood function are not important when searching for the maximum. The grid search algorithm is well suited for the two-dimensional NS model, with appropriate grid spacing and constraints. However, the search space can become unwieldy for the four-dimensional NCV model. In addition, standard grid search does not allow imposing the dynamical constraint between position and velocity over time in both the NCV and NCVP model. A maximum *a posteriori* (MAP) estimate of the emitter state could be derived from the point-mass filter, which keeps track of prior information through weights associated with each grid point. However, for N grid points, the weight update is $O(N^2)$, as opposed to the $O(N)$ update of the particle filter [116]. In addition, the particle filter searches the state space more efficiently than the point-mass filter through a "dynamically-sized" grid. In the following subsections, a Kalman and particle filter implementation for direct geolocation are presented.

3.5.4 Kalman Filter

A Kalman filter allows smoothing the measurement information from each integration interval with the dynamical constraints on the emitter state. In addition, instead of searching the whole state space like in naive grid search, the previous emitter state and covariance can be used constrain the search space to a smaller region for each integration interval in order to reduce computational effort, although with the added possible risk of divergence.

The standard Kalman filter algorithm for direct geolocation, denoted as KF1, is described in the sequel. Let $\bar{\eta}(k)$ and $\hat{\eta}(k)$ be the *a priori* and *a posteriori* emitter state estimate with error covariance $\bar{P}(k)$ and $P(k)$, respectively, at time k. An iteration of the Kalman filter has two steps: prediction and update. For the nth dynamics model, the prediction step is given by

$$\bar{\eta}(k) = F_n \hat{\eta}(k-1)$$
$$\bar{P}(k) = F_n P(k-1) F_n^{\mathrm{T}} + Q_n.$$

The update step is given by

$$\hat{\eta}(k) = (I - K(k))\,\bar{\eta}(k) + K(k)\,y(k)$$
$$P(k) = (I - K(k))\,\bar{P}(k),$$

where $K(k) = \bar{P}(k)\left[\bar{P}(k) + R(k)\right]^{-1}$ is the Kalman gain and $y(k)$ is the emitter state measurement with noise covariance $R(k)$. In order to compute $y(k)$, consider a set of N_p sampling points $\mathcal{P}(k) = \{\eta_i(k) \,|i = 1, \ldots, N_p\}$, which is generated by either grid or Monte-Carlo sampling. A naive grid sampling method generates points with a user-selected resolution and the square-root of the diagonal elements of $\bar{P}(k)$ multiplied by a user-selected scale factor to set the limits of the grid in each state dimension. A more sophisticated grid sampling also considers the correlation between state dimensions as in the LAMBDA method for integer estimation [117]. The Monte-Carlo sampling method chooses each point such that $\eta_i(k) \sim \mathcal{N}\left(\bar{\eta}(k), \bar{P}(k)\right)$. Once all the sampling points are chosen, then the measurement $y(k)$ is given by the most likely point, i.e.

$$y(k) = \arg\max_{\eta \in \mathcal{P}(k)} L(\eta | z(k)),$$

where $z(k)$ represents the all of the samples from all receivers for the kth integration interval. Again, as in grid search, the choice of likelihood function may slightly affect the state estimate.

The covariance $R(k)$ is determined by computing the Cramer-Rao lower bound (CRLB) at the current state estimate $\bar{\eta}(k)$. Linearization about the raw sample measurements is impractical, so the CRLB is computed with respect to range and range-rate difference of arrival (R/RR-DOA) pseudo-measurements, denoted by ρ_{ij} and $\dot{\rho}_{ij}$ for receivers i and j [118]. Using the pseudo-measurements is a reasonable approximation for slow emitter dynamics, or, in other words, R/RR-DOA measurements are a nearly sufficient statistic for the raw sample measurements. Let $\xi(k)$ be the vector of R/RR-DOA measurements from all possible receiver combinations. The model for $\xi(k)$ is given by

$$\xi(k) = h(\eta(k)) + w(k),$$

where h is a non-linear function that maps the emitter state space to the R/RR-DOA measurement space and $w(k)$ is AWGN with covariance $W(k)$. Choosing $W(k)$ is somewhat arbitrary, although for simplicity, $W(k)$ is set to be diagonal with error variances σ_{ρ}^2 and $\sigma_{\dot{\rho}}^2$ for each range and range-rate DOA measurement, respectively [102, 101]. In reality, the measurements are correlated since the same raw samples are used in cross-correlation pairs that share a common receiver, as shown quite thoroughly in [99]. Regardless of how $W(k)$ is chosen, the CRLB is given by

$$R(k) = \left(H^{\mathrm{T}}(k) W^{-1}(k) H(k)\right)^{-1},$$

where

$$H(k) = \left.\frac{\partial h}{\partial \eta}\right|_{\eta=\bar{\eta}(k)}.$$

In the present work, a central difference method is used to approximate the linearization matrix $H(k)$ for simplicity and generalization of the implementation, although the R/RR-DOA linearizations with respect to typical state parameterizations such as NS are well documented in the literature for interested readers [96].

A modified Kalman filter algorithm, denoted as KF2, uses the same prediction step as KF1 and a particle-filter-like update step. After the prediction step, a set of particles are created with Monte-Carlo sampling i.e. $\eta_i(k) \sim \mathcal{N}\left(\bar{\eta}(k), \bar{P}(k)\right)$. Each particle is associated with a weight $w_i(k)$ determined by the likelihood function such that

$$\begin{aligned}
\bar{w}_i(k) &= \hat{L}\left(\eta_i(k) \,|z(k)\right), \text{ and} \\
w_i(k) &= \left(\sum_{i=1}^{N_p} \bar{w}_i(k)\right)^{-1} \bar{w}_i(k).
\end{aligned}$$

Note that normalized likelihood function $\hat{L}(\eta|z)$ is used here exclusively because $L(\eta|z)$ tends to be far too peaky, as discussed previously. Then, the emitter estimate and error covariance are given by the conditional mean and covariance

$$\begin{aligned}
\hat{\eta}(k) &= \sum_{i=1}^{N_p} w_i(k)\,\eta_i(k), \text{ and} \\
P(k) &= \sum_{i=1}^{N_p} w_i(k)\left(\eta_i(k) - \hat{\eta}(k)\right)\left(\eta_i(k) - \hat{\eta}(k)\right)^{\mathrm{T}}.
\end{aligned}$$

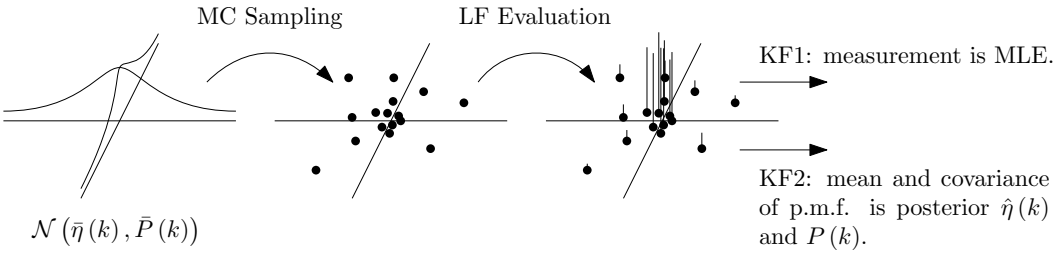

Figure 3.2: Diagram showing the difference between the measurement update of KF1 and KF2 with Monte-Carlo sampling.

Schematically, the difference between the measurement update of KF1 and KF2 with Monte-Carlo sampling is shown in Fig 3.2. The main advantage of the KF2 formulation is that it avoids the computation of $R(k)$ and the errors introduced with the CRLB approximation. However, information is lost when collapsing the particles into a mean and covariance, which is avoided in the full particle filter implementation in the following subsection.

3.5.5 Particle Filter

A particle filter algorithm, denoted as PF, is presented in order to address the two previously mentioned limitations of the grid search algorithm—large search space and lack of dynamical constraints. Note that Sidi also recognized these limitations when using a particle filter to extend Weiss's original direct-geolocation grid search algorithm to the dynamic emitter model [35]. The bootstrap filter is used in the sequel to implement a sequential Monte-Carlo sampling method for Bayesian filtering [119].

Let $\eta(k)$ and $z(k)$ be the emitter state and measurement vector, respectively, at time k. Now, consider a set of N_p sampling points or particles

and associated weights at time k, $\{(\eta_i(k), w_i(k)) | i = 1, \ldots, N_p\}$ and the set of all measurements up to time k be given by $Z(k) = \{z(1), \ldots, z(k)\}$. Then, an approximate posterior probability density function (pdf) is given by

$$p(\eta(k)|Z(k)) = \sum_{i=1}^{N_p} w_i(k)\,\delta(\eta(k) - \eta_i(k)).$$

The bootstrap filter is initialized with particles sampled from an initial proposal pdf $\eta_i(0) \sim p_0(\eta)$ and weights set to $w_i(0) = 1/N_p$ for all i. Typically $p_0(\eta)$ is a uniform distribution over a constrained region of the state space (such as the region used in grid search). An iteration of the bootstrap filter has three steps: prediction, update, and resampling. In the prediction step, process noise samples for each particle are generated from a process noise pdf $v_i(k) \sim N(0, Q_n)$ and the state transition function is used to produce the *a priori* state particles $\bar{\eta}_i(k+1) = F_n\eta_i(k) + v_i(k)$, for the nth dynamics model. Before the update, the index k is incremented. In the update step, the weights for each particle is updated based on the likelihood of the measurement such that

$$\bar{w}_i(k) = \hat{L}(\bar{\eta}_i(k)|z(k))\, w_i(k-1), \text{ and}$$
$$w_i(k) = \left(\sum_{i=1}^{N_p} \bar{w}_i(k)\right)^{-1} \bar{w}_i(k).$$

Note that the normalized likelihood function $\hat{L}(\eta|z)$ is used here exclusively because experimental results have shown $L(\eta|z)$ is far too "peaky" in practice due to the long integration time. Finally, in the resampling step, if the number of effective particles,

$$\hat{N}_{eff} = \left(\sum_{i=1}^{N_p} w_i^2(k)\right)^{-1},$$

is less than some fixed threshold $N_t \leq N$, then the particles are resampled such that

$$\eta_i(k) \sim \sum_{i=1}^{N_p} w_i(k)\,\delta\left(\eta(k) - \bar{\eta}_i(k)\right),$$

and the weights are reset to $w_i(k) = {}^1/_{N_p}$. Otherwise, $\eta_i(k) = \bar{\eta}_i(k)$ and the weights are given by $w_i(k)$ above. Note that in the original implementation of the bootstrap filter, resampling occurs at every iteration so that the threshold is effectively set to $N_t = N$. As Weiss and Sidi have already shown the efficacy of the grid search and particle filter algorithm for direct geolocation in simulation, the following section will apply the presented algorithms to experimental data.

3.6 Experiments

In this section, three field experiments are documented to show the capability of the emitter localization system. Each experiment uses the GPS signals transmitted at 1575.42 MHz (L1) to estimate the receiver position and clock offset. For the WSMR and UTEN scenario, the emitter was a commercial off-the-shelf GPS jammer that transmitted a 30 MHz chirp waveform at GPS L1. For the UAV scenario, the emitter was an Ettus E100 universal software radio peripheral (USRP) that was programmed to transmit a GPS chipping sequence in the 900 MHz amateur radio band. The scenarios for each experiment are summarized in Table 3.2.

The post-processing emitter localization workflow is described in Fig. 3.3. For each experiment, the raw sample data from each dual-antenna tightly-

coupled receiver is recorded to file. First, the reference signal from each antenna is processed by a software-defined GPS receiver. Then, a reference antenna is chosen, with its absolute coordinates fixed either by pseudorange observations or a geodetically-referenced marker, and carrier-phase differential processing for the other antennas allows estimating to within centimeters the baseline with respect to the reference antenna. For each receiver, the local clock offset time history is estimated using carrier-smoothed pseudorange observations from a single satellite and the position of the antenna and the chosen satellite. In order to significantly reduce the computational effort required during cross-correlation, the time delay due to clock effects is accounted for separately in a pre-processing stage, which generates cross-correlated "subaccumulations" for each receiver pair. The pre-processing stage requires choosing two parameters, a subaccumulation integration interval T_{sub} and maximum geometric time offset $\bar{\tau}_{\text{sub}}$, which must satisfy

$$T_{\text{sub}} < \frac{\lambda_c}{4v_{\max}} \text{ and } \bar{\tau}_{\text{sub}} > \frac{b_{\max}}{c},$$

where $v_{\max}$ is the maximum emitter velocity and $b_{\max}$ is the largest baseline between a receiver pair. Since the geometric component of the time delay is unknown at the pre-processing stage, multiple subaccumulations are generated per receiver pair, each testing a hypothetical geometric TDOA at the sampling interval resolution T_s, up to $\pm\bar{\tau}_{\text{sub}}$. Therefore, the number of subaccumulations generated per receiver pair is given by

$$N_{\text{sub}} = 2\text{ceil}\left(f_s\bar{\tau}_{\text{sub}}\right) + 1.$$

Name	Emitter Band	Emitter Dynamics	Receiver Configuration
WSMR	GPS L1	NCV and NCVP	4 stationary
UTEN	GPS L1	NCVP	2 stationary
UAV	900 MHz	NS	1 stationary, 1 airborne

Table 3.2: Summary of experiment scenarios for emitter localization.

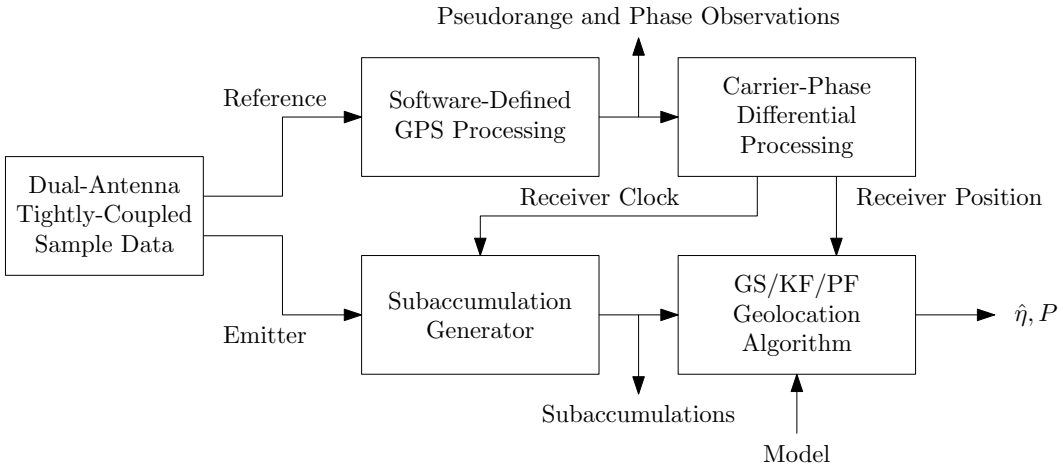

Figure 3.3: Emitter localization post-processing workflow.

Note that if T_{sub} is too small, then there are no computational savings from using the pre-processing stage since the cross-correlation at mostly unnecessary sample offsets must be evaluated. A more sophisticated and computationally-efficient scheme would use a caching mechanism for the subaccumulations within the geolocation algorithm instead of the separate pre-processing stage to avoid evaluating (or re-evaluating) unnecessary subaccumulations. Finally, the geolocation algorithm coherently combines the subaccumulations up to the integration interval T and produces estimates of the emitter state and error covariance.

3.6.1 WSMR Experiment

The U.S. Department of Homeland Security invited the Radionavigation Laboratory and Cornell University to participate in a GPS jamming test exercise at White Sands Missile Range (WSMR) in New Mexico during the summer of 2012. Of the numerous jamming scenarios, several involved driving a commercial 2.5 W off-the-shelf chirp jammer along WSMR's main highway to emulate real-world jamming incidents along major U.S. highways, the most notable being the Newark jammer [24, 25]. Four stationary receivers tuned to GPS L1 were set up along Route 7 as shown in Fig. 3.4, with the vehicle carrying the jammer traveling southbound at about $19\,\mathrm{m/s}$. In this experiment, each receiver was an Ettus USRP N200 driven by a free-running oven-controlled crystal oscillator. The receivers recorded 16-bit complex-baseband samples with sampling interval $T_s = 110\,\mathrm{ns}$ or, equivalently, with sampling rate $f_s \approx 9.09\,\mathrm{MS/s}$. The emitter and reference signal are derived from the same upward-facing hemispherical-gain GPS antenna. As in the UAV experiment, the received GPS signals were used to estimate the antenna positions and the local clock offset time history using both pseudorange and carrier phase observables. The centimeter-accurate antenna baselines are listed in Table 3.3. The subaccumulation integration interval, maximum geometric time offset,

Receiver Pair	Baseline Length (m)
1-2	603.27 m
1-3	653.62 m
1-4	543.86 m
2-3	1205.56 m
2-4	808.93 m
3-4	584.66 m

Table 3.3: Baseline lengths of receiver pairs for the WSMR experiment.

and number of cross-correlation offsets were set to

$$T_{\text{sub}} = \frac{19\,\text{cm}}{4 \times 25\,{}^{\text{m}}\!/_{\text{s}}} \approx 2\,\text{ms},$$
$$\bar{\tau}_{\text{sub}} = \frac{1200\,\text{m}}{3 \times 10^{8}\,{}^{\text{m}}\!/_{\text{s}}} \approx 4\,\mu\text{s},$$
$$N_{\text{sub}} = 2 \times \text{ceil}\left(\frac{4\,\mu\text{s}}{110\,\text{ns}}\right) + 1 = 75.$$

The unscaled raw received power from each receiver as the jammer passed through the network during a 200-second interval is shown in Fig. 3.5. Given the raw subaccumulations produced by the pre-processing stage, the cross-correlated complex ambiguity function (CAF) for each receiver pair can be generated, as shown in Fig. 3.6. The truck was carrying a high-grade inertial navigation system, whose position and velocity estimates are used as truth in order to evaluate the performance of the various algorithms, as detailed in the rest of this section.

In the following set of results, only the cross-correlation pairs between receivers 1, 3, and 4 are considered. The subset allows analyzing geolocation performance with both good and poor emitter-receiver geometry as the emitter moves through the network. Each geolocation algorithm is applied to this

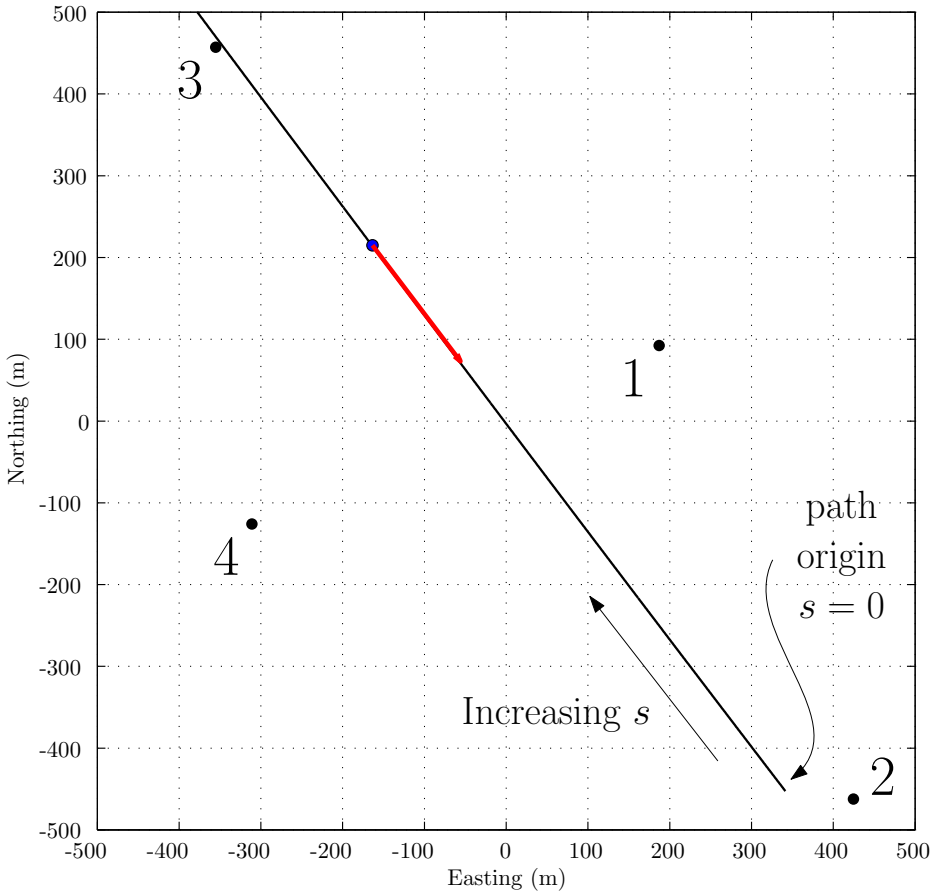

Figure 3.4: Receiver network layout, denoted by numbered black dots, for the WSMR experiment. Route 7 is indicated by the black path, and the "truth" truck position and velocity at a particular instant in time is denoted by the blue dot and red arrow, respectively. Note that the truck's speed is 19.8 $^{\mathrm{m}}/_{\mathrm{s}}$.

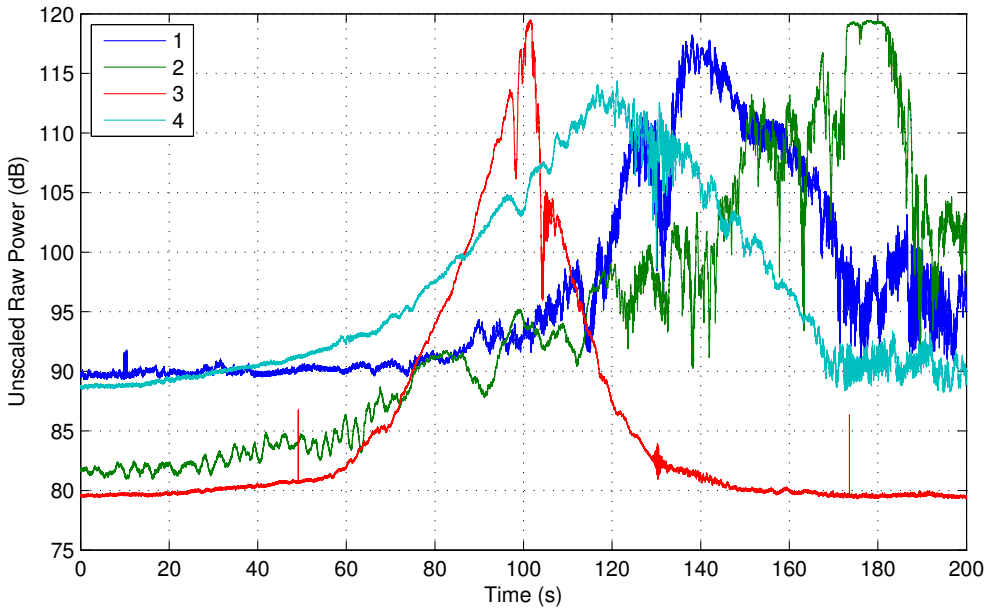

Figure 3.5: Unscaled raw power in decibels measured at each receiver of each 2 ms subaccumulation for the WSMR experiment. Note that the received power at each antenna is loosely correlated with the distance to the jammer. Differences in cable loss and amplifier gain settings yield different raw receiver noise power values, which must be accounted for in the standard likelihood function $L(\eta|z)$. Note that the noise power at receiver 2 and 3 is about 10 dB less than the other two receivers. For the normalized likelihood function $\hat{L}(\eta|z)$, the raw power with all its variations pictured above is used to normalize the cross-correlation of each pair.

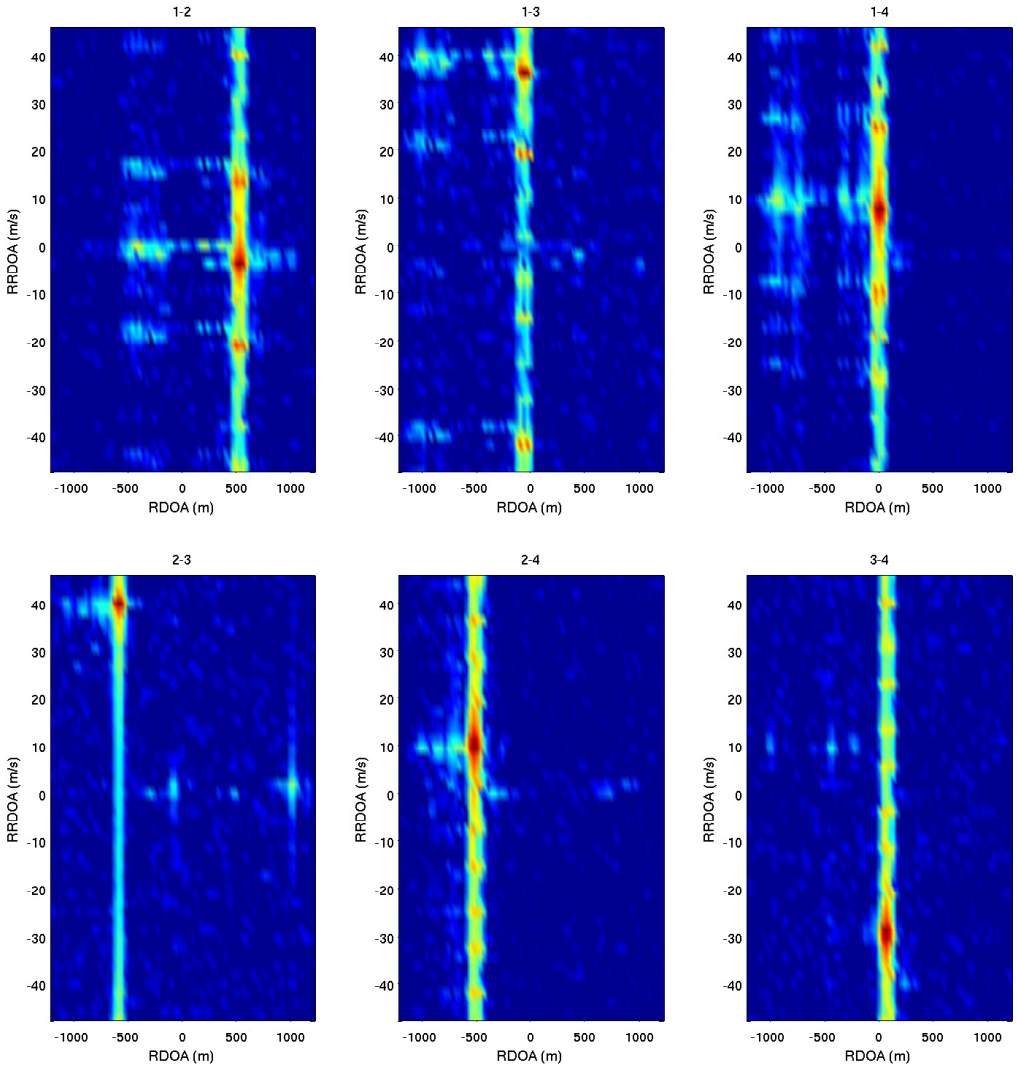

Figure 3.6: Cross-correlated complex ambiguity function, in decibels, for a subset of receiver pairs in the WSMR experiment. Red and blue indicate the strongest and weakest cross-correlation magnitude, respectively. The color scale for each pair is set so that maximum point is red and anything below the mean of the grid is blue. The subaccumulation interval is 2 ms and the Fourier transform interval is 100 ms, so that the image above is generated with 75 RDOA offsets and 50 subaccumulation samples per offset. The instant in time pictured above corresponds to the time of the snapshot in Fig. 3.4. Note that the non-ideal chirp signal structure and multipath results in strong peaks away from the true R/RR-DOA.

subset of the WSMR dataset using the NCVP model with $q = 10\,\mathrm{m}^2/\mathrm{s}^3$ and $T = 100\,\mathrm{ms}$. Note that in order for the cross-correlation to remain coherent, the one-sigma position uncertainty should remain well within one wavelength of the transmitted emitter signal, assuming that the NCVP model actually reflects reality. With the chosen parameters, this condition is met, i.e.

$$\sqrt{q\frac{T^3}{3}} \ll \lambda_c \Rightarrow 5.8\,\mathrm{cm} < 19\,\mathrm{cm}.$$

For the GS algorithm, the state space was constrained to $s \in \begin{bmatrix} 0, & 2000 \end{bmatrix}$ m for the path position with 300 grid points and $\dot{s} \in \begin{bmatrix} -22, & -13 \end{bmatrix}$ m/s for the path velocity with 30 grid points. The GS estimation error using both the standard and normalized likelihood function, along with the predicted standard deviation of the error, is shown in Fig. 3.7. The predicted standard deviation is based on the CRLB approximation with $\sigma_\rho = 100\,\mathrm{m}$ and $\sigma_{\dot{\rho}} = 1\,\mathrm{m/s}$, which were chosen to conservatively fit the experimental error instead of being based on a theoretical thermal noise derivation. In fact, the thermal noise errors are insignificant due to the high emitter power in this experiment, and other sources such as clock synchronization and grid sampling errors dominate. Note that the normalized likelihood function performs slightly better for naive grid search. The KF1 and KF2 algorithms were initialized by using the truth state η_1 and simulated Gaussian noise, i.e. $\bar{\eta}_1 = \mathcal{N}\left(\eta_1, \bar{P}_1\right)$, where

$$\bar{P}_1 = \begin{bmatrix} 20^2 & 0 \\ 0 & 2^2 \end{bmatrix},$$

although in practice, an initial grid search would be used to "acquire" the target. In addition, KF1 and KF2 both used Monte-Carlo sampling of $\hat{L}\left(\eta|z\right)$

for the measurement update with 100 points. The PF algorithm is initialized by uniformly sampling the GS-constrained state space with 10,000 particles (note that the grid resolution was 300×30). Then, after the first resampling step, the number of particles is reduced to 100. The computational complexity for KF1, KF2, and PF are about the same since they all use the same number of sampling points for the likelihood function per integration interval i.e. $N_p = 100$, although, the PF resampling stage can dominate the total run time for short integration intervals. The threshold for the resampling stage is 70% of the total number of particles. For the KF2 and PF algorithms, the scaling factor for $\hat{L}(\eta|z)$ is set to $\gamma = 2$ and $\gamma = 1$, respectively, which was observed to provide the best overall performance. It was observed that if γ is too large, then the particle filter tends to diverge more often, despite having smaller "formal" errors.

Since KF1, KF2, and PF require Monte-Carlo sampling and random initialization, the estimation error was averaged over ten trials for each algorithm, as shown in Fig. 3.8. The estimates of each algorithm for a single trial are shown in Fig. 3.9. Note that KF1 seems to outperform the other algorithms with its lower true estimation error, although KF1 predicts the estimation error to be about twice as large. The reason for KF1's better performance is due to better knowledge of the measurement quality through the CRLB approximation. Therefore, towards the end, when the emitter-receiver geometry becomes quite poor, KF1 is able to depend less on new measurements and use the dynamics model to propagate the information from prior measurements. KF2 and PF must sample the likelihood function with enough points so that

any significant spreading, indicating poor measurement quality, is adequately captured. Also, note that KF2 and PF perform quite similarly as expected due to the update step being nearly identical, with KF2 having slightly better true position estimation error, likely due to KF2 having a deterministic prediction step. Finally, note that the consistency of the estimators (i.e. the difference between predicted and true errors) are well within an order of magnitude.

The effect of the number of particles on the performance of the PF algorithm is shown in Fig. 3.10. Note that the true error decreases with increasing number of particles as expected, especially in position and in the region with poor emitter-receiver geometry. However, the predicted position error increases with more particles, although the predicted velocity error does not seem to be sensitive to the number of particles. Note that the PF algorithm with $N_p = 1,000$ performs similarly to the KF1 algorithm with $N_p = 100$.

Using KF1 and KF2 in practice can lead to problems since accurate initialization requires good emitter-receiver geometry and high received signal strength. Otherwise, false targets may be acquired, and the Kalman filters will not converge to the true emitter state. The effect of emitter-receiver geometry on estimability for the NCVP model is shown in Fig. 3.11. In addition, the GS and PF sample points are plotted together to visualize how the algorithms cope with the changing geometry as the truck moves through the receiver network, shown in Fig. 3.12. The start time was chosen such that truck had already passed receiver 3 and was well within the receiver network at about $s \approx 1000\,\text{m}$. A standard particle filter will be more robust to rejecting false targets than KF1 and KF2, although parallel Kalman filters could be used in

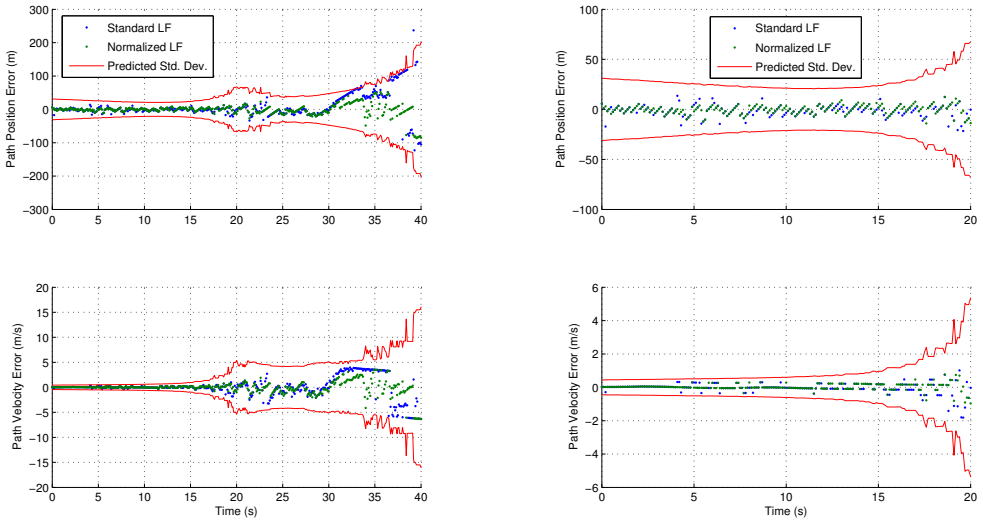

(a) Estimation performance over 40-second interval of the WSMR dataset.

(b) Zoomed-in region before emitter-receiver geometry leads to poor estimability.

Figure 3.7: GS algorithm performance with NCVP model for the WSMR experiment with both the standard and normalized likelihood function. Predicted standard deviation based on CRLB approximation is tuned to match true errors.

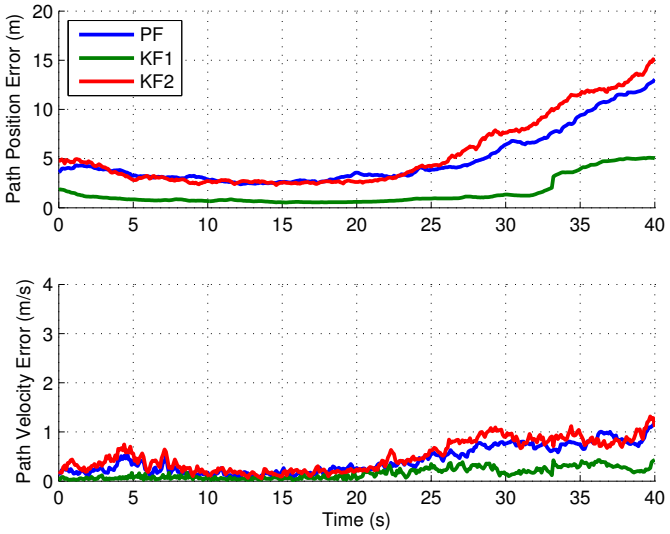

(a) True one-sigma estimation error averaged over ten Monte-Carlo trials.

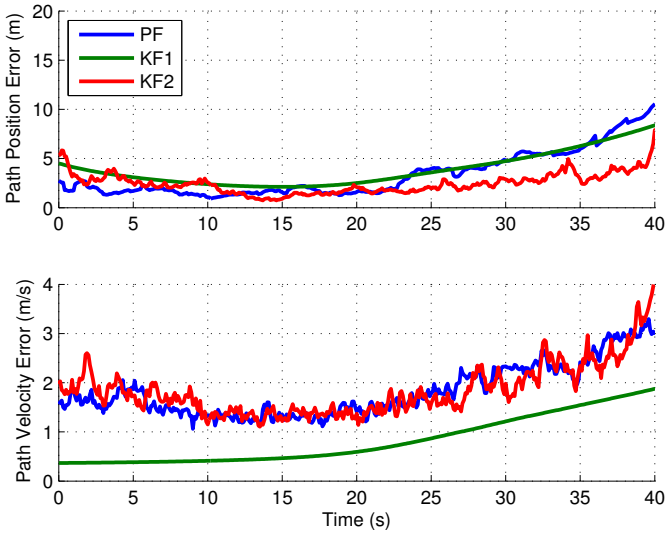

(b) Predicted one-sigma estimation error for a single trial.

Figure 3.8: KF1, KF2, and PF algorithm performance with NCVP model for the WSMR experiment.

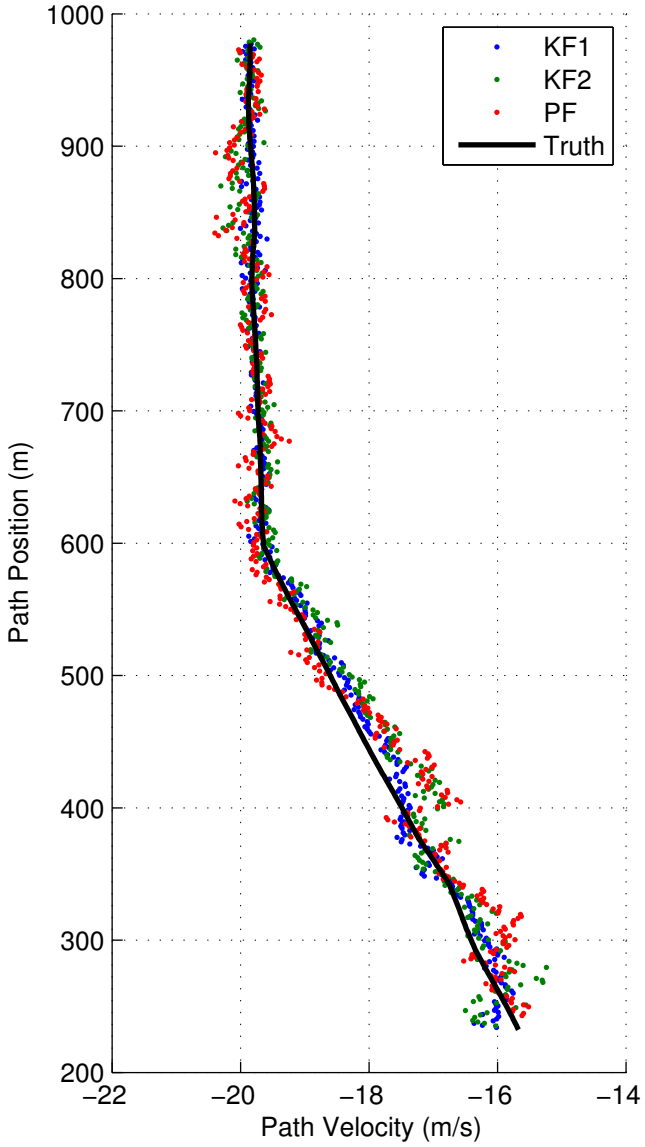

Figure 3.9: Estimated path position and velocity for KF1, KF2, and PF algorithms with NCVP model from a single trial over a 40-second interval of the WSMR dataset.

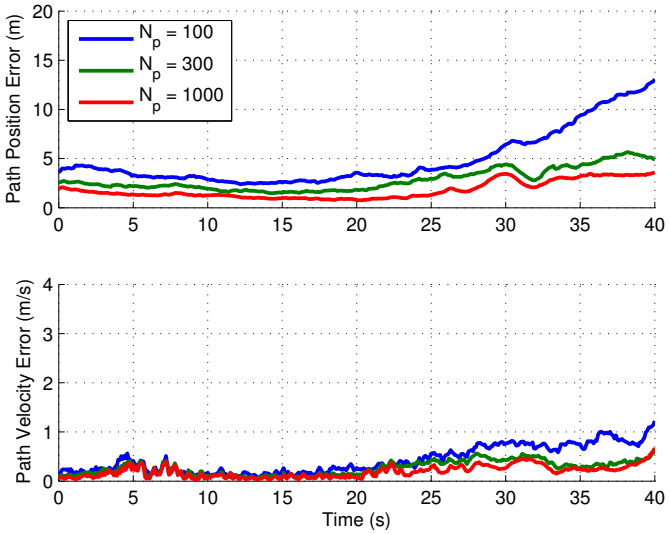

(a) True one-sigma estimation error averaged over ten Monte-Carlo trials.

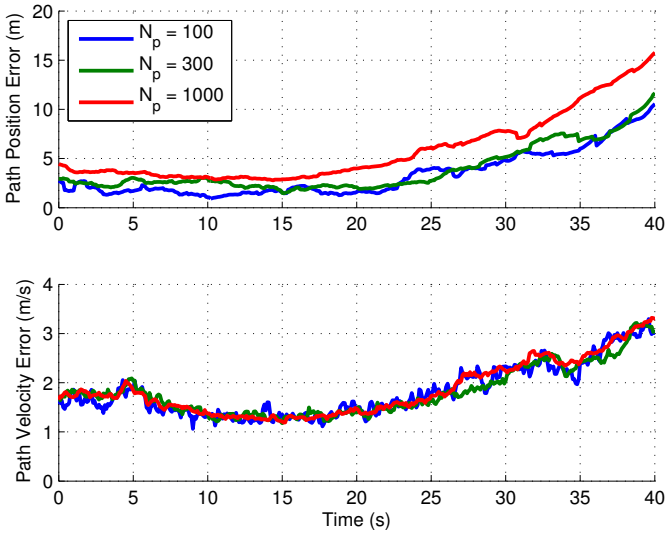

(b) Predicted one-sigma estimation error for a single trial.

Figure 3.10: PF algorithm performance with NVCP model and varying number of particles for the WSMR experiment.

a multiple hypothesis framework at the expense of computational complexity.

The PF algorithm is applied to the WSMR dataset using the NCV model with $q \in \{10,\, 30\}$ $\mathrm{m^2/s^3}$ and $T = 100\,\mathrm{ms}$. The algorithm is initialized by uniformly sampling the state space with 10,000 particles constrained to $x \in \begin{bmatrix}-500, & 0\end{bmatrix}$ m, $y \in \begin{bmatrix}0, & 500\end{bmatrix}$ m, $\dot{x} \in \begin{bmatrix}5, & 20\end{bmatrix}$ m/s, and $\dot{y} \in \begin{bmatrix}-20, & 5\end{bmatrix}$ m/s. Then, after the first resampling step, the number of particles is reduced to 1,000. The normalized likelihood function was used with $\gamma = 2$. The position and velocity estimation error of the NCV model with $q = 10$ and $q = 30$ compared to the NCVP model with $q = 10$ and $N_p = 100$ is shown in Fig. 3.13. The velocity estimation error is about the same for all models, however, the position estimation error is about twice as large for the NCV model with $q = 10$ than the other two. Again, the algorithm's estimates are consistent to within an order of magnitude. In the case of three receivers, the path constraint was not expected to change the estimation performance, although the NCVP model can achieve the same performance as the NCV model with 10 times fewer particles. In addition, only two receivers are required for observability with the NCVP model, although estimability will suffer when compared to three receivers. A final way to decrease computational effort is to lengthen the integration interval T so that fewer particle filter updates occur, either by coherent or non-coherent integration.

Finally, additional results from the WSMR experiment using data from two and four receivers are presented. First, the estimability for two and four receivers are shown in Figs. 3.14 and 3.15, respectively. Note in Fig. 3.4 that receiver pair $(2, 3)$ has the longest baseline, and most of the constant-

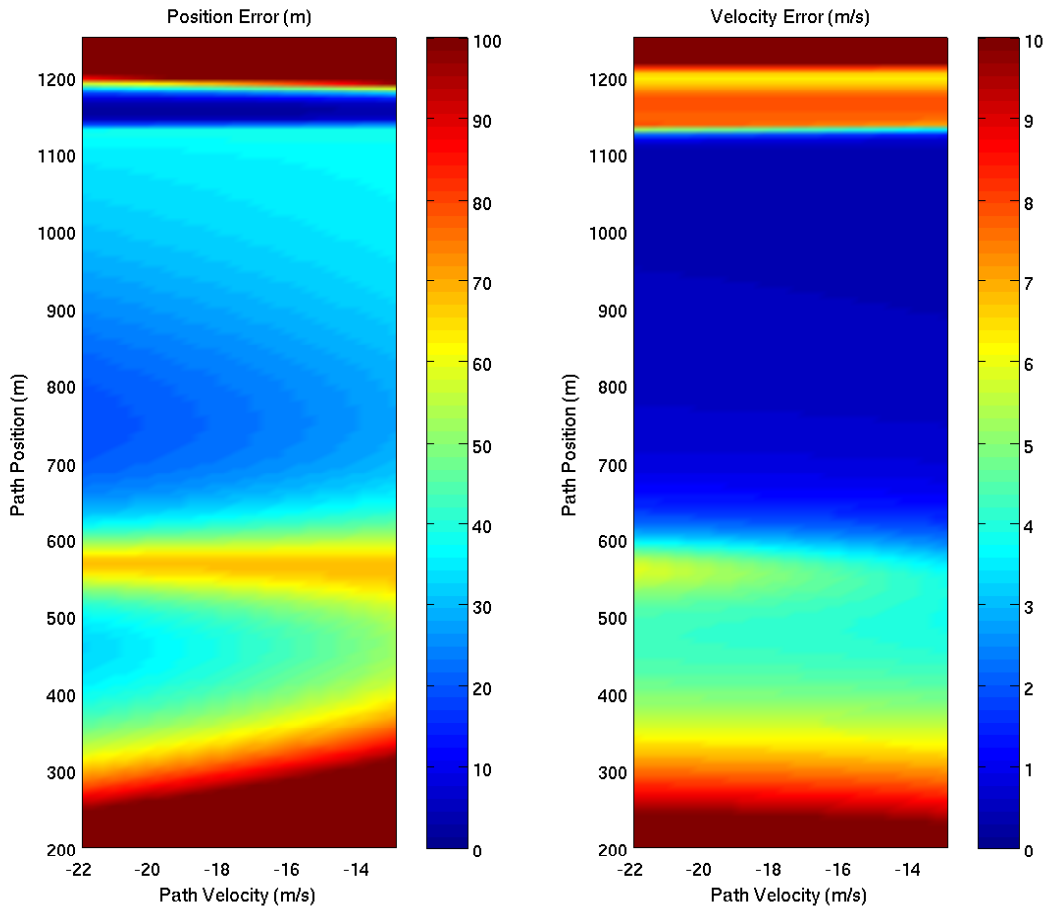

Figure 3.11: Predicted path position and velocity error using CRLB approximation with $\sigma_\rho = 100\,\mathrm{m}$ and $\sigma_{\dot{\rho}} = 1\,\mathrm{m/s}$ for NCVP model and using only receivers 1, 3, and 4 as shown in Fig. 3.4. Note that for $s < 250\,\mathrm{m}$ or $s > 1200\,\mathrm{m}$, the emitter-receiver geometry yields poor estimability, with error exceeding the maximum value of the color scale.

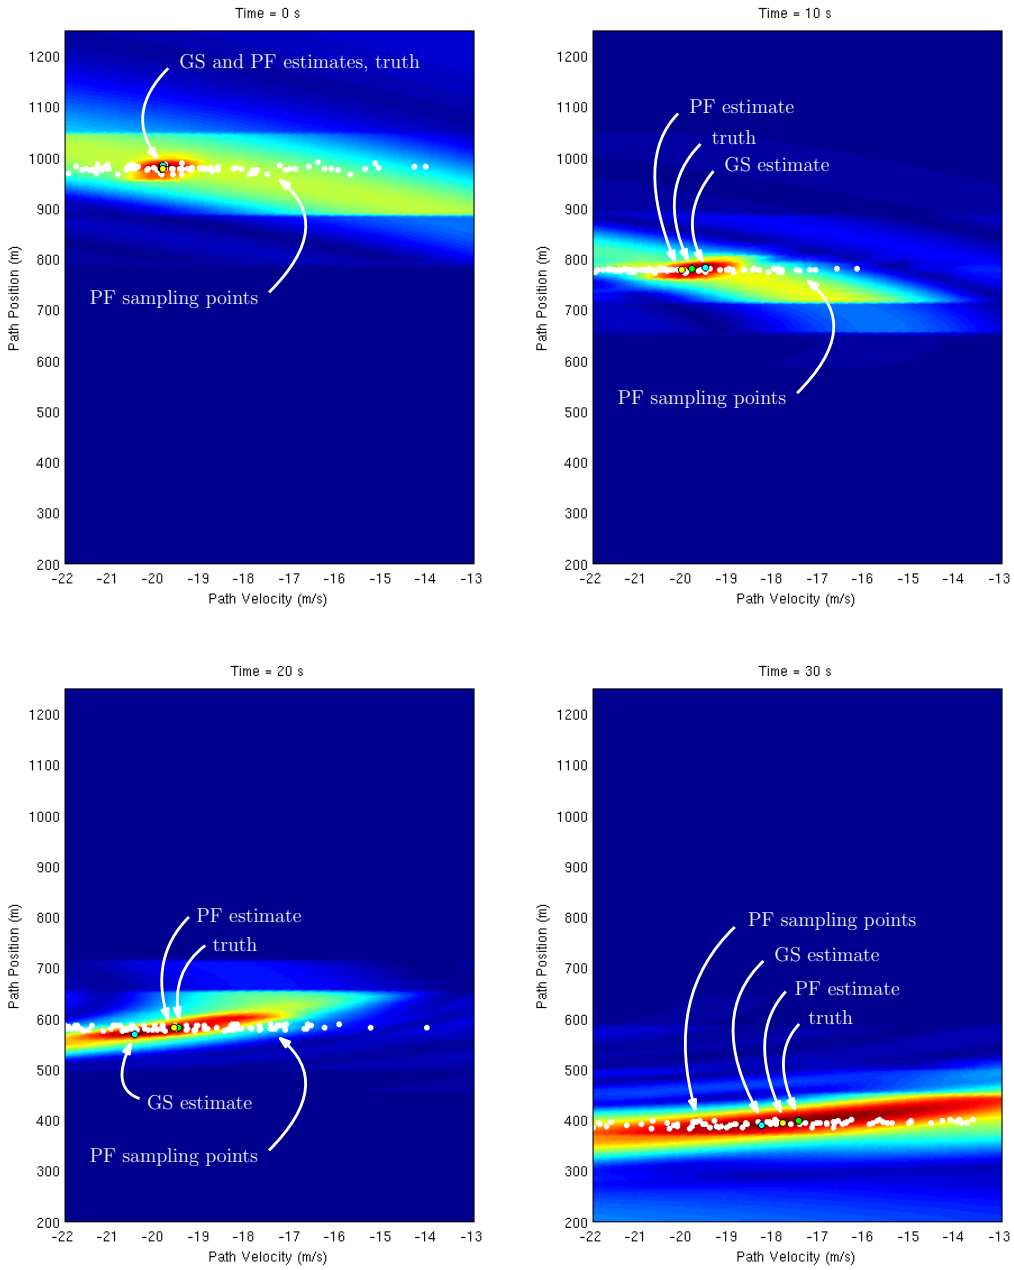

Figure 3.12: GS and PF NCVP state estimate and sample points for the WSMR experiment at four instants in time. A portion of the 300×30 GS points form the background color plot with blue and red for lower and higher values of the normalized likelihood function, respectively. The 100 PF points are indicated by small white dots.

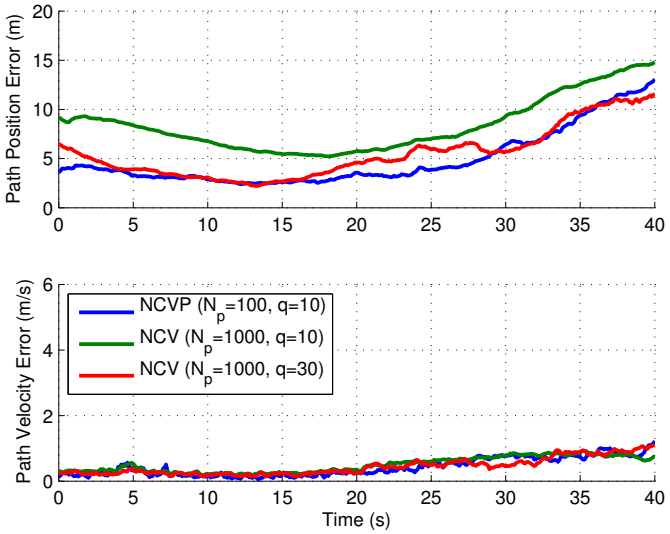

(a) True one-sigma estimation error averaged over ten Monte-Carlo trials.

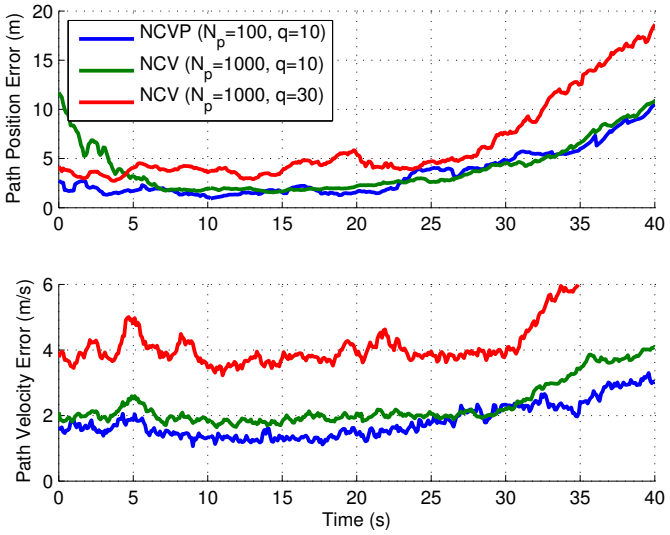

(b) Predicted one-sigma estimation error for a single trial.

Figure 3.13: PF algorithm performance comparing NCV and NCVP models for the WSMR experiment.

TDOA hyperbolas generated by the pair cross the road at approximately right angles. As expected, for $s < 1100\,\mathrm{m}$, the predicted position and velocity error is relatively uniform, even for just the single receiver pair. The number of receivers required for emitter localization in a given search space can be reduced, which additionally reduces computational and network complexity, with strategic placement of receivers. The GS algorithm is applied to both cases as shown in Figs. 3.16 and 3.17, which provides an upper bound on the estimation error for all the algorithms as GS estimates are not smoothed by the dynamics model.

3.6.2 UTEN Experiment

The UT emitter localization network (UTEN) consists of two fixed receivers denoted as CSR and ARL in Fig. 3.18. The network was first introduced in [44] and demonstrated localization of static emitters in an amateur radio band using three static receivers (an addition mobile receiver denoted as MBL was used). With three static receivers, it is possible to estimate the two-dimensional position and velocity of an emitter (i.e. the NCV model is observable). However, in order to locate jammers in the GPS band, a continuous monitoring network is required due to the infrequent nature of GPS jamming, so the MBL receiver could not be used. In order to localize an emitter with only two receivers, the emitter is constrained to move along the two major highways in the area.

From December 2011 to January 2012, the CSR and ARL receivers were tasked to continuously sample the GPS spectrum at $10\,{}^{\mathrm{MS}}/_{\mathrm{s}}$ through both an

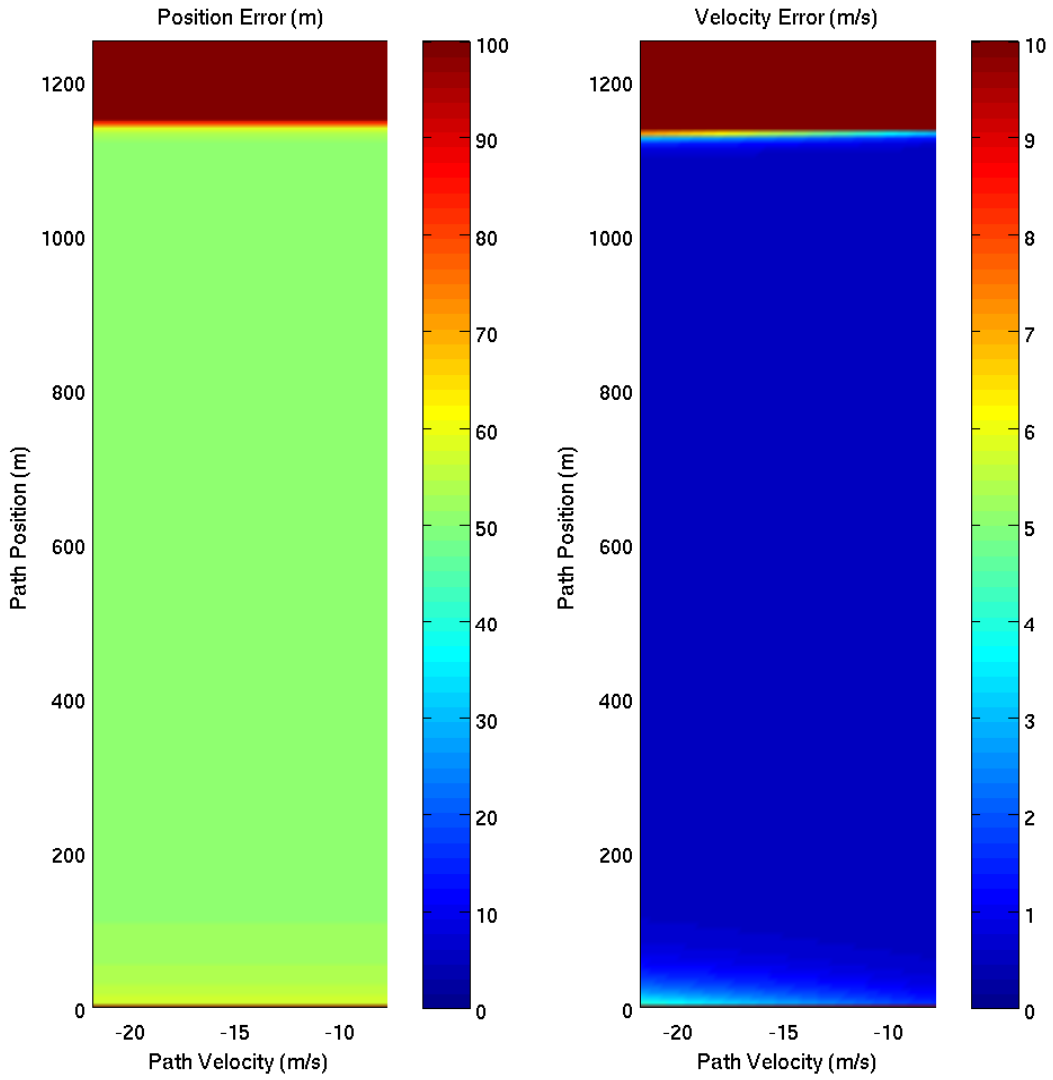

Figure 3.14: Predicted path position and velocity error using CRLB approximation and using only receivers 2 and 3 as shown in Fig. 3.4.

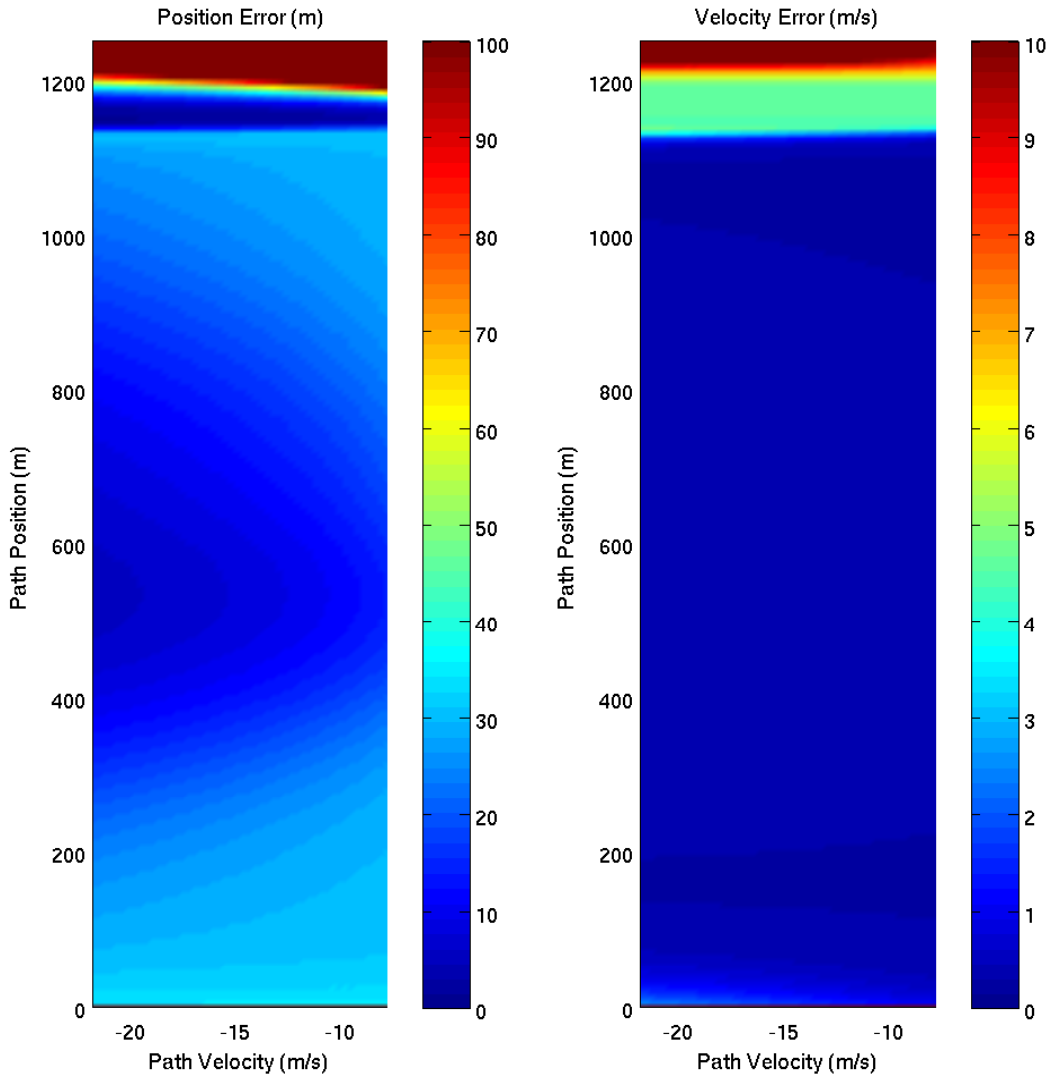

Figure 3.15: Predicted path position and velocity error using CRLB approximation and using all four receivers as shown in Fig. 3.4.

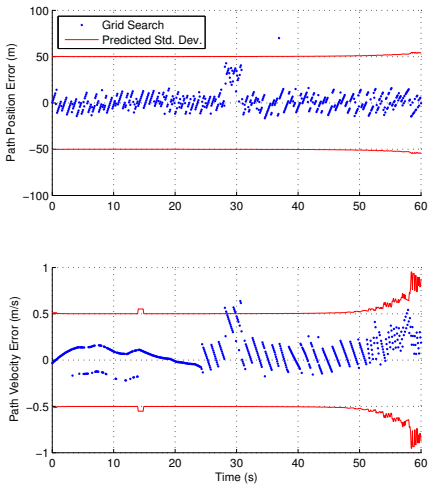

(a) Estimation performance with receivers 2 and 3.

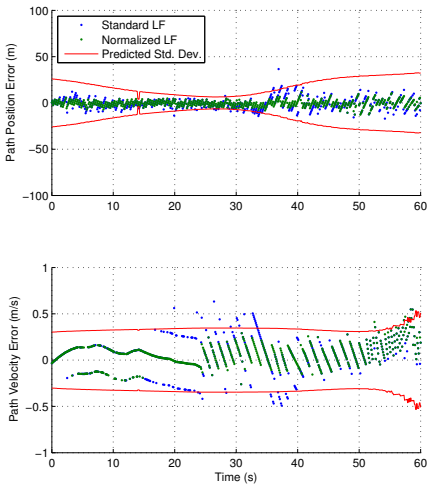

(b) Estimation performance with all receivers.

Figure 3.16: GS algorithm performance with NCVP model for the WSMR experiment with different receiver combinations. Note that for two receivers, the standard and normalized likelihood function yield the same result. Predicted standard deviation based on CRLB approximation is tuned to match true errors.

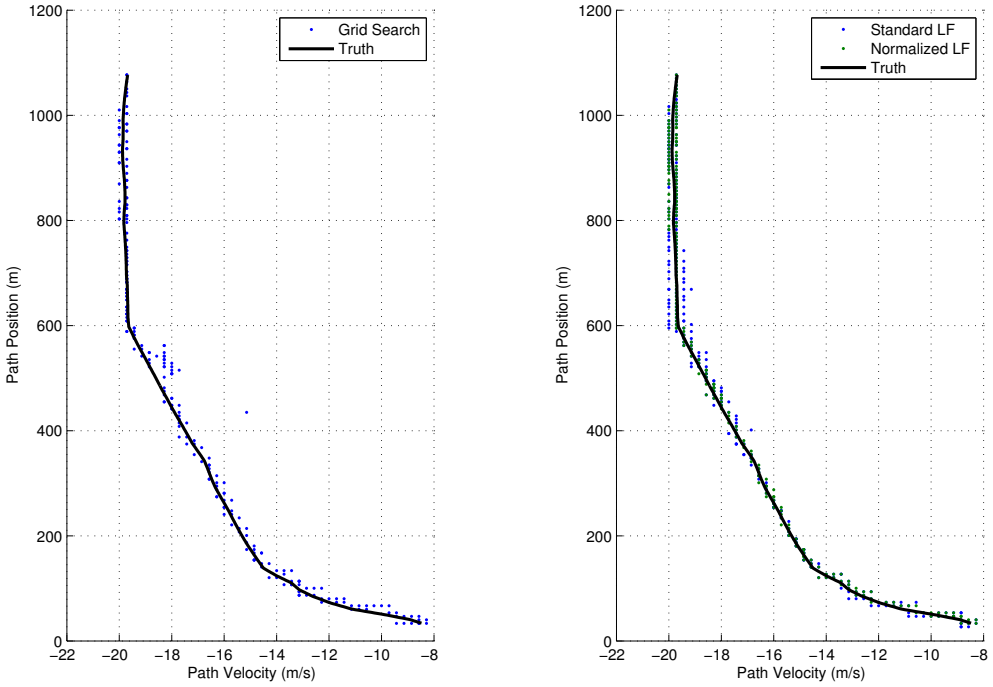

(a) Estimated NCVP state with receivers 2 and 3.

(b) Estimated NCVP state with all receivers.

Figure 3.17: Estimated path position and velocity for GS algorithm with different receiver combinations compared to truth for WSMR dataset.

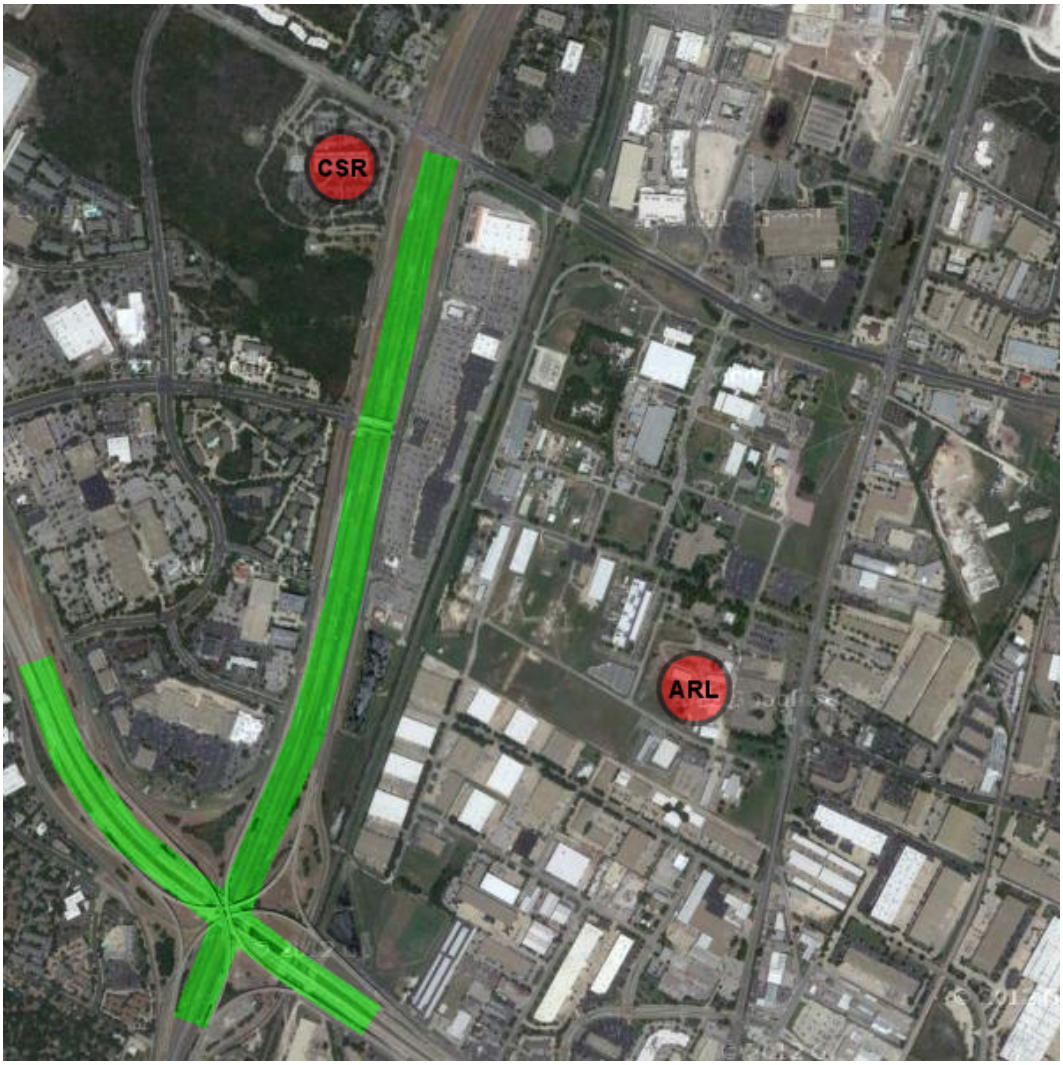

Figure 3.18: Map of UTEN with jamming localization area highlighted in green and receiver locations denoted by red markers.

emitter and reference antenna. When the received power exceeded a threshold, the receiver recorded a two-minute buffer of samples in memory to disk. One of these jamming events, which had a chirp signal commonly observed from PPDs, is analyzed in the sequel. The PF algorithm is initialized with a flat prior over the NCVP state space and begins to track the jammer as shown in Figs. 3.19, 3.20, and 3.21. Further analysis is not available for this experiment due to the loss of the recorded data and lack of time to maintain and operate UTEN.

3.6.3 UAV Experiment

The UAV experiment was designed to mimic applications where a stationary base and dynamic rover platform work together to locate a target. The sensor payload used in the following airborne experiments was a self-contained prototype dual-input sensor shown in Fig. 3.22. The dual-frequency Stereo board from Nottingham Scientific Ltd. was chosen to be the sensor's RF frontend. An Ettus N200 USRP driven by an oven-controlled crystal oscillator (OXCO) was considered, but the weight, power consumption, and form factor precluded the higher-quality combination from being selected. The Stereo was designed for prototyping purposes in emitter localization applications since the frontend has a common clock that drives a GPS L1-only MAX2769-based channel for sensor synchronization and an L-band ($800-2400\,\text{MHz}$) MAX2112-based channel. The L1-only samples were 2-bit real and the L-band samples were 3-bit complex. The Stereo's sampling rate was configured to $5.3\,\text{MHz}$ (and therefore generated data at $5.3\,{}^{\text{MB}}/_{\text{s}}$). The baseband samples were trans-

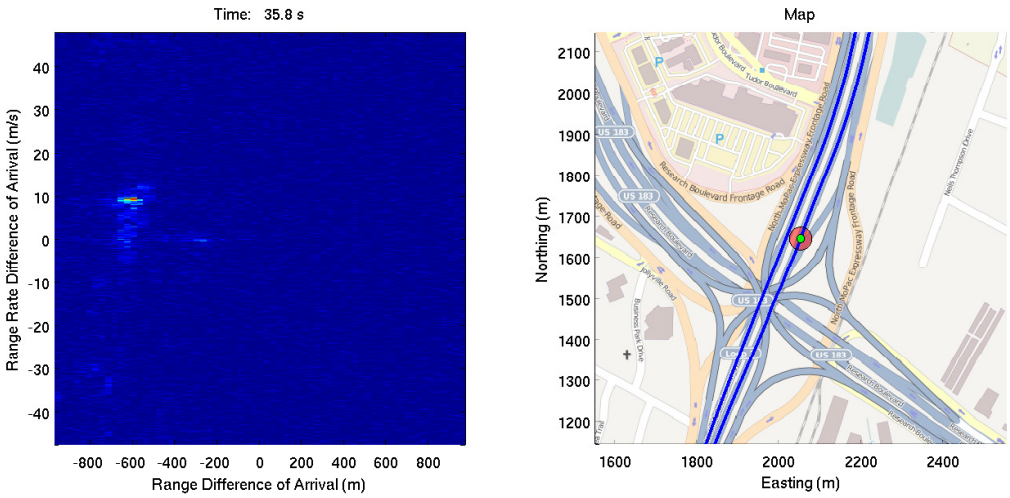

(a) Low SNR, some spreading in the ambiguity function.

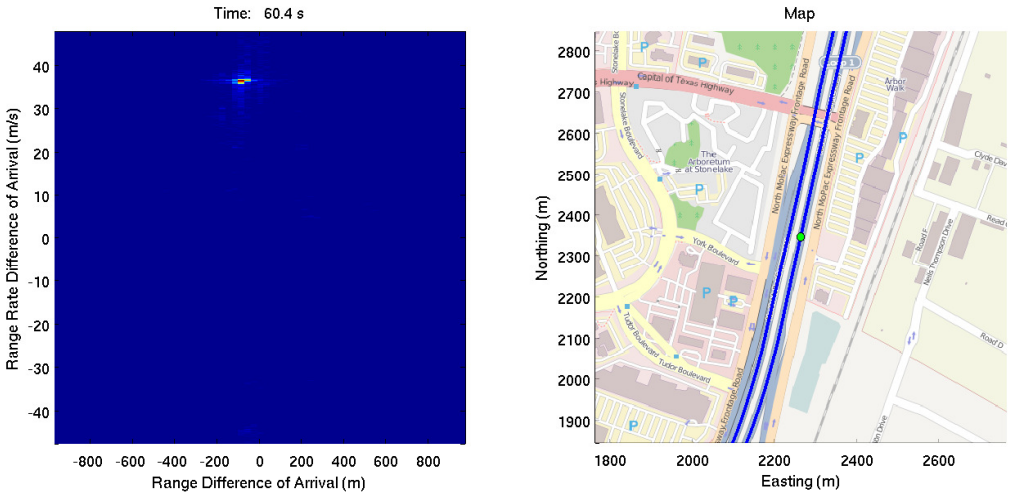

(b) High SNR, compact ambiguity function.

Figure 3.19: Particle filter tracking a northbound jammer on the Mopac highway for two different time intervals for the UTEN experiment. The left panel shows the ambiguity function, and the right panels shows the jammer on a road map with the red circle indicating 1-sigma deviation of the estimate.

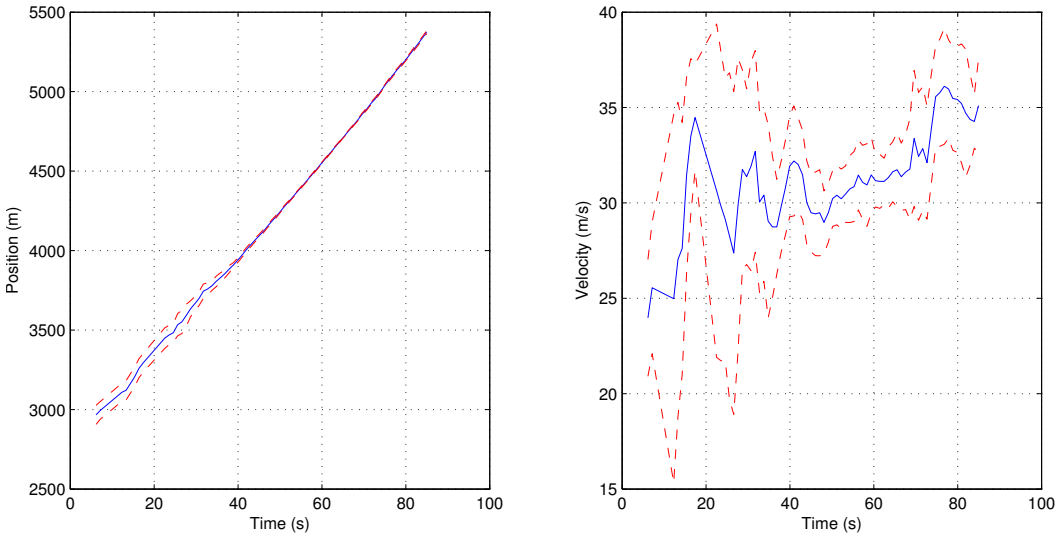

Figure 3.20: Particle filter's estimate of the jammer state over time with red dashed lines indicating 1-sigma deviation estimates for the UTEN experiment. The left panel shows the position state, and the right panel shows the velocity state (both along the highway).

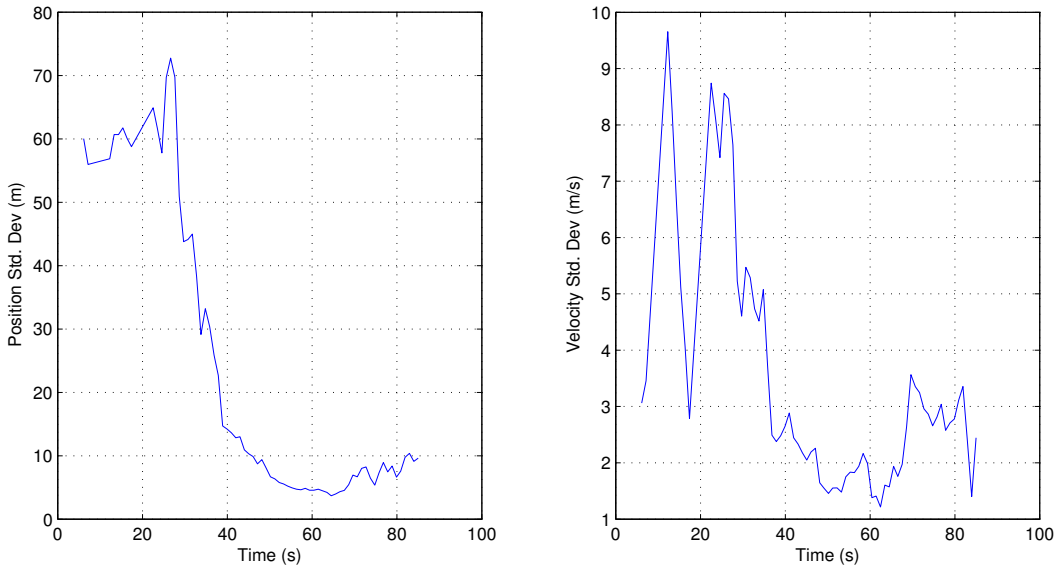

Figure 3.21: Particle filter's estimate of 1-sigma deviation over time for UTEN experiment.

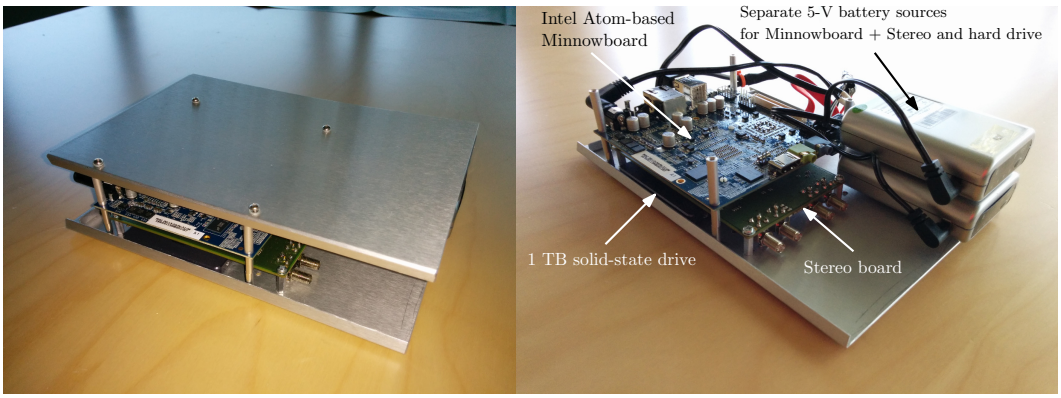

Fully-assembled Stereo-based EMLOC sensor

Stereo-based EMLOC sensor with top cover removed

Figure 3.22: The prototype portable self-contained dual-input UAV sensor.

mitted over USB to an Intel Minnowboard single-board computer (SBC). The Minnowboard was selected over smaller and more powerful ARM-based SBCs since only Intel-compiled proprietary firmware for the Stereo was available. The SBC could transmit the digital samples over a wireless network connection for real-time cross-correlation or store the samples to a Samsung solid-state drive (SSD) for post-processing. Lastly, two 5 V rechargeable battery packs (25 Wh each) were included to power the payload for about two hours. The total weight of the payload including antennas was approximately 1 kg.

The 3DR DIY quadcopter with Pixhawk autopilot, shown in Fig. 3.23, was chosen to carry the sensor payload after the loss of the lab's Hornet Mini. With no payload, the entire UAV draws a total of 20 A in hover. The UAV's four motors are rated at 20 A each. However, with the payload, the total amperage draw increases to 45 A in hover, which is still within each individual motor's rating (unless the UAV is grossly unbalanced). The total flight time with one 6600 mAh battery was about five minutes.

Figure 3.23: 3DR quadcopter based on DIY Quad Kit.

The setup was taken to a model aircraft field in November 2014 to test locating the emitter while the UAV is airborne. The USRP E100 emitter was configured to transmit an 8 MHz spread-spectrum signal at 902 MHz. The UAV was commanded to fly a box formation around the emitter and did not exceed a speed of 3.6 $^{\mathrm{m}}/_{\mathrm{s}}$. Unfortunately, due to a procedural error, an independent "truth" measurement of the emitter location was not made. For cross-correlation-based geolocation, at least two sensors are required, so a static sensor using the same frontend and emitter antenna as the UAV sensor was set up nearby. The differential position and time solution between the UAV and static sensor was computed with a carrier-phase differential GPS (CDGPS) algorithm, shown in Fig. 3.24. The carrier-phase residuals, shown in Fig. 3.25, are zero-mean indicating a high probability of ambiguity convergence.

Each geolocation algorithm was applied to a 30-second interval of the collected experimental data. The UAV path partially encircles the emitter during this interval in order to provide good emitter-receiver geometry. The subaccumulation integration interval, maximum geometric time offset, and number of cross-correlation offsets were set to

$$
\begin{aligned}
T_{\mathrm{sub}} &= \frac{33.3\,\mathrm{cm}}{4 \times 8\,{}^{\mathrm{m}}/_{\mathrm{s}}} \approx 10\,\mathrm{ms}, \\
\bar{\tau}_{\mathrm{sub}} &= \frac{100\,\mathrm{m}}{3 \times 10^{8}\,{}^{\mathrm{m}}/_{\mathrm{s}}} \approx 333\,\mathrm{ns}, \\
N_{\mathrm{sub}} &= 2 \times \mathrm{ceil}\left(\frac{333\,\mathrm{ns}}{189\,\mathrm{ns}}\right) + 1 = 5.
\end{aligned}
$$

The CAF is not useful for analysis here since the UAV dynamics breaks the assumption of constant T/FDOA over any useful coherent integration interval.

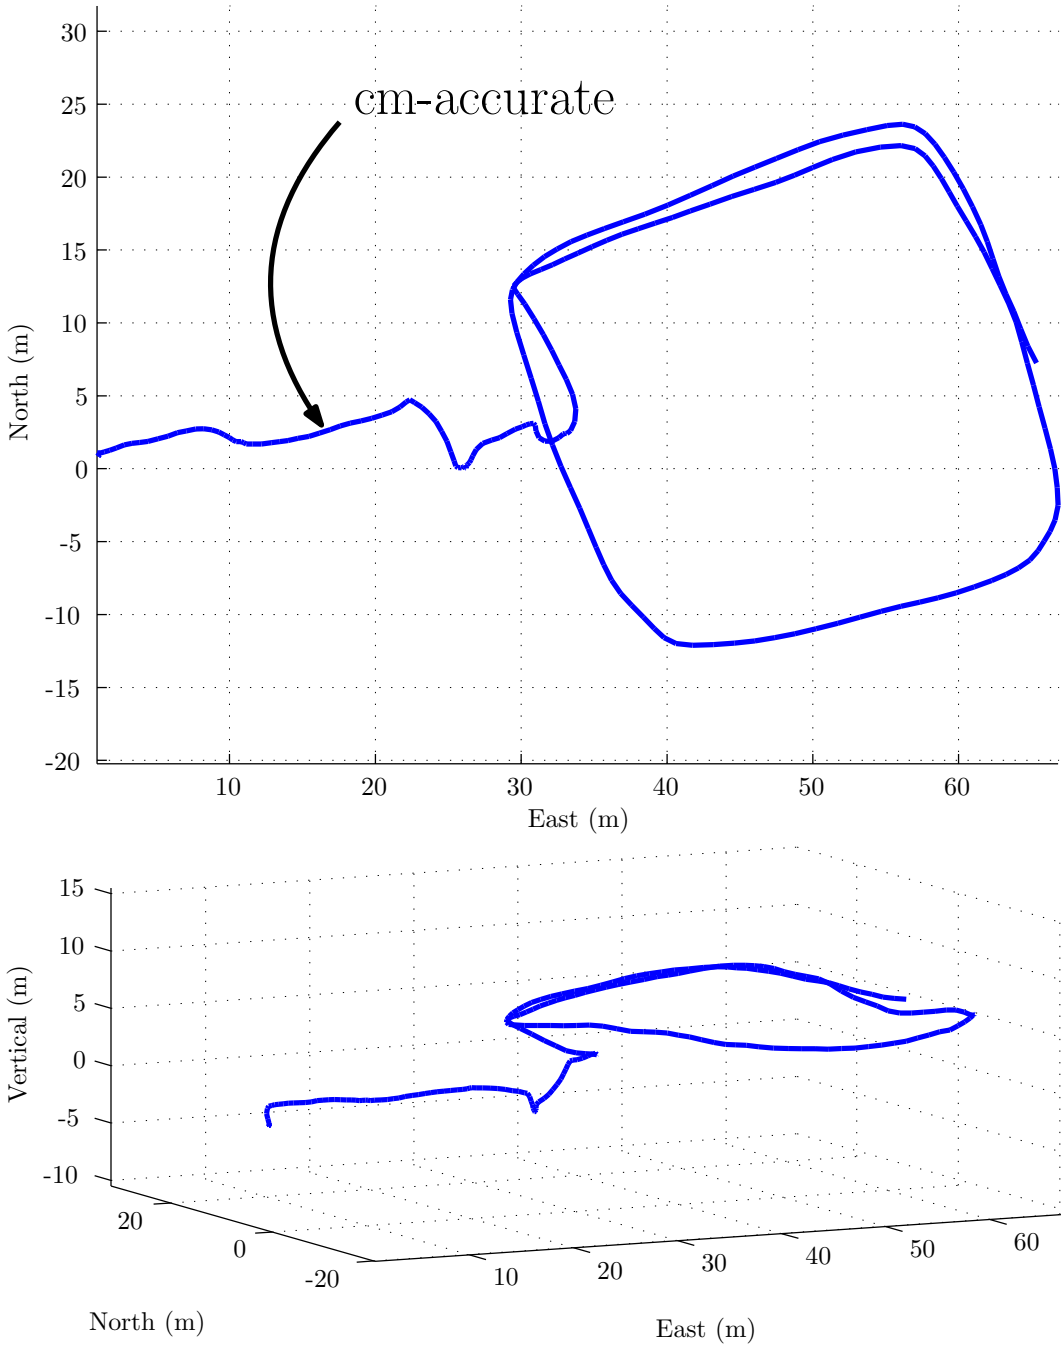

Figure 3.24: UAV path derived from CDGPS at model aircraft field. The origin represents the static sensor.

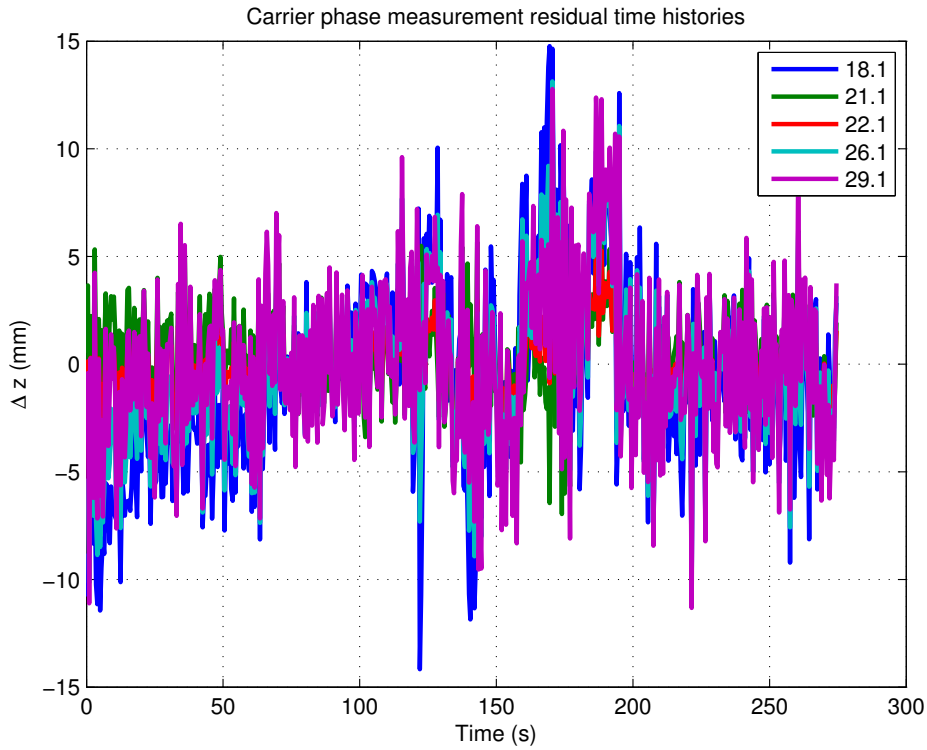

Figure 3.25: Carrier-phase residuals for UAV path at model aircraft field.

The GS algorithm with the standard likelihood function is applied to the data for various coherent integration intervals. Non-coherent averaging is used to accumulate information from each integration interval in order to improve the overall estimability of the estimate. The value of the log-likelihood function of each grid point in each coherent interval within the total 30-second interval is summed to produce the radar-like images shown in Fig. 3.26. Note that longer coherent integration intervals lead to narrower likelihood peaks due to the moving UAV receiver, an effect that is similar to synthetic aperture radar. However, longer integration intervals require increased grid resolution to ensure the peak is actually found with GS. Note that the PF algorithm with $q = 0$ and no resampling effectively performs non-coherent averaging over the initial particle set because the final particle weight is the product of the likelihood function values from the current and previous integration intervals. The results of the PF algorithm are shown in Fig. 3.27 with $T = 1\,\mathrm{s}$, $N_p = \{100, 300\}$, $\gamma = 10$, and $q = \{10^{-1}, 10^{-2}\}$, where "truth" is given by the GS estimate with 15-second coherent integration. The best performance from this set of parameters seems to be $N_p = 300$ and $q = 10^{-2}$, resulting in decimeter-level accuracy.

To verify the accuracy of the system, the USRP emitter was placed near an independently-verified location on the Woolrich Labs roof and the UAV with integrated sensor traversed a marked pathway (not under its own power due to limited space on the roof) as shown in Fig. 3.28. The accuracy of the system (as measured by the distance from the pixel with the most power to the known location) in the single run is approximately 15 cm of error over

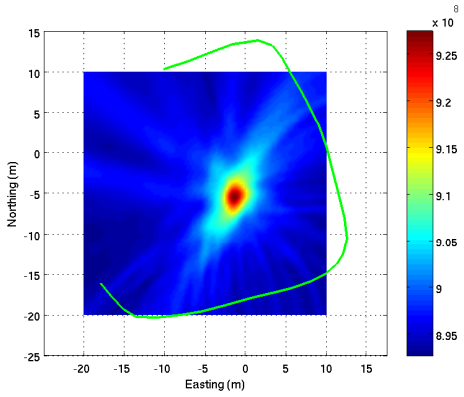

(a) 1-second coherent integration

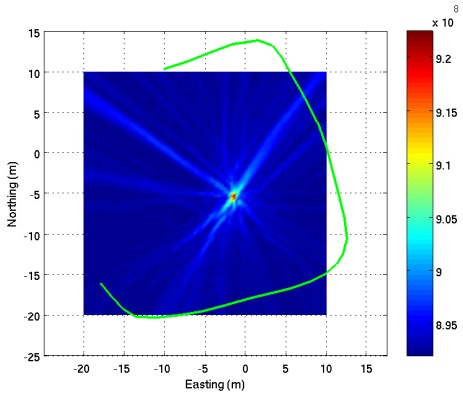

(b) 5-second coherent integration

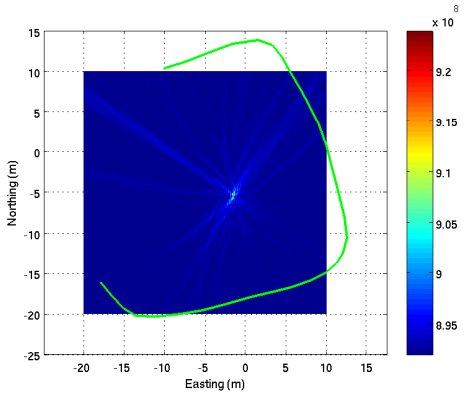

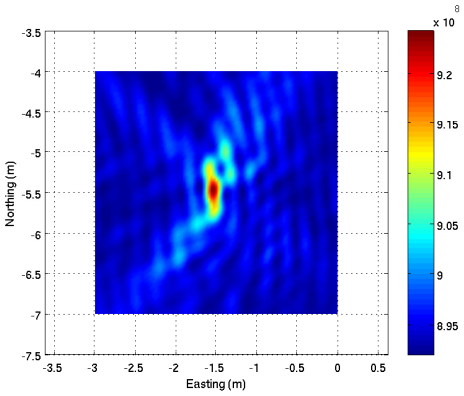

(c) 15-second coherent integration (zoomed in on right)

Figure 3.26: GS emitter localization results at the model aircraft field with non-coherent averaging. The green trace is the UAV path over a 30-second interval. In the radar-like images, red indicates the strongest accumulation of log-likelihood values, while blue is the weakest.

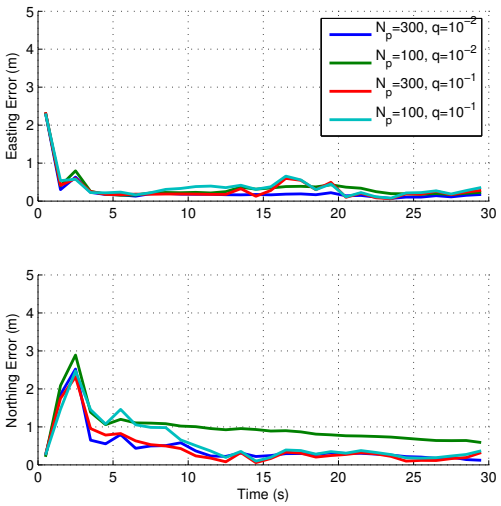

(a) True one-sigma estimation error averaged over ten Monte-Carlo trials.

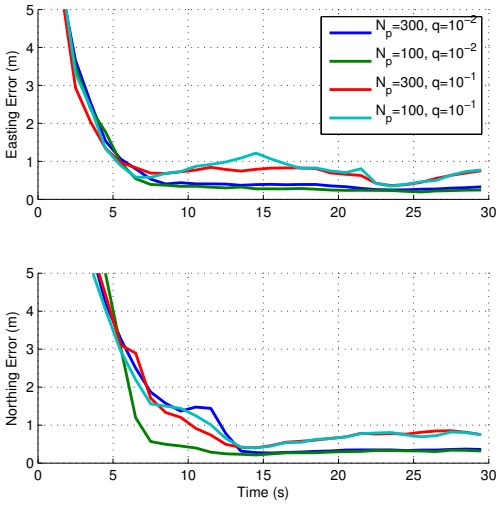

(b) Predicted one-sigma estimation error for a single trial.

Figure 3.27: PF algorithm performance with NS model for the UAV experiment.

Figure 3.28: Quadcopter on pathway in Woolrich Labs roof tests.

different coherent integration intervals as shown in Fig. 3.29. The estimated accuracy is on the same order as the total errors of the "truth" location.

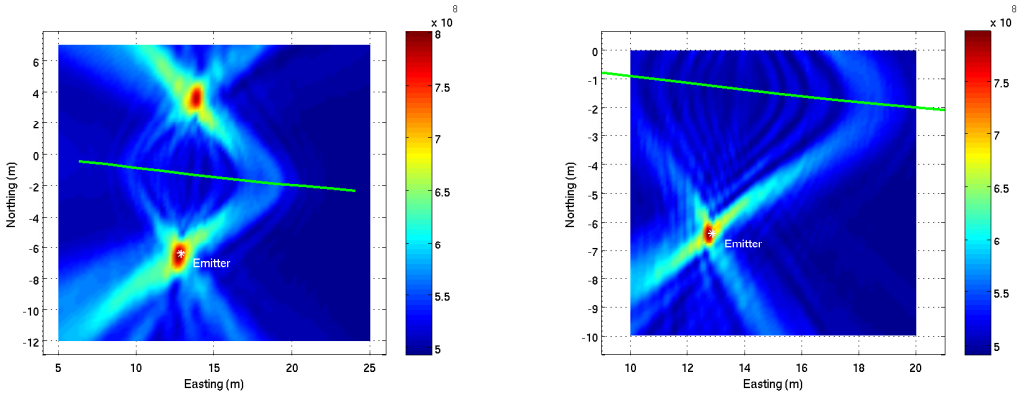

(a) 2-second coherent interval (11.5 cm error) (b) 5-second coherent interval (16.7 cm error)

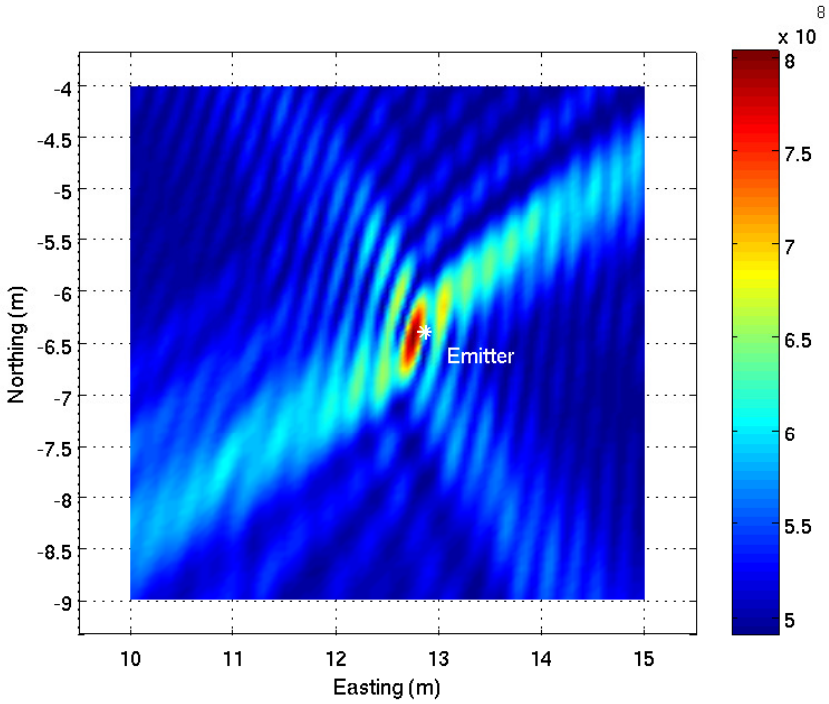

(c) 10-second coherent interval (12.9 cm error)

Figure 3.29: GS emitter localization results on the roof with non-coherent averaging. The green trace is the UAV path over a 10-second interval. The "truth" emitter location is indicated by a white star. In the radar-like images, red indicates the strongest accumulation of log-likelihood values, while blue is the weakest.

Chapter 4

Conclusions

A detection framework has been developed to detect spoofing attacks in maritime environments based solely on Doppler log, gyrocompass, and potentially-spoofed GPS measurements. Although more sophisticated spoofing detection techniques such as the dual-antenna defense are much more effective, the framework was developed to be easily implementable in ECDIS software currently available on all ships. The framework is based on a dynamics model that captures the essential features of the environmental disturbances, which are dominated by ocean currents and wind. Although this dissertation focused on the maritime dynamics model, the framework can be easily applied to an inertial measurement unit or clock model, which both have drift parameters governed by Gauss-Markov processes, but is left for future work. In addition, this framework could be incorporated into a broader probabilistic secure perception framework that includes both secure physical layer sensing and secure estimation in a layered defense to GPS spoofing.

Although the detection framework was limited to a NIS detection statistic, future work could explore more powerful statistics that consider the goal-oriented nature of the attack. This dissertation derived the performance of the detection framework, which is captured by the integrity risk or probability

of HMI. The sampling time of the framework was optimized by minimizing the maximum integrity risk given a set of possible attack profiles. In the present work, the attack profiles for the maritime case were restricted due to ease of implementation in the spoofer to linear ramps, which are close to the worst-case fault profiles derived using the methods in [67]. Just as aviation has developed rigorous integrity risk standards for GPS faults, maritime regulatory authorities can use the detection framework analysis to compute the best possible integrity risk given reasonable values for real-world disturbance and attack parameters and the maximum acceptable continuity risk. Lastly, Monte-Carlo simulations verified the theoretical integrity risk of the detection framework and an unprecedented experiment demonstrated the feasibility of conducting a spoofing attack on an actual vessel.

A passive RF emitter localization system has been developed originally to find GPS jammers posing a threat to critical infrastructure. The system can find emitters of unknown waveform, regardless of intention, for both the war-fighter and civil enforcement of protected RF bands. The system avoided making an intermediate T/FDOA estimate based on only short intervals for which the constant T/FDOA approximation is valid and instead used direct geolocation and long coherent integration techniques, thus, in theory, improving the system's estimation performance with weak emitters and multipath. Three field experiments in various emitter and receiver configurations demonstrated the capability of the system in practice. Given the emphasis on efficiency of implementation, the system presented in this dissertation has already shown real-time performance with the modern computational resources of desktop

computers and can be readily deployed in a real-time configuration with sufficient network capacity.

Although the present work considers only a single emitter, future work could investigate the implications of multiple emitters on the estimation architecture in computational and algorithmic complexity. Blindly applying the single-emitter direct geolocation algorithms to the multiple emitter problem leads to problems. The GS algorithm will clearly only lock on to the strongest emitter. However, the PF algorithm may develop a multimodal distribution with the peak of the modes related to the SNR of the emitter signals. Nonetheless, without a sufficiently large number of particles, the resampling step will prune unlikely modes, which is important for rejecting multipath-induced errors. Even if the particle filter sustains tracking a multimodal distribution, a clustering algorithm will be required to associate each particle with a particular emitter before making an estimate. A properly-dimensioned particle filter that duplicates the state space for each additional emitter may be able to track multiple emitters nearly optimally, but the particle filter's curse of dimensionality strikes as the computational complexity is $O\left(M^N\right)$, where M is the number of particles per emitter and N is the number of emitters.

A potential path forward may consider a bank of particle filters, where the likelihood function of each subsequent filter in the bank is masked by some function of the particles of the previous filters. Note that in steady-state each subsequent filter should track a weaker emitter, and the total computational complexity is given by $O\left(MN\right)$. Finally, future work could develop motion planning algorithms for dynamic receiver platforms based on an information-

seeking control law. For example, the control law developed in [120] for two-step geolocation, which generates feedback from the particle filter approximation of the emitter state probability distribution, could be easily applied to the direct-geolocation PF algorithm.

Appendix A

Generalized Sensor Deception Detection

This appendix considers an alternative formulation of the sensor deception detection framework presented in Chapter 2. The formulation in the sequel allows consideration of other detection statistics than the normalized innovation squared (NIS) statistic and other deception profiles than linear ramps.

A.1 Sensor Deception Model

Consider a discrete-time linear dynamics and measurement model

$$x(k+1) = Fx(k) + Gu(k) + v(k)$$
$$z(k) = Hx(k) + f(k) + w(k),$$

where k is the time index, $x(k)$ is the state vector of length n_x, $u(k)$ is the control vector of length n_u, $z(k)$ is the measurement vector of length n_z, $v(k) \sim \mathcal{N}(0, Q)$ is the white process noise vector, $w(k) \sim \mathcal{N}(0, R)$ is the white measurement noise vector, and $f(k)$ is the deterministic deception vector. The matrices F, G, and H complete the description of the linear time-invariant system model. Note that $f(k)$ is analogous to the fault vector in the fault-detection problem. The discrete-time model has an implied sampling

time T_s. In this work, the discrete-time linear model suffices for analysis of sensor deception in navigation and clock applications.

Although not necessarily useful at this point, the equations for the defender's estimator are presented here to introduce notation. The optimal sequential estimator for the system model under deception-free conditions (i.e. $\forall k \;\; f(k) = 0$) is the discrete-time Kalman filter. The recursion equations for the *a priori* and *a posteriori* estimates of the Kalman filter, $\bar{x}(k)$ and $\hat{x}(k)$, respectively, are given by

$$\bar{x}(k+1) = F\hat{x}(k) + Gu(k),$$
$$\hat{x}(k) = \bar{x}(k) + K(k)\nu(k),$$

where $\nu(k) = z(k) - H\bar{x}(k)$ is the innovation and $K(k)$ is the Kalman gain. The recursion equations for the covariance of the *a priori* and *a posteriori* estimation error, $\bar{P}(k)$ and $P(k)$, respectively, are given by

$$\bar{P}(k) = FP(k-1)F^T + Q$$
$$P(k) = (I - K(k)H)\bar{P}(k).$$

Note that the Kalman gain is given by

$$K(k) = \bar{P}(k)H^{\mathrm{T}}S^{-1}(k),$$

where

$$S(k) = H\bar{P}(k)H^{\mathrm{T}} + R.$$

In the sequel, it is assumed that the estimation error covariances have reached their steady-state values (which can be found by solving a discrete-time algebraic Riccati equation), and the index k is dropped from P, $\bar{P}$, S, and K.

In addition, the estimation biases induced by a non-zero fault vector $f(k)$ are needed to analyze the effect of sensor deception. The recursion equations for the *a priori* and *a posteriori* estimation error, $\bar{\epsilon}(k) = x(k) - \bar{x}(k)$ and $\hat{\epsilon}(k) = x(k) - \hat{x}(k)$, respectively, are given by

$$\begin{aligned}\bar{\epsilon}(k+1) &= F\hat{\epsilon}(k) + v(k) \\ \hat{\epsilon}(k) &= (I - KH)\bar{\epsilon}(k) - K(w(k) + f(k)).\end{aligned}$$

Then, the expected value of the estimation errors and innovation are given by

$$\begin{aligned}\mathbb{E}[\bar{\epsilon}(k+1)] &= F\mathbb{E}[\hat{\epsilon}(k)] \\ \mathbb{E}[\hat{\epsilon}(k)] &= (I - KH)\mathbb{E}[\bar{\epsilon}(k)] - Kf(k) \\ \mathbb{E}[\nu(k)] &= f(k) + H\mathbb{E}[\bar{\epsilon}(k)],\end{aligned}$$

which are clearly biased for non-zero fault vectors.

A.2 Batch Residual and Filter Innovations

The equivalence of the batch linear estimation residual to the Kalman filter innovations is useful in proving optimality of certain detection statistics. Although the equivalence is intuitively obvious, a proof is provided for completeness. For compactness, subscripts are used to indicate the time index, i.e. $x_k = x(k)$. In addition, without loss of generality, $\forall k\, u(k) = 0$. First, consider the one-step cost function

$$\begin{aligned}J(x_{k-1}, x_k) = &(x_{k-1} - \hat{x}_{k-1})^{\mathrm{T}} P_{k-1}^{-1} (x_{k-1} - \hat{x}_{k-1}) \\ &+ (x_k - Fx_{k-1})^{\mathrm{T}} Q^{-1} (x_k - Fx_{k-1}) \\ &+ (z_k - Hx_k)^{\mathrm{T}} R^{-1} (z_k - Hx_k).\end{aligned} \tag{A.1}$$

The previous state x_{k-1} is eliminated from (A.1) by finding its optimal value in terms of x_k and $\hat{x}_{k-1}$, also known as the smoother equation. The optimal value $x_{k-1}^{\star}$ is found by setting the gradient of the cost function with respect to x_{k-1} to zero,

$$\frac{\partial J\left(x_{k-1}, x_k\right)}{\partial x_{k-1}} = P_{k-1}^{-1}\left(x_{k-1} - \hat{x}_{k-1}\right) - F^{\mathrm{T}} Q^{-1}\left(x_k - F x_{k-1}\right) = 0. \tag{A.2}$$

Letting $\bar{x}_k = F\hat{x}_{k-1}$ and rearranging (A.2) yields

$$\begin{aligned} x_{k-1}^{\star} &= \left(P_{k-1}^{-1} + F^{\mathrm{T}} Q^{-1} F\right)^{-1}\left(P_{k-1}^{-1}\hat{x}_{k-1} + F^{\mathrm{T}} Q^{-1} x_k\right) \\ &= \hat{x}_{k-1} + \left(P_{k-1}^{-1} + F^{\mathrm{T}} Q^{-1} F\right)^{-1} F^{\mathrm{T}} Q^{-1}\left(x_k - \bar{x}_k\right) \end{aligned} \tag{A.3}$$

Considering the first two terms in (A.1) and noting the optimality condition in (A.2) yields

$$\begin{aligned} J'\left(x_{k-1}, x_k\right) &= \left(x_{k-1} - \hat{x}_{k-1}\right)^{\mathrm{T}} P_{k-1}^{-1}\left(x_{k-1} - \hat{x}_{k-1}\right) \\ &\quad + \left(x_k - F x_{k-1}\right)^{\mathrm{T}} Q^{-1}\left(x_k - F x_{k-1}\right) \\ &= \left(x_{k-1} - \hat{x}_{k-1}\right)^{\mathrm{T}} P_{k-1}^{-1}\left(x_{k-1} - \hat{x}_{k-1}\right) \\ &\quad + \left(x_k - \bar{x}_k - F\left(x_{k-1} - \hat{x}_{k-1}\right)\right)^{\mathrm{T}} Q^{-1}\left(x_k - F x_{k-1}\right) \\ &= \left(x_{k-1} - \hat{x}_{k-1}\right)^{\mathrm{T}} P_{k-1}^{-1}\left(x_{k-1} - \hat{x}_{k-1}\right) \\ &\quad - \left(x_{k-1} - \hat{x}_{k-1}\right)^{\mathrm{T}} F^{\mathrm{T}} Q^{-1}\left(x_k - F x_{k-1}\right) \\ &\quad + \left(x_k - \bar{x}_k\right)^{\mathrm{T}} Q^{-1}\left(x_k - F x_{k-1}\right) \\ &= \left(x_{k-1} - \hat{x}_{k-1}\right)^{\mathrm{T}} \cancelto{0}{\frac{\partial J\left(x_{k-1}, x_k\right)}{\partial x_{k-1}}} \\ &\quad + \left(x_k - \bar{x}_k\right)^{\mathrm{T}} Q^{-1}\left(x_k - F x_{k-1}\right) \\ &= \left(x_k - \bar{x}_k\right)^{\mathrm{T}} Q^{-1}\left(x_k - \bar{x}_k - F\left(x_{k-1} - \hat{x}_{k-1}\right)\right). \end{aligned} \tag{A.4}$$

Substituting (A.3) in (A.4) and judicious use of the matrix inversion lemma yields

$$\begin{aligned}
J''(x_k) &= J'\left(x_{k-1}^{\star}, x_k\right) \\
&= (x_k - \bar{x}_k)^{\mathrm{T}} Q^{-1} (x_k - \bar{x}_k) \\
&\quad - (x_k - \bar{x}_k)^{\mathrm{T}} Q^{-1} F^{\mathrm{T}} \left(P_{k-1}^{-1} + F^{\mathrm{T}} Q^{-1} F\right)^{-1} F^{\mathrm{T}} Q^{-1} (x_k - \bar{x}_k) \\
&= (x_k - \bar{x}_k)^{\mathrm{T}} \left(F P_{k-1} F^{\mathrm{T}} + Q\right)^{-1} (x_k - \bar{x}_k) \\
&= (x_k - \bar{x}_k)^{\mathrm{T}} \bar{P}_k^{-1} (x_k - \bar{x}_k),
\end{aligned}$$

where

$$\bar{P}_k = F P_{k-1} F^{\mathrm{T}} + Q.$$

The one-step cost can be rewritten as

$$\begin{aligned}
J'''(x_k) &= J\left(x_{k-1}^{\star}, x_k\right) \\
&= (x_k - \bar{x}_k)^{\mathrm{T}} \bar{P}_k^{-1} (x_k - \bar{x}_k) + (z_k - H x_k)^{\mathrm{T}} R^{-1} (z_k - H x_k).
\end{aligned}$$

Let the following be defined as

$$\begin{aligned}
P_k^{-1} &= \bar{P}_k^{-1} + H^{\mathrm{T}} R^{-1} H \\
S_k &= H \bar{P}_k H^{\mathrm{T}} + R \\
\nu_k &= z_k - H \bar{x}_k \\
\hat{x}_k &= \bar{x}_k + P_k H^{\mathrm{T}} R^{-1} \nu_k.
\end{aligned}$$

Then, with appropriate substitutions,

$$\begin{aligned}
J'''(x_k) &= (x_k - \bar{x}_k)^{\mathrm{T}} \bar{P}_k^{-1} (x_k - \bar{x}_k) \\
&\quad + (\nu_k - H(x_k - \bar{x}_k))^{\mathrm{T}} R^{-1} (\nu_k - H(x_k - \bar{x}_k)) \\
&= (x_k - \bar{x}_k)^{\mathrm{T}} \left(\bar{P}_k^{-1} + H^{\mathrm{T}} R^{-1} H\right) (x_k - \bar{x}_k) \\
&\quad - 2\nu_k^{\mathrm{T}} R^{-1} H (x_k - \bar{x}_k) + \nu_k^{\mathrm{T}} R^{-1} \nu_k \\
&= \left(x_k - \hat{x}_k + P_k H^{\mathrm{T}} R^{-1} \nu_k\right)^{\mathrm{T}} P_k^{-1} \left(x_k - \hat{x}_k + P_k H^{\mathrm{T}} R^{-1} \nu_k\right) \\
&\quad - 2\nu_k^{\mathrm{T}} R^{-1} H (x_k - \bar{x}_k) + \nu_k^{\mathrm{T}} R^{-1} \nu_k \\
&= (x_k - \hat{x}_k)^{\mathrm{T}} P_k^{-1} (x_k - \hat{x}_k) \\
&\quad + 2\nu_k^{\mathrm{T}} R^{-1} H ((x_k - \hat{x}_k) - (x_k - \bar{x}_k)) \\
&\quad + \nu_k^{\mathrm{T}} \left(R^{-1} + R^{-1} H P_k H^{\mathrm{T}} R^{-1}\right) \nu_k \\
&= \cdots - 2\nu_k^{\mathrm{T}} R^{-1} H P_k H^{\mathrm{T}} R^{-1} \nu_k \\
&\quad + \nu_k^{\mathrm{T}} \left(R^{-1} + R^{-1} H P_k H^{\mathrm{T}} R^{-1}\right) \nu_k \\
&= \cdots + \nu_k^{\mathrm{T}} \left(R^{-1} - R^{-1} H \left(\bar{P}_k^{-1} + H^{\mathrm{T}} R^{-1} H\right)^{-1} H^{\mathrm{T}} R^{-1}\right) \nu_k.
\end{aligned}$$

A final application of the matrix inversion lemma yields

$$J'''(x_k) = (x_k - \hat{x}_k)^{\mathrm{T}} P_k^{-1} (x_k - \hat{x}_k) + \nu_k^{\mathrm{T}} S_k^{-1} \nu_k. \tag{A.5}$$

Now consider the batch linear estimation cost function with prior information, dynamical, and measurement constraints:

$$\begin{aligned}
J_0(x_0, \ldots, x_M) &= (x_0 - \hat{x}_0)^{\mathrm{T}} P_0^{-1} (x_0 - \hat{x}_0) \\
&\quad + \sum_{k=1}^{M} (x_k - F x_{k-1})^{\mathrm{T}} Q^{-1} (x_k - F x_{k-1}) \\
&\quad + \sum_{k=1}^{M} (z_k - H x_k)^{\mathrm{T}} R^{-1} (z_k - H x_k).
\end{aligned} \tag{A.6}$$

Recursively applying the transformations of the one-step cost function yields

$$\begin{aligned} J_L\left(x_L, \ldots, x_M\right) = & \sum_{k=1}^{L} \nu_k^{\mathrm{T}} S_k^{-1} \nu_k + \left(x_L - \hat{x}_L\right)^{\mathrm{T}} P_L^{-1} \left(x_L - \hat{x}_L\right) \\ & + \sum_{k=L+1}^{M} \left(x_k - F x_{k-1}\right)^{\mathrm{T}} Q^{-1} \left(x_k - F x_{k-1}\right) \\ & + \sum_{k=L+1}^{M} \left(z_k - H x_k\right)^{\mathrm{T}} R^{-1} \left(z_k - H x_k\right), \end{aligned}$$

until only the last state remains

$$J_M\left(x_M\right) = \sum_{k=1}^{M} \nu_k^{\mathrm{T}} S_k^{-1} \nu_k + \left(x_M - \hat{x}_M\right)^{\mathrm{T}} P_M^{-1} \left(x_M - \hat{x}_M\right). \tag{A.7}$$

Then, the optimal value of x_M is by inspection $x_M^\star = \hat{x}_M$, and the rest of the optimal states $x_k^\star$ for $0 \leq k < M$ can be recursively solved with the smoother equation (A.3), which, all together, form the solution to the batch estimation problem. The irreducible part of (A.6) is the residual of the batch estimation problem and is given by the sum of the NIS,

$$J_{\text{resid}} = J_0\left(x_0^\star, \ldots, x_M^\star\right) = J_M\left(x_M^\star\right) = \sum_{k=1}^{M} \nu_k^{\mathrm{T}} S_k^{-1} \nu_k. \tag{A.8}$$

A.3 Detection Statistic

The detection framework developed in this work borrows concepts from GPS integrity monitoring in aviation applications [66, 121] and the fault detection literature [62], which are applied here to the "fraud detection" problem. If the fault profile can be parameterized by a "jump" magnitude and a start time for the jump, then the generalized likelihood ratio (GLR) framework allows

close to optimal probability of detection. In GLR, the most likely start time is determined from a bank of matched filters tuned to different start times. Then, a decision between the hypotheses is made based on the matched filter correlation from the most likely start time. The bank of filters is usually truncated to a sliding window of most recent possible start times to avoid linear computational growth in time. GLR is avoided in this work due to the framework's complexity, which can make evaluating integrity risk difficult. Sequential decision procedures such as the sequential probability ratio test (SPRT) optimally minimize the time-to-detect for a fixed probability of detection and false alarm, which for some applications may be a suitable proxy for integrity risk. However, the "indifference region" of the SPRT, where no decision is made based on the current time decision statistic, thus requiring more samples, can make evaluating the integrity risk difficult. Instead, a fixed sample size (FSS) decision procedure, as outlined in [71], is considered in this work.

In the framework, a detection test over a window of M samples decides between two hypotheses—a null hypothesis H_0 indicating nominal operating conditions, and an alternative hypothesis H_1 indicating a deception attack is underway i.e. non-zero $f(k)$. Let $q(l)$ be a test statistic that monitors the presence of sensor deception every M samples such that

$$q(l) \underset{H_0}{\overset{H_1}{\gtrless}} \lambda,$$

where λ is a threshold chosen to maintain a constant false-alarm rate and l is a positive integer. At $k = 0$, the null hypothesis is assumed to be true, and

at some later time $k > k_0$, a transition to the alternative hypothesis occurs, although a transition is not necessarily guaranteed to occur. Additionally, if a deception attack begins, it is assumed the attack will continue until either hazardous conditions occur or the attack is detected. The detection statistic $q(l)$ must remain below a threshold in order to assume the null hypothesis. The threshold λ satisfies

$$P(q(l) > \lambda | H_0) = \frac{T_d}{M_F}$$

to maintain the prescribed false-alarm rate M_F, where $T_d = MT_s$ is the detection interval. Note that the probability distribution of the detection statistic under the null hypothesis, and therefore λ, is independent of time.

Various forms for the detection statistic are considered given different assumptions for the fault profile $f(k)$. To simplify the forms in this section, k_0 is assumed to be known and zero and the fault profile is known to cause hazardous conditions for $k > M$ (i.e. $k_L = M + 1$), so that the defender must make a decision on whether a deception attack occurred with only the first M measurements. For an arbitrary fault profile unknown to the defender, the optimal detection statistic for a generalized likelihood ratio test (GLRT) over a batch of M measurements is the sum of the NIS, i.e.

$$q_1 = \sum_{k=1}^{M} \nu^T(k) S^{-1} \nu(k) \sim \chi^2(Mn_z, \delta_1).$$

Note that q_1 is distributed as non-central chi-squared with Mn_z degrees of freedom and non-centrality parameter

$$\delta_1 = \sum_{k=1}^{M} \left\| S^{-1/2} \mathbb{E}[\nu(k)] \right\|_2^2.$$

The proof of optimality in the GLRT sense is straightforward after noting that the residual to the dynamic batch estimation problem is the sum of NIS as shown previously, which Joerger failed to realize in [66]. The proof begins by considering the likelihood ratio

$$\Lambda\left(Z\right) = \frac{p\left(Z|H_1\right)}{p\left(Z|H_0\right)} = \frac{p\left(Z|X_1, F\right)}{p\left(Z|X_0\right)},$$

where

$$Z = \left\{z\left(1\right), \ldots, z\left(M\right)\right\}$$

is the set of all measurements,

$$X_i = \left\{x_i\left(0\right), \ldots, x_i\left(M\right)\right\}$$

is the set of all dynamical states for hypothesis H_i, and

$$F = \left\{f\left(1\right), \ldots, f\left(M\right)\right\}$$

is the set of all fault vectors under the spoofing hypothesis H_1. Since X_0, X_1, and F are unknown, the GLRT approach is to replace them with their best estimates, i.e.

$$\Lambda_{\text{GLR}}\left(Z\right) = \frac{p\left(Z|\hat{X}_1, \hat{F}\right)}{p\left(Z|\hat{X}_0\right)}.$$

For each hypothesis, there are $\left(M+1\right)n_x$ dynamical and Mn_z measurement constraints. Note that for H_1, the number of unknown degrees of freedom is equal to the number of constraints. Therefore, the residual of the best estimate is exactly zero for the observable batch linear estimation problem,

i.e. $p\left(Z|\hat{X}_1, \hat{F}\right) = 1$. For H_0, the probability is easily written in terms of NIS as

$$p\left(Z|\hat{X}_0\right) = \frac{1}{\sqrt{\left((2\pi)^{n_z} \det S\right)^M}} \exp\left(-\frac{1}{2}\sum_{k=1}^{M} \nu^T(k) S^{-1} \nu(k)\right).$$

The log-likelihood ratio is given by

$$\begin{aligned} L(Z) &= \log \Lambda_{\mathrm{GLR}}(Z) \\ &= \frac{M}{2}\left(n_z \log 2\pi + \log\det S\right) + \frac{1}{2} q_1, \end{aligned}$$

which shows that q_1 is equivalent to $L(Z)$, up to an additive constant and scaling.

However, in a deception attack, the fault profile is not arbitrary—it is designed by the attacker to eventually cause hazardously misleading conditions. If the fault profile is known to the defender, then rearranging the generalized log-likelihood ratio yields

$$\begin{aligned} L(Z) &= \log \Lambda_{\mathrm{GLR}}(Z) \\ &= \log p\left(Z|\hat{X}_1, F\right) - \log p\left(Z|\hat{X}_0\right) \\ &= -\frac{1}{2}\sum_{k=1}^{M} \left(\nu(k) - \mathbb{E}\left[\nu(k)\right]\right)^{\mathrm{T}} S^{-1}\left(\nu(k) - \mathbb{E}\left[\nu(k)\right]\right) \\ &\quad + \frac{1}{2}\sum_{k=1}^{M} \nu^T(k) S^{-1} \nu(k) \\ &= \frac{1}{2}\sum_{k=1}^{M} \mathbb{E}\left[\nu(k)\right]^{\mathrm{T}} S^{-1} \nu(k) - \frac{1}{2}\sum_{k=1}^{M} \mathbb{E}\left[\nu(k)\right]^{\mathrm{T}} S^{-1} \mathbb{E}\left[\nu(k)\right]. \end{aligned}$$

Now, it is clear that the detection statistic

$$q_2 = \sum_{k=1}^{M} \xi^T(k) S^{-1/2} \nu(k) \sim \mathcal{N}(\mu_2, 1)$$

is optimal in the GLRT sense, where

$$\mu_2 = \sum_{k=1}^{M} \xi^T(k) S^{-1/2} \mathbb{E}[\nu(k)]$$

and $\xi(k)$ are normalized weights determined by the fault profile. Note that the form of q_2 is a matched filter, where $\xi(k)$ is matched to the expected value of the normalized innovations $S^{-1/2}\mathbb{E}[\nu(k)]$. In order to set the variance of q_2 to unity, the weights must satisfy $\sum_{k=1}^{M} \xi^T(k)\xi(k) = 1$, so

$$\xi(k) = \frac{S^{-1/2}\mathbb{E}[\nu(k)]}{\sqrt{\sum_{k=1}^{M} \left\| S^{-1/2}\mathbb{E}[\nu(k)] \right\|_2^2}}.$$

Note that if the fault profile is expressed as $f(k) = \alpha \tilde{f}(k)$, where α is an unknown scalar and $\tilde{f}(k)$ is a known normalized fault profile, then q_2 is a uniformly most powerful detection statistic for $\alpha > 0$. Simply squaring q_2 allows optimal fault detection when positive and negative α are equally likely, where $q_2^2 \sim \chi^2(1, \mu_2^2)$. Alternatively, if the fault profile is expressed as $f(k) = \beta(k)\hat{d}$, where $\beta(k)$ are known positive scalars and $\hat{d}$ is an unknown unit direction vector, then the defender could use a detection statistic

$$q_3 = \left\| \sum_{k=1}^{M} w(k) S^{-1/2} \nu(k) \right\|_2^2 \sim \chi^2(n_z, \delta_3),$$

where

$$\delta_3 = \left\| \sum_{k=1}^{M} w(k) S^{-1/2} \mathbb{E}[\nu(k)] \right\|_2^2$$

and $w(k)$ is a normalized set of weights such that $\sum_{k=1}^{M} w(k)^2 = 1$. Note that for $n_z = 1$, q_3 is equivalent to q_2^2, but no claims of the optimality of q_3 are made for $n_z > 1$ in general. However, for systems where the measurement coordinates (typically vehicle position) have identical decoupled dynamics, and all possible direction vectors $\hat{d}$ are equally likely, then q_3 is a good heuristic with $w(k) = \|\xi(k)\|$. Finally, note that for $M = 1$ (i.e. $T_s = T_d$), q_3 is equivalent to q_1. By using only one measurement per detection interval, the detection framework resembles a dead-reckoning consistency check. In such a check, the *a priori* state $\bar{x}$ and innovation covariance S represents the dead-reckoned state and uncertainty after propagating for T_d seconds, respectively. The dead-reckoned state is reset by ingesting the measurement if the detection test decides H_0. However, while the detection framework prefers T_s to be large to improve detection performance for the slowest possible attack profile, the system's controller prefers T_s to be small in order for the control error to remain small. In one implementation requiring two separate estimators, the controller could ingest measurements at a different rate than the detection framework, as was considered in Chapter 2.

A.4 Worst-Case Fault Profile

If the deception attack is restricted to a family of fault profiles (such as linear or quadratic ramps where $\tilde{f}(k)$ or $\beta(k)$ are known), then the defender should choose statistics like q_2^2 or q_3 with appropriately chosen weights. However, in reality, the attacker is allowed to choose any fault profile that will cause

hazardous conditions and is smooth enough so that the control error remains small. Within this set of fault profiles and assuming the attacker knows the defender's choice of weights, the attacker will choose the worst-case (from the defender's perspective) fault profile $\bar{f}(k)$, which maximizes the integrity risk of the detection framework. For an uncoordinated attack, where the defender cannot assume alignment of the detection window with the deception attack, then a reasonable strategy for the weights is a uniform distribution over the detection window, i.e. $w(k) = 1/\sqrt{M}$ in order to maximize detection for all possible window alignments. Similarly, for the attacker, a reasonable strategy is to evenly distribute the magnitude of the normalized expected innovation $S^{-1/2}\mathbb{E}[\nu(k)]$ over the duration of the deception attack in order to minimize detection for all possible attack alignments. For $k_0 = 0$ and k_L specified and defining hazardous conditions as $\|H\mathbb{E}[\hat{\epsilon}(k)]\|_2 \geq L$, the worst-case fault profile $\bar{f}(k)$ is the solution to the following optimal control problem

$$
\begin{aligned}
\min_{f(k)} \quad & \max_k \left\| S^{-1/2}\mathbb{E}[\nu(k)] \right\|_2 \\
\text{s.t.} \quad & H\mathbb{E}[\hat{\epsilon}(k_L)] = L\hat{d} \\
& \mathbb{E}[\hat{\epsilon}(0)] = 0 \\
0 < k \leq k_L \quad & \mathbb{E}[\hat{\epsilon}(k)] = (I - KH)F\mathbb{E}[\hat{\epsilon}(k-1)] - Kf(k) \\
0 < k < k_L \quad & \mathbb{E}[\nu(k)] = f(k) + HF\mathbb{E}[\hat{\epsilon}(k-1)] \\
0 < k < k_L \quad & \|H\mathbb{E}[\hat{\epsilon}(k)]\|_2 \leq L
\end{aligned}
\tag{A.9}
$$

The minimax cost function equalizes the magnitude of the normalized innovation for all k, i.e. $\forall k \;\; \left\| S^{-1/2}\mathbb{E}[\nu(k)] \right\|_2 = C$. Note that choosing to minimize $\sum_k \left\| S^{-1/2}\mathbb{E}[\nu(k)] \right\|_2^2$ instead of (A.9) can be shown to be nearly identical to

the worst-case fault profiles produced by [67]. For one-dimensional ship dynamics with parameters $T_d = 200\,\mathrm{s}$, $\sigma_d = 0.02\,\mathrm{m/s^{1.5}}$, $T_s = 5\,\mathrm{s}$, and $\sigma_p = 3\,\mathrm{m}$, the resulting fault profiles for the minimax and two-norm costs with $L = 200\,\mathrm{m}$ and $k_L = 400$ are shown in Fig. A.1. Note that the profiles are very similar to the ramp-like modulations proposed in Chapter 2.

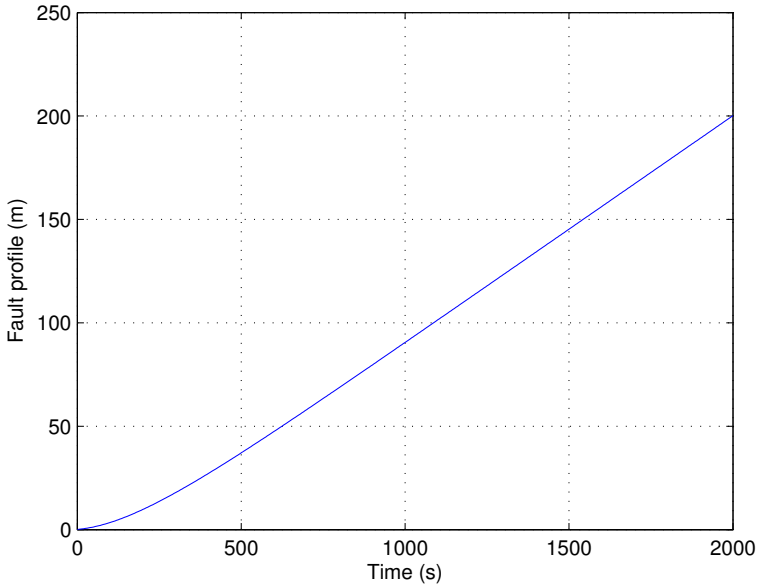

(a) Minimax worst-case fault profile.

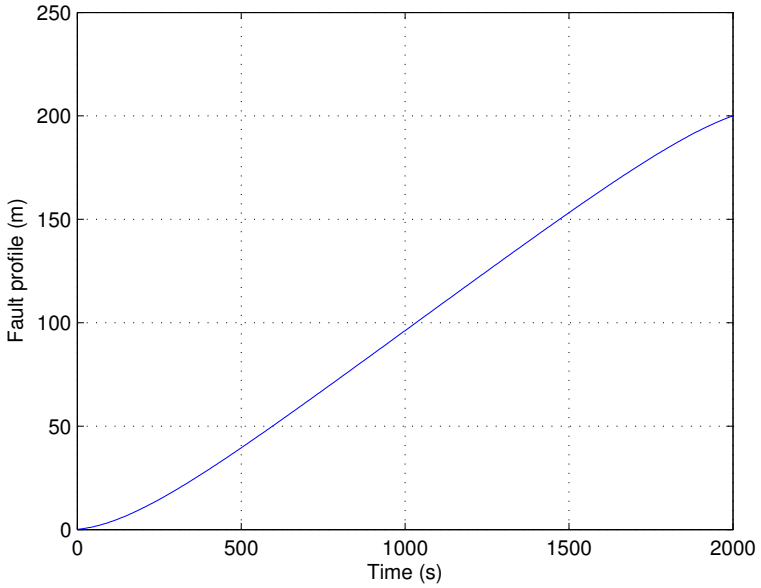

(b) Two-norm worst-case fault profile.

Figure A.1: Worst-case fault profile for one-dimensional ship dynamics using two-norm and minimax innovation heuristics.

Bibliography

[1] T. I. Fossen, *Guidance and Control of Ocean Vehicles*. New York: John Wiley and Sons, 1994.

[2] A. A. Cárdenas, S. Amin, and S. Sastry, "Research challenges for the security of control systems," in *Proceedings of the 3rd conference on Hot topics in security*. USENIX Association, 2008, pp. 1–6.

[3] A. A. Cardenas, S. Amin, and S. Sastry, "Secure control: Towards survivable cyber-physical systems," in *Distributed Computing Systems Workshops, 2008. ICDCS'08. 28th International Conference on*. IEEE, 2008, pp. 495–500.

[4] A. Cardenas, S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry, "Challenges for securing cyber physical systems," in *Workshop on future directions in cyber-physical systems security*, 2009.

[5] A. Teixeira, S. Amin, H. Sandberg, K. H. Johansson, and S. S. Sastry, "Cyber security analysis of state estimators in electric power systems," in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 5991–5998.

[6] S. Amin, X. Litrico, S. S. Sastry, and A. M. Bayen, "Stealthy deception attacks on water scada systems," in *Proceedings of the 13th ACM inter-*

national conference on Hybrid systems: computation and control. ACM, 2010, pp. 161–170.

[7] L. Xie, Y. Mo, and B. Sinopoli, "False data injection attacks in electricity markets," in *Proceedings of IEEE Smart Grid Communications (SmartGridComm) Conference*, 2010, pp. 226–231.

[8] Y. Mo, E. Garone, A. Casavola, and B. Sinopoli, "False data injection attacks against state estimation in wireless sensor networks," in *Proceedings of IEEE Decision and Control (CDC) Conference*, 2010, pp. 5967–5972.

[9] A. Gupta, C. Langbort, and T. Basar, "Optimal control in the presence of an intelligent jammer with limited actions," in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 1096–1101.

[10] Y. Liu, P. Ning, and M. K. Reiter, "False data injection attacks against state estimation in electric power grids," *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 1, p. 13, 2011.

[11] K. L. Morrow, E. Heine, K. M. Rogers, R. B. Bobba, and T. J. Overbye, "Topology perturbation for detecting malicious data injection," in *System Science (HICSS), 2012 45th Hawaii International Conference on*. IEEE, 2012, pp. 2104–2113.

[12] A. Teixeira, D. Pérez, H. Sandberg, and K. H. Johansson, "Attack models and scenarios for networked control systems," in *Proceedings of the 1st*

international conference on High Confidence Networked Systems. ACM, 2012, pp. 55–64.

[13] W. Zeng and M.-Y. Chow, "Optimal tradeoff between performance and security in networked control systems based on coevolutionary algorithms," *IEEE Transactions on Industrial Electronics*, vol. 59, no. 7, pp. 3016–3025, 2012.

[14] S. Cui, Z. Han, S. Kar, T. T. Kim, H. V. Poor, and A. Tajer, "Coordinated data-injection attack and detection in the smart grid: A detailed look at enriching detection solutions," *Signal Processing Magazine, IEEE*, vol. 29, no. 5, pp. 106–115, 2012.

[15] F. Pasqualetti, F. Dorfler, and F. Bullo, "Attack detection and identification in cyber-physical systems," *Automatic Control, IEEE Transactions on*, vol. 58, no. 11, pp. 2715–2729, Nov 2013.

[16] H. Fawzi, P. Tabuada, and S. Diggavi, "Secure estimation and control for cyber-physical systems under adversarial attacks," *Automatic Control, IEEE Transactions on*, vol. 59, no. 6, pp. 1454–1467, June 2014.

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

[18] D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Drone hack: Spoofing attack demonstration on a civilian unmanned aerial vehicle," *GPS*

World, Aug. 2012.

[19] A. J. Kerns, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "Unmanned aircraft capture and control via GPS spoofing," *Journal of Field Robotics*, vol. 31, no. 4, pp. 617–636, 2014.

[20] D. Shepard, T. E. Humphreys, and A. Fansler, "Evaulation of the vulnerability of Phasor Measurement Units to GPS spoofing," in *Sixth Annual IFIP WG 11.10 International Conference on Critical Infrastructure Protection*, Washington, DC, Mar. 2012.

[21] A. H. Rutkin, "'Spoofers' use fake GPS signals to knock a yacht off course," *MIT Technology Review*, Aug. 2013.

[22] National PNT Advisory Board, "Jamming the Global Positioning System - A national security threat: Recent events and potential cures," Nov. 2010.

[23] R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. Humphreys, "Signal characteristics of civil GPS jammers," in *Proceedings of the ION GNSS Meeting*, 2011.

[24] S. Pullen and G. Gao, "GNSS jamming in the name of privacy," *Inside GNSS*, vol. 7, no. 2, Mar./Apr. 2012.

[25] J. C. Grabowski, "Personal Privacy Jammers: Locating Jersey PPDs Jamming GBAS Safety-of-Life Signals," *GPS World*, vol. 23, no. 4, pp. 28–37, April 2012.

[26] Department of Homeland Security official, private communication, Sept. 2011.

[27] Federal Communications Commission, "Public Notice DA-05-1776," June 2005.

[28] T. E. Humphreys, "The GPS dot and its discontents: Privacy vs. GNSS integrity," *Inside GNSS*, vol. 7, no. 2, Mar./Apr. 2012.

[29] L. Scott, "J911: Fast Jammer Detection," *GPS World*, vol. 21, no. 11, pp. 32–37, 2010.

[30] A. Brown, D. Reynolds, D. Roberts, and S. Serie, "Jammer and interference location system," in *Proceedings of the ION GPS Meeting*. Nashville, TN: Institute of Navigation, Sept. 1999, pp. 137–142.

[31] A. Proctor, C. Curry, J. Tong, R. Watson, M. Greaves, and P. Cruddace, "Protecting the UK infrastructure," *Inside GNSS*, vol. 6, no. 5, Sep./Oct. 2011.

[32] O. Isoz, A. T. Balaei, and D. Akos, "Interference detection and localization in the GPS L1 band," in *Proceedings of the ION International Technical Meeting*. San Diego, CA: Institute of Navigation, Jan. 2010, pp. 925–929.

[33] J. Lindstrom, D. M. Akos, O. Isoz, and M. Junered, "GNSS interference detection and localization using a network of low-cost front-end modules," in *Proceedings of the ION GNSS Meeting*. Institute of Navigation, 2007.

[34] A. Weiss, "Direct geolocation of wideband emitters based on delay and doppler," *Signal Processing, IEEE Transactions on*, vol. 59, no. 6, pp. 2513–2521, June 2011.

[35] A. Sidi and A. Weiss, "Delay and doppler induced direct tracking by particle filter," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 50, no. 1, pp. 559–572, January 2014.

[36] K. B. Deshpande, G. S. Bust, C. R. Clauer, H. Kim, J. E. Macon, T. E. Humphreys, J. A. Bhatti, S. B. Musko, G. Crowley, and A. T. Weatherwax, "Initial GPS scintillation results from CASES receiver at South Pole, Antarctica," *Radio Science*, vol. 47, no. 5, 2012.

[37] C. R. Clauer, H. Kim, K. Deshpande, Z. Xu, D. Weimer, S. Musko, G. Crowley, C. Fish, R. Nealy, T. E. Humphreys, J. A. Bhatti, and A. J. Ridley, "Autonomous adaptive low-power instrument platform (AAL-PIP) for remote high latitude geospace data collection," *Geoscientific Instrumentation, Methods and Data Systems*, vol. 3, pp. 211–227, 2014.

[38] M. Psiaki, B. O'Hanlon, J. Bhatti, D. Shepard, and T. Humphreys, "GPS spoofing detection via dual-receiver correlation of military signals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2250–2267, 2013.

[39] B. W. O'Hanlon, M. L. Psiaki, T. E. Humphreys, J. A. Bhatti, and D. P. Shepard, "Real-time GPS spoofing detection via correlation of encrypted

signals," *Navigation, Journal of the Institute of Navigation*, vol. 60, no. 4, pp. 267–278, 2013.

[40] E. G. Lightsey, T. E. Humphreys, J. A. Bhatti, A. J. Joplin, B. W. O'Hanlon, and S. P. Powell, "Demonstration of a space capable miniature dual frequency GNSS receiver," *Navigation, Journal of the Institute of Navigation*, vol. 61, no. 1, pp. 53–64, 2014.

[41] J. Bhatti and T. Humphreys, "Hostile control of surface vessels via counterfeit GPS signals: Demonstration and detection," *Navigation, Journal of the Institute of Navigation*, 2015, (In preparation.).

[42] J. Bhatti, B. Ledvina, and T. Humphreys, "Analysis and experimental results of direct geolocation techniques," *Navigation, Journal of the Institute of Navigation*, 2015, (In preparation.).

[43] T. E. Humphreys, J. Bhatti, T. Pany, B. Ledvina, and B. O'Hanlon, "Exploiting multicore technology in software-defined GNSS receivers," in *Proceedings of the ION GNSS Meeting*. Savannah, GA: Institute of Navigation, 2009, pp. 326–338.

[44] J. A. Bhatti, T. E. Humphreys, and B. M. Ledvina, "Development and demonstration of a TDOA-based GNSS interference signal localization system," in *Proceedings of the IEEE/ION PLANS Meeting*, April 2012, pp. 1209–1220.

[45] T. E. Humphreys, J. Bhatti, and B. M. Ledvina, "The GPS Assimilator: Upgrading receivers via benign spoofing," *Inside GNSS*, vol. 5, no. 4, pp.

50–58, June 2010.

[46] R. Mitch, R. Dougherty, M. Psiaki, S. Powell, B. O'Hanlon, J. Bhatti, and T. E. Humphreys, "Know your enemy: Signal characteristics of civil GPS jammers," *GPS World*, Jan. 2012.

[47] National Transportation Safety Board, "Marine accident report: Grounding of the Panamanian passenger ship Royal Majesty on Rose and Crown Shoal near Nantucket, Massachusetts June 10, 1995," National Transportation Safety Board, Tech. Rep., 1997.

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

[49] A. Grant, "GPS jamming and the impact on maritime navigation," *Journal of Navigation*, vol. 62, no. 2, 2009.

[50] International Marine Contractors Association, "Guidelines for the design and operation of dynamically positioned vessels," 2007, http://www.imca-int.com/media/73055/imcam103.pdf.

[51] M. Thomas, J. Norton, A. Jones, A. Hopper, N. Ward, P. Cannon, N. Ackroyd, P. Cruddace, and M. Unwin, "Global navigation space systems: reliance and vulnerabilities," *The Royal Academy of Engineering, London*, 2011.

[52] M. Caccia, M. Bibuli, R. Bono, and G. Bruzzone, "Basic navigation, guidance and control of an unmanned surface vehicle," *Autonomous Robots*, vol. 25, no. 4, pp. 349–365, 2008.

[53] L. Elkins, D. Sellers, and W. R. Monach, "The autonomous maritime navigation (AMN) project: Field tests, autonomous and cooperative behaviors, data fusion, sensors, and vehicles," *Journal of Field Robotics*, vol. 27, no. 6, pp. 790–818, 2010.

[54] L. Paull, S. Saeedi, M. Seto, and H. Li, "AUV navigation and localization: A review," *Oceanic Engineering, IEEE Journal of*, vol. 39, no. 1, pp. 131–149, 2014.

[55] D. K. Meduna, S. M. Rock, and R. S. McEwen, "Closed-loop terrain relative navigation for AUVs with non-inertial grade navigation sensors," in *Autonomous Underwater Vehicles (AUV), 2010 IEEE/OES*. IEEE, 2010, pp. 1–8.

[56] D. Meduna, S. M. Rock, and R. McEwen, "AUV terrain relative navigation using coarse maps," in *Unmanned Untethered Submersible Technology Conference*, 2009.

[57] U.S. Coast Guard; U.S. Department of Homeland Security, "Terminate long range aids to navigation (Loran-C) signal," Federal Register, Jan. 2010.

[58] M. Narins, M. Lombardi, P. Enge, B. Peterson, S. Lo, Y. H. Chen, and D. Akos, "The need for a robust precise time and frequency alternative

to global navigation satellite systems," *Journal of Air Traffic Control*, vol. 55, no. 1, 2012.

[59] D. P. Shepard, T. E. Humphreys, and A. A. Fansler, "Evaluation of the vulnerability of phasor measurement units to GPS spoofing attacks," *International Journal of Critical Infrastructure Protection*, vol. 5, no. 3-4, pp. 146–153, 2012.

[60] GPS Directorate, "Systems engineering and integration Interface Specification IS-GPS-200G," 2012, http://www.gps.gov/technical/icwg/.

[61] European Union, "European GNSS (Galileo) open service signal in space interface control document," 2010, http://ec.europa.eu/enterprise/policies/satnav/galileo/open-service/.

[62] A. S. Willsky, "A survey of design methods for failure detection in dynamic systems," *Automatica*, vol. 12, no. 6, pp. 601–611, 1976.

[63] M. Basseville, "Detecting changes in signals and systems—a survey," *Automatica*, vol. 24, no. 3, pp. 309–326, 1988.

[64] P. M. Frank, "Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy: A survey and some new results," *Automatica*, vol. 26, no. 3, pp. 459–474, 1990.

[65] J. Chen and R. J. Patton, *Robust model-based fault diagnosis for dynamic systems*. Springer Publishing Company, Incorporated, 2012.

[66] M. Joerger and B. Pervan, "Kalman filter-based integrity monitoring against sensor faults," *Journal of Guidance, Control, and Dynamics*, vol. 36, pp. 349–361, 2013.

[67] S. Khanafseh, N. Roshan, S. Langel, F. Cheng-Chan, M. Joerger, and B. Pervan, "GPS spoofing detection using RAIM with INS coupling," in *Proceedings of the IEEE/ION PLANS Meeting*, May 2014.

[68] A. Wald, "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, June 1945.

[69] H. L. V. Trees, *Detection, Estimation, and Modulation Theory*. Wiley, 2001.

[70] R. K. Mehra and J. Peschon, "An innovations approach to fault detection and diagnosis in dynamic systems," *Automatica*, vol. 7, no. 5, pp. 637–640, 1971.

[71] L. Pelkowitz and S. Schwarts, "Asymptotically optimum sample size for quickest detection," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-23, no. 2, pp. 263–272, March 1987.

[72] N. Bowditch, *The American Practical Navigator*. Bethesda, Maryland: National Imagery and Mapping Agency, 2002.

[73] GPS World staff, "Hemisphere GPS Offers Vector Compass Products for Marine Applications," *GPS World*, Oct. 2012, http://gpsworld.com/hemisphere-gps-offers-vector-compass-products-for-marine-applications.

[74] *Radar Navigation and Maneuvering Board Manual*, 7th ed., National Imagery and Mapping Agency, Bethesda, Maryland, 2001, http://msi.nga.mil/MSISiteContent/StaticFiles/NAV_PUBS/RNM/310ch5.pdf.

[75] A. Norris, *ECDIS and Positioning*, ser. Integrated bridge systems. Nautical Institute, 2010.

[76] eNav International, "Totem ECDIS and GPS Spoofing," June 2013, http://www.enav-international.com/news/id5774-Totem_ECDIS_and_GPS_Spoofing.html.

[77] M. H. Lützhöft and S. W. Dekker, "On your watch: Automation on the bridge," *Journal of Navigation*, vol. 55, no. 1, pp. 83–96, 2002.

[78] F. Kendoul, "Survey of advances in guidance, navigation, and control of unmanned rotorcraft systems," *Journal of Field Robotics*, vol. 29, no. 2, pp. 315–378, 2012.

[79] R. D. Luce and H. Raiffa, *Games and Decisions: Introduction and Critical Survey*. Dover, 1989.

[80] D. A. Blackwell, *Theory of games and statistical decisions*. Courier Dover Publications, 1979.

[81] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley and Sons, 2001.

[82] B. M. Ledvina, W. J. Bencze, B. Galusha, and I. Miller, "An in-line anti-spoofing module for legacy civil GPS receivers," in *Proceedings of the ION International Technical Meeting*, San Diego, CA, Jan. 2010.

[83] K. D. Wesson, D. P. Shepard, J. A. Bhatti, and T. E. Humphreys, "An evaluation of the vestigial signal defense for civil GPS anti-spoofing," in *Proceedings of the ION GNSS Meeting*, Portland, OR, 2011.

[84] V. Dehghanian, J. Nielsen, and G. Lachapelle, "GNSS spoofing detection based on receiver C/N_0 estimates," in *Proceedings of the ION GNSS Meeting*. Nashville, Tennessee: Institute of Navigation, 2012.

[85] K. D. Wesson, B. L. Evans, and T. E. Humphreys, "A combined symmetric difference and power monitoring GNSS anti-spoofing technique," in *IEEE Global Conference on Signal and Information Processing*, 2013.

[86] D. S. D. Lorenzo, J. Gautier, J. Rife, P. Enge, and D. Akos, "Adaptive array processing for GPS interference rejection," in *Proceedings of the ION GNSS Meeting*. Long Beach, CA: Institute of Navigation, Sept. 2005.

[87] P. Y. Montgomery, T. E. Humphreys, and B. M. Ledvina, "A multi-antenna defense: Receiver-autonomous GPS spoofing detection," *Inside GNSS*, vol. 4, no. 2, pp. 40–46, April 2009.

[88] A. Broumandan, A. Jafarnia-Jahromi, V. Dehgahanian, J. Nielsen, and G. Lachapelle, "GNSS spoofing detection in handheld receivers based

on signal spatial correlation," in *Proceedings of the IEEE/ION PLANS Meeting*. Myrtle Beach, SC: Institute of Navigation, April 2012.

[89] K. D. Wesson, M. P. Rothlisberger, and T. E. Humphreys, "Practical cryptographic civil GPS signal authentication," *Navigation, Journal of the Institute of Navigation*, vol. 59, no. 3, pp. 177–193, 2012.

[90] T. E. Humphreys, "Detection strategy for cryptographic GNSS anti-spoofing," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 2, pp. 1073–1090, 2013.

[91] S. Lo, D. DeLorenzo, P. Enge, D. Akos, and P. Bradley, "Signal authentication," *Inside GNSS*, vol. 0, no. 0, pp. 30–39, Sept. 2009.

[92] M. L. Psiaki, B. W. O'Hanlon, J. A. Bhatti, and T. E. Humphreys, "Civilian GPS spoofing detection based on dual-receiver correlation of military signals," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.

[93] B. O'Hanlon, M. Psiaki, J. Bhatti, and T. Humphreys, "Real-time spoofing detection using correlation between two civil GPS receiver," in *Proceedings of the ION GNSS Meeting*. Nashville, Tennessee: Institute of Navigation, 2012.

[94] M. L. Psiaki, B. W. O'Hanlon, S. P. Powell, J. A. Bhatti, K. D. Wesson, T. E. Humphreys, and A. Schofield, "GNSS spoofing detection using two-antenna differential carrier phase," in *Proceedings of the ION GNSS+ Meeting*. Tampa, FL: Institute of Navigation, 2014.

[95] K. Becker, "Passive localization of frequency-agile radars from angle and frequency measurements," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 35, no. 4, pp. 1129–1144, Oct 1999.

[96] D. Musicki, R. Kaune, and W. Koch, "Mobile emitter geolocation and tracking using tdoa and fdoa measurements," *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1863–1874, March 2010.

[97] A. Mikhalev, E. Hughes, and R. Ormondroyd, "Comparison of hough transform and particle filter methods of passive emitter geolocation using fusion of tdoa and aoa data," in *Information Fusion (FUSION), 2010 13th Conference on*, July 2010, pp. 1–8.

[98] R. Thompson, E. Cetin, and A. Dempster, "Unknown source localization using rss in open areas in the presence of ground reflections," in *Position Location and Navigation Symposium (PLANS), 2012 IEEE/ION*, April 2012, pp. 1018–1027.

[99] K. G. Gromov, "GIDL: Generalized interference detection and localization system," Ph.D. dissertation, Stanford University, March 2002.

[100] N. El Gemayel, S. Koslowski, F. Jondral, and J. Tschan, "A low cost tdoa localization system: Setup, challenges and results," in *Positioning Navigation and Communication (WPNC), 2013 10th Workshop on*, March 2013, pp. 1–4.

[101] S. Stein, "Algorithms for ambiguity function processing," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 29, no. 3, pp.

588–599, June 1981.

[102] C. Knapp and G. Carter, "The generalized correlation method for estimation of time delay," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 24, no. 4, pp. 320–327, 1976.

[103] Y. Chan and K. Ho, "A simple and efficient estimator for hyperbolic location," *IEEE Transactions on Signal Processing*, vol. 42, no. 8, pp. 1905–1915, Aug. 1994.

[104] K. Ho and W. Xu, "An accurate algebraic solution for moving source location using tdoa and fdoa measurements," *Signal Processing, IEEE Transactions on*, vol. 52, no. 9, pp. 2453–2463, Sept 2004.

[105] A. Mikhalev and R. Ormondroyd, "Comparison of hough transform and particle filter methods of emitter geolocation using fusion of tdoa data," in *Positioning, Navigation and Communication, 2007. WPNC '07. 4th Workshop on*, March 2007, pp. 121–127.

[106] A. Rihaczek, *Principles of high-resolution radar*. McGraw-Hill, 1969.

[107] M. Psiaki and S. Mohiuddin, "Modeling, analysis, and simulation of GPS carrier phase for spacecraft relative navigation," *Journal of Guidance Control and Dynamics*, vol. 30, no. 6, p. 1628, 2007.

[108] P. Misra and P. Enge, *Global Positioning System: Signals, Measurements, and Performance*, revised second ed. Lincoln, Massachusetts: Ganga-Jumana Press, 2012.

[109] S. Stein, "Differential delay/doppler ml estimation with unknown signals," *Signal Processing, IEEE Transactions on*, vol. 41, no. 8, pp. 2717–2719, Aug 1993.

[110] W. Bajwa, K. Gedalyahu, and Y. Eldar, "Identification of parametric underspread linear systems and super-resolution radar," *IEEE Transactions on Signal Processing*, vol. 59, no. 6, pp. 2548–2561, June 2011.

[111] K. M. Pesyna, Jr., K. D. Wesson, R. W. Heath, Jr., and T. E. Humphreys, "Extending the reach of GPS-assisted femtocell synchronization and localization through tightly-coupled opportunistic navigation," in *IEEE GLOBECOM Workshop*, 2011.

[112] K. M. Pesyna Jr., Z. M. Kassas, J. A. Bhatti, and T. E. Humphreys, "Tightly-coupled opportunistic navigation for deep urban and indoor positioning," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2011.

[113] K. D. Wesson, K. M. Pesyna, Jr., J. A. Bhatti, and T. E. Humphreys, "Opportunistic frequency stability transfer for extending the coherence time of GNSS receiver clocks," in *Proceedings of the ION GNSS Meeting*. Portland, Oregon: Institute of Navigation, 2010.

[114] M. Richards, "Coherent integration loss due to white gaussian phase noise," *Signal Processing Letters, IEEE*, vol. 10, no. 7, pp. 208–210, July 2003.

[115] A. Weiss and A. Amar, "Direct geolocation of stationary wideband radio signal based on time delays and doppler shifts," in *Statistical Signal Processing, 2009. SSP '09. IEEE/SP 15th Workshop on*, Aug 2009, pp. 101–104.

[116] F. Gustafsson, "Particle filter theory and practice with positioning applications," *Aerospace and Electronic Systems Magazine, IEEE*, vol. 25, no. 7, pp. 53–82, July 2010.

[117] P. Teunissen, P. De Jonge, and C. Tiberius, "The LAMBDA method for fast GPS surveying," in *Proceedings of International Symposium on GPS Technology Applications*, vol. 29. Bucharest, Romania: Union of Romanian Geodesy, Sept. 1995, pp. 203–210.

[118] D. Torrieri, "Statistical theory of passive location systems," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. AES-20, no. 2, pp. 183–198, March 1984.

[119] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," *Radar and Signal Processing, IEE Proceedings F*, vol. 140, no. 2, pp. 107–113, Apr 1993.

[120] G. Hoffmann and C. Tomlin, "Mobile sensor network control using mutual information methods and particle filters," *Automatic Control, IEEE Transactions on*, vol. 55, no. 1, pp. 32–47, Jan 2010.

[121] R. G. Brown, *Global Positioning System: Theory and Applications*. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1996,

vol. 2, ch. 5: Receiver Autonomous Integrity Monitoring, pp. 143–168.

Vita

Jahshan A. Bhatti is pursuing a Ph.D. in the Department of Aerospace Engineering and Engineering Mechanics at the University of Texas at Austin, where he also received his M.S. in 2011 and B.S. in 2009. He grew up in McAllen, TX, where he graduated from Nikki Rowe High School in 2006. He is a member of the UT Radionavigation Laboratory. His research interests are in the development of small satellites, software-defined radio applications, space weather, and GNSS security and integrity.

Permanent address: jahshan@utexas.edu

This dissertation was typeset with LaTeX by the author.