

**Decentralized Collaborative Navigation in
Limited Observability Environments with
Low Earth Orbit Satellite Signals of Opportunity between
Aerial and Ground Vehicles**

by

Christian D. Moomaw

A thesis submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Auburn, Alabama
August 5, 2023

Keywords: signals of opportunity, collaborative navigation, data fusion

Copyright 2023 by Christian D. Moomaw

Approved by

Scott M. Martin, Chair, Assistant Professor of Mechanical Engineering
David M. Bevly, McNair Distinguished Professor of Mechanical Engineering
Chad G. Rose, Assistant Professor of Mechanical Engineering

Approved for public release; distribution unlimited.
Public release case number 23-2667.

Abstract

A decentralized collaborative navigation algorithm known as inverse covariance intersection (ICI) is studied in the context of a group of vehicles navigating using opportunistic Doppler measurements. Signals of opportunity (SOOPs) have been extensively studied for applications requiring reliable position, velocity, and timing (PVT) information in conditions with potentially degraded GNSS performance. Doppler measurements derived from SOOPs can be used for positioning when GNSS signals are unavailable, but the resulting position estimate accuracy from Doppler-only techniques is unacceptably poor for many use cases. Collaborative techniques can leverage high-quality peer-to-peer range measurements to constrain PVT estimate error growth for each vehicle in a collaborating group.

A navigator employing a tightly-coupled Doppler-inertial extended Kalman filter (EKF) is developed. Its performance is analyzed using Monte Carlo techniques and simulated Doppler measurements from a collection of satellites in low earth orbit (LEO). Peer-to-peer range measurements are integrated using techniques including the well-known covariance intersection (CI), ICI, and a centralized EKF. The performance gains of each method are presented as compared to non-cooperating vehicles. Additionally, the two decentralized navigators are each compared to the centralized navigator, which represents a reasonable best case. The proposed ICI-based navigator is shown using a Monte Carlo test to achieve 62% of the position error reduction of an ideal centralized navigator in the average case, compared to 28% for the well-studied CI-based technique. The proposed ICI navigator is tested with experimentally-collected ranges from ultrawideband transceivers and is shown remain functional in the presence of faulty range measurements.

Acknowledgments

To my family, for being a reliable and loving foundation,

To Zachary and William, for weathering the last year or two of neglected friendship,

To Dr. Martin, for guidance,

To Dr. Bevly, for opportunity,

To Dr. Rose, for insight and encouragement,

To Ben, for friendly and frank scrutiny,

To Dan, for getting the UAV test off the ground,

*To Sam, Thomas, William, Madeline, J.B., Sam, Walt, Kathleen, and Gabe, for creating a
sense of workplace camaraderie that I may never experience again,*

To Dr. Burch, for forthright mentorship,

To Dr. Ragan, for showing me the meaning of service to others,

*To those whose work I build upon, for providing an interesting and challenging
starting point,*

To everyone in my life, for making me who I am today,

To Elizabeth, for the lifetime we have to look forward to together,

I owe a great debt of gratitude.

Table of Contents

Abstract	ii
Acknowledgments	iii
1 Introduction	1
1.1 Prior Art	1
1.2 Contributions of This Work	3
1.3 Overview of Subsequent Chapters	3
2 Signals of Opportunity	5
2.1 Opportunistic Navigation	5
2.2 Useful SOOPs	7
3 Doppler Navigation	10
3.1 The Doppler Effect	10
3.2 Historical Doppler Navigation	12
3.3 Application to an EKF Navigator	12
3.4 Navigator Implementation	17
4 LEO Performance Heuristics	19
4.1 Overhead Satellite Counts	20
4.2 Doppler Geometry Index	24

5	Solo Navigator Performance Study	30
5.1	Simulation Framework	30
5.2	Monte Carlo Testing	33
5.3	Test Scenario	34
5.4	Test Results	35
6	Collaborative Navigation	42
6.1	Peer-to-Peer Ranging	42
6.2	Data Fusion	44
6.2.1	Centralized Collaboration	45
6.2.2	Decentralized Collaboration with CI	47
6.2.3	Decentralized Collaboration with ICI	50
6.2.4	Adapting the Navigator for ICI	51
7	Collaborative Navigator Performance Study	58
8	Test Track Experiments	63
8.1	Modifications to the Navigators	63
8.2	Test Scenario	69
8.3	Results	74
8.3.1	With SOOPs	74
8.3.2	Without SOOPs	76
8.3.3	Simulated UWB Measurements	77
9	Conclusions and Future Work	83
9.1	Summary	83

9.2	Opportunistic Doppler Navigation	84
9.3	Collaborative Navigation Simulation Results	84
9.4	Collaborative Navigation Experimental Results	85
9.5	Future Work	86
	References	88
	Appendices	97
A	Additional Monte Carlo Position Error Results	98
B	Additional Test Track Results	109
	B.1 Circling Trajectory Results	109
	B.2 Additional Data	109

List of Figures

4.1	Expected overhead satellite map for Iridium.	21
4.2	Expected overhead satellite map for Globalstar.	21
4.3	Expected overhead satellite map for ORBCOMM.	21
4.4	Expected total overhead satellite map for the three LEO constellations.	22
4.5	Expected number of overhead Iridium satellites at varying latitude.	23
4.6	Expected number of overhead Globalstar satellites at varying latitude.	23
4.7	Expected number of overhead ORBCOMM satellites at varying latitude.	23
4.8	Expected total number of overhead satellites at varying latitude.	24
4.9	Evolution of batched DGI over time for various satellite geometries.	28
4.10	Latitude-dependent DGI characteristics with moderate batching times.	29
4.11	Latitude-dependent DGI characteristics with long batching times.	29
5.1	High-level simulation components and associated data flow.	31
5.2	The test trajectory consisting of four UAVs.	35
5.3	Position error for four non-collaborating navigators at 30° N.	36
5.4	Position error magnitude for four non-collaborating navigators at the equator.	38
5.5	Position error magnitude for four non-collaborating navigators at 30° N.	38
5.6	Position error magnitude for four non-collaborating navigators at 60° N.	39
5.7	Position error magnitude for four non-collaborating navigators without SOOP measurements.	39
5.8	Summarized navigator positioning error magnitude.	40

5.9	Comparison of time-to-DGI values and Monte Carlo 3σ position errors. . . .	41
6.1	An informal illustration of the CI algorithm.	48
6.2	An informal illustration of the ICI algorithm.	51
7.1	Comparison of collaborative positioning techniques at the equator.	60
7.2	Comparison of collaborative positioning techniques at 30° N.	60
7.3	Comparison of collaborative positioning techniques at 60° N.	61
7.4	Comparison of collaborative positioning techniques with no LEO measurements.	61
8.1	MKZ roof-mounted test equipment.	70
8.2	Kia roof-mounted test equipment.	70
8.3	UAV-mounted test equipment.	71
8.4	Test vehicles during final checkouts before testing.	71
8.5	View from the driver's seat of the Kia before starting a test run.	72
8.6	Ground track of test run with a stationary hovering UAV.	73
8.7	Ground track of test run with a circling UAV.	73
8.8	Position error magnitude with no collaboration.	75
8.9	Overall position error magnitude for various collaborative methods.	75
8.10	UWB measurement error.	76
8.11	MKZ dead reckoned ground track.	79
8.12	Kia dead reckoned ground track.	79
8.13	UAV dead reckoned ground track.	79
8.14	Position error magnitude with no collaboration and no SOOP measurements.	80
8.15	Overall position error magnitude with no SOOP measurements.	80
8.16	Position error magnitude with CI and no SOOP measurements.	81
8.17	Position error magnitude with ICI and no SOOP measurements.	81

8.18 Overall position error magnitude with simulated UWB ranges.	82
8.19 Overall position error magnitude with simulated UWB ranges and no SOOP measurements.	82
A.1 Position error for a navigator using GPS and no collaboration.	99
A.2 Position error for a navigator using GPS and CI.	99
A.3 Position error for a navigator using GPS and ICI.	100
A.4 Position error for a navigator using GPS and centralized collaboration. . . .	100
A.5 Position error for a navigator using IMU and no collaboration.	101
A.6 Position error for a navigator using IMU and CI.	101
A.7 Position error for a navigator using IMU and ICI.	102
A.8 Position error for a navigator using IMU and centralized collaboration. . . .	102
A.9 Position error for a navigator using LEO at the equator and no collaboration.	103
A.10 Position error for a navigator using LEO at the equator and CI.	103
A.11 Position error for a navigator using LEO at the equator and ICI.	104
A.12 Position error for a navigator using LEO at the equator and centralized col- laboration.	104
A.13 Position error for a navigator using LEO at 30° N and no collaboration. . . .	105
A.14 Position error for a navigator using LEO at 30° N and CI.	105
A.15 Position error for a navigator using LEO at 30° N and ICI.	106
A.16 Position error for a navigator using LEO at 30° N and centralized collaboration.	106
A.17 Position error for a navigator using LEO at 60° N and no collaboration. . . .	107
A.18 Position error for a navigator using LEO at 60° N and CI.	107
A.19 Position error for a navigator using LEO at 60° N and ICI.	108
A.20 Position error for a navigator using LEO at 60° N and centralized collaboration.	108
B.1 Position error magnitude with no collaboration (see Fig. 8.8).	111

B.2	Overall position error magnitude for various collaborative methods (see Fig. 8.9).	111
B.3	UWB measurement error (see Fig. 8.10).	112
B.4	MKZ dead reckoned ground track (see Fig. 8.11).	112
B.5	Kia dead reckoned ground track (see Fig. 8.12).	113
B.6	UAV dead reckoned ground track (see Fig. 8.13).	113
B.7	Position error magnitude with no collaboration and no SOOP measurements (see Fig. 8.14).	114
B.8	Overall position error magnitude with no SOOP measurements (see Fig. 8.15).	114
B.9	Position error magnitude with CI and no SOOP measurements (see Fig. 8.16).	115
B.10	Position error magnitude with ICI and no SOOP measurements (see Fig. 8.17).	115
B.11	NED position errors with LEO SOOPs and no collaboration.	116
B.12	NED velocity errors with LEO SOOPs and no collaboration.	116
B.13	NED position errors with LEO SOOPs and CI.	117
B.14	NED velocity errors with LEO SOOPs and CI.	117
B.15	NED position errors with LEO SOOPs and ICI.	118
B.16	NED velocity errors with LEO SOOPs and ICI.	118
B.17	NED position errors with no SOOPs and no collaboration.	119
B.18	NED velocity errors with no SOOPs and no collaboration.	119
B.19	NED position errors with no SOOPs and CI.	120
B.20	NED velocity errors with no SOOPs and CI.	120
B.21	NED position errors with no SOOPs and ICI.	121
B.22	NED velocity errors with no SOOPs and ICI.	121
B.23	Removing a UWB unit from its mounting hardware.	122

List of Tables

3.1	Nominal H matrix positioning values for three emitter classes (scaled for easy comparison).	16
3.2	SOOP Doppler measurement noise parameters.	17
5.1	IMU model parameters.	32
7.1	Comparison of collaborative positioning accuracies.	62

Chapter 1

Introduction

Since their inception, global navigation satellite systems (GNSSs) such as GPS have demonstrated the importance of the global availability of high-quality positioning measurements. In the intervening decades, common causes of performance degradation of GNSSs such as foliage occlusion or jamming (among many others) have been extensively studied. More recently, signals of opportunity (SOOPs) have seen increasing study as a supplementary source of positioning information when GNSS is degraded or unavailable [1]. The quality of the measurements available from SOOP sources is, however, generally poorer than those from GNSSs. Consequently, the positioning performance resulting from the use of SOOP measurements is unacceptably poor for many use cases. Collaborative navigation techniques utilizing peer-to-peer range measurements among a group of users can enhance the viability of SOOPs as a GNSS fallback navigation technique, and thus is the focus of this work.

1.1 Prior Art

One of the earliest examples of SOOP-based positioning is presented in [2], which used carrier phase odometry from terrestrial AM radio stations to estimate the relative position between a static base receiver and a roving receiver. Research on SOOPs is relatively sparse in the literature from this point until around 2009, when works such as [3] examined the potential for using the Iridium satellite network to augment GPS performance. In the

years since there has been a steady stream of works published on SOOP navigation, with the Autonomous Systems Perception, Intelligence, & Navigation (ASPIN) Laboratory being particularly prolific publishers in the field [4–11]. In particular, [11] is pertinent as it demonstrates obtaining Doppler measurements opportunistically from the Iridium and ORBCOMM satellite networks, though only a stationary navigator is considered. Obtaining useful opportunistic measurements is also presented in [12] for the second-generation Iridium network, along with a discussion of dilution of precision metrics for satellite Doppler measurements. As for Globalstar (yet another satellite communications network), [13] describes some of the foundational components one would need to opportunistically track the network’s signals, but a working implementation of an opportunistic Globalstar receiver has yet to be demonstrated. A more general discussion of Doppler positioning can be found in [14], which focuses on using Doppler measurements derived from GPS.

An early example of collaborative navigation using peer-to-peer ranges is given in [15], which utilizes a centralized EKF and makes planar assumptions for the collaborating navigators. The development of the centralized navigator concept is further refined in [16], which includes with UAVs moving with a full six degrees of freedom. The use of covariance intersection (CI) for decentralized collaborative navigation is discussed extensively in [17, 18].

The inverse covariance intersection (ICI) algorithm which is of interest to this work is rigorously defined in [19, 20] and is implemented for measurement integrity monitoring in [21]. Interest in the ICI algorithm has increased in the literature during the past 2–3 years. At the time of writing, there are a number of works published within the past year which adapt and modify ICI to be better suited to distributed positioning applications. The algorithm is adapted to distributed sensor networks in [22–24], though in these cases the application is a tracking problem, in which a network of sensors is estimating the state of a common tracked object. Most recently — and most similarly to this work — [25] uses ICI to improve the positioning performance of UAVs, though it does so by implementing a LiDAR-based tracking system to estimate UAV positions.

Only one study was found in the literature applying ICI to a fully distributed network of independent navigators collaborating using peer-to-peer range measurements [26].

1.2 Contributions of This Work

The application of ICI to the problem of peer-to-peer cooperative navigation is not well studied in the literature. In particular, the integration of ICI with SOOP-based Doppler navigation has yet to be explored. The key contributions of this work, as well as the most relevant chapter for each, are as follows:

- Development of a mathematically simple and computationally efficient measure for describing the effect of emitter geometry on Doppler navigation performance (Chapter 4).
- A modified EKF-based navigator which permits the use of inverse covariance intersection for decentralized fusion of peer-to-peer range measurements (Chapter 6).
- Monte Carlo comparison of the proposed navigator to existing decentralized navigation techniques, and to a centralized collaborative navigator (Chapter 7).
- Integration of the proposed navigator with experimental peer-to-peer range measurements logged from ultrawideband (UWB) transceivers affixed to a set of aerial and ground vehicles (Chapter 8).

1.3 Overview of Subsequent Chapters

The remaining chapters in this work incrementally build the mathematical and computer models necessary to study a network of collaborating SOOP navigators. Chapter 2 introduces the concept of opportunistic navigation and provides rationale for the selection of signals used in this work. Chapter 3 demonstrates how Doppler measurements can be used for positioning and how the pattern of a tightly-coupled GNSS-INS EKF navigator can be modified to perform Doppler positioning. Chapter 4 develops a computationally efficient

performance metric for a Doppler-based navigator given a set of available measurements. Chapter 5 presents the simulation architecture used to test the navigator design of interest and presents Monte Carlo simulation results for non-cooperating Doppler navigators. Chapter 6 introduces fundamentals of peer-to-peer ranging and collaborative navigation, describing existing well-known data fusion methods as well as the novel ICI-centric architecture, with corresponding Monte Carlo simulation results given in Chapter 7. Chapter 8 applies the collaborative navigators to an experimentally collected data set. Chapter 9 summarizes the major findings of this work and outlines opportunities for further study.

Chapter 2

Signals of Opportunity

This chapter introduces the concept of signals of opportunity (SOOPs) and of opportunistic navigation more broadly. The difficulties inherent to SOOP navigation as compared to GNSSs are examined, and the criteria used to select the SOOPs of interest for this work are presented.

2.1 Opportunistic Navigation

In considering the broad collection of methods for measuring an object's position, the methods can be sorted into two broad categories: *coordinated navigation* and *opportunistic navigation*. In a coordinated navigation scheme, the navigator depends on the existence and functioning of some kind of infrastructure which was built for the purposes of providing navigation information. The nature of such navigation infrastructure can range in scale and complexity from the proximity Bluetooth beacons present in public venues like sporting arenas [27] to GNSSs like GPS and Galileo. In any case, the navigator leverages its *a priori* knowledge of the infrastructure's design to determine location. In contrast, opportunistic navigation is primarily characterized by a lack of intentionally deployed navigation infrastructure. This includes the opportunistic radio frequency (RF) navigation concepts that are a focus of this work, but ideas such as celestial navigation, geomagnetic navigation, and many vision-based approaches can also be described as opportunistic, since they only leverage preexisting information in the environment. The information used to opportunistically

determine position may be natural (stars) or artificial (communications satellites), but the common property is that the information being leveraged in the environment was not created for the express purpose of navigation.

Part of the appeal of opportunistic methods is their generally low cost of implementation. No infrastructure needs to be deployed, and in cases where non-navigation infrastructure (such as LEO satellite networks) is leveraged, no payment is made to the relevant network operator. In contrast, they are often inferior to coordinated navigation systems in terms of positioning accuracy. This is unsurprising, especially given that the purpose of designing and deploying navigation infrastructure for a coordinated navigation system is to create a known and predictable environment which is capable of yielding precise and accurate measurements. The deployment of GNSSs involves the tremendous expense of building and maintaining extensive networks of ground stations and satellites — and for good reason. A GNSS’s ground and space segments combine to create a precisely tuned and easily predictable RF environment for the navigator, which is a large part of how GNSSs can achieve excellent positioning performance on a planetary scale. Conversely, knowledge of the environment is a challenge for opportunistic navigation. Depending on the features which are of interest to the navigator, the environment may change in unexpected ways. The geomagnetic field shifts over time and maps used for vision navigation become outdated. Similarly for opportunistic radionavigation, the signals of interest can change behavior or fall out of service without notice.

Given such challenges, one might conclude that opportunistic navigation is not an area worthy of study. Indeed, a large proportion of the world’s navigation needs are adequately served by GNSSs and other coordinated systems. Where opportunistic techniques fall short in terms of predictability, they can excel in *availability*. GNSS performance degrades markedly without a clear view of the sky [28, 29] and fares poorly in the presence of interference [30]. Bluetooth beacons and similar technologies are only effective for confined areas. An opportunistic RF navigator can target signals in the environment which are available without

a clear view of the sky, or which have a stronger received signal strength to better resist interference. Similarly, vision-based approaches can work anywhere that is mappable, and celestial navigation can even be effective in deep space environments, far from *any* built infrastructure.

This chapter focuses on a specific type of opportunistic navigation known as *signal of opportunity* navigation. Signals of opportunity (SOOPs) can include any RF signal in the environment that is not intended for navigation. This excludes navigation signals such as GNSSs, eLORAN, VOR and DME, and the like, but includes almost everything else: terrestrial AM and FM radio, digital television, cellular network signals, weather station telemetry links, and satellite communication networks. Though it is relatively easy to enumerate a list of signals that *are* SOOPs, finding signals which are *useful* SOOPs is more challenging.

2.2 Useful SOOPs

For a particular SOOP to be of use to an opportunistic navigator, the navigator must be able to derive one or more *observables* from the signal. An observable is a quantity that is related to the states that the navigator is estimating (such as position or velocity, among others). Common observables include pseudorange, pseudorange rate, and angle of arrival, though others can be formulated. Determining what observables, if any, can be obtained from a particular SOOP, and (more importantly), establishing *how* to obtain those observables, is one of the largest barriers to effective implementation of a SOOP-based navigator. The signals of interest are often proprietary, with little publicly available documentation describing their structure. Empirical trial-and-error with live sky signal recordings are often the only way that a suitable signal tracking algorithm can be developed. Progress made in tracking SOOPs is vulnerable to regression, as any signal which is operated by a private company utilizing proprietary protocols is subject to change in structure suddenly and without notice.

The degree to which this is a risk worth considering varies greatly from one SOOP to another. Examining historical trends and considering the volume of consumer hardware using a given signal can be useful as a proxy for signal structure stability. Taken together, these considerations are often enough to narrow the field of candidate SOOPs from nearly every signal imaginable down to a small number of well-suited prospects.

Three signals form the basis of the work presented here: Iridium, Globalstar, and ORBCOMM. All three are low earth orbit (LEO) satellite communications networks. Iridium and Globalstar primarily offer satellite mobile phone services, while ORBCOMM offers non-realtime low-bandwidth data transmission services which are often used for collecting data from autonomous sensing equipment. These signals were selected for their broad availability across the earth (unlike terrestrial sources like cellular networks), their high received signal power relative to other satellite signals (such as those in geostationary orbit), and their generally stable and well-studied signal structures (unlike newer LEO constellations such as Starlink).

This work focuses on pseudorange rate observables derived from these signals. The most impactful advantage this choice offers over using pseudoranges is that Doppler shift is a simpler quantity to measure for an arbitrary signal than time of arrival (TOA). Users of Doppler measurements need not be concerned with clock biases of the transmitter and receiver, but rather they only need to account for the relative clock *drift*. Minimizing the impact of timing errors is especially helpful given that many non-GNSS satellites do not carry onboard atomic clocks. The major downside of using Doppler for navigation is that the pseudorange rate is not a strong function of position. Nonetheless it *is* a function of position, and thus it *can* be used for navigation. This topic is discussed in more detail in Chapter 3. A summary of rationale for selecting each signal is given below.

Iridium The Iridium network is perhaps the most widely studied space SOOP. Iridium’s LEO network boasts continuous coverage anywhere on earth, and the network has been

operational for long enough that the downlink signals are generally well understood [3, 12, 31, 32].

Globalstar Similarly to Iridium, Globalstar operates a global satellite phone network, but Globalstar's service footprint does not cover the entire earth. However, Globalstar does offer continuous coverage of at least one satellite overhead at any given time for the majority of the earth's land area. It is markedly less studied as a SOOP compared to Iridium, but it is included here as a plausible signal that could be leveraged in the future.

ORBCOMM Unlike the other two networks, ORBCOMM does not offer continuous satellite coverage; rather, their network is configured such that most locations rarely go longer than 10-20 minutes without a satellite pass. As such, ORBCOMM offers fewer measurements overall, but it has a well-understood signal structure and has been studied as a SOOP in the past [11, 33].

Chapter 3

Doppler Navigation

The use of Doppler measurements, and the pseudorange rate observables that they yield, is central to the collaborative navigation ideas presented in this work. This chapter introduces the concept of using Doppler measurements as a means of estimating position.

3.1 The Doppler Effect

The Doppler effect describes the signal propagation phenomenon by which relative motion between an emitter and a receiver causes a change in the apparent frequency of a signal as observed by the receiver. This property of propagating waves is responsible both for the characteristic *eeeeownn* sound of a passing car as heard by a pedestrian, as well as the redshifting of light from extremely distant stars as observed from earth [34]. The effect applies equally well to both mechanical and electromagnetic waves.

More formally, the Doppler effect means that the frequency of a received signal is shifted by an amount which is approximately proportional to the component of the relative velocity between the emitter and receiver which is along the line of sight. The Doppler frequency can be expressed as

$$\Delta f = -\frac{\dot{r}}{c}f_0 \tag{3.1}$$

where Δf is the change in observed frequency relative to transmission frequency f_0 , c is the wave speed, and \dot{r} is the range rate between the transmitter and the receiver. Range rate is given by

$$\dot{r} = (\mathbf{v}_{\text{RX}} - \mathbf{v}_{\text{TX}}) \cdot \mathbf{u}_{\text{TX} \rightarrow \text{RX}} \quad (3.2)$$

where \mathbf{v} describes the velocity of either the transmitter (TX) or receiver (RX) expressed in a common reference frame, and \mathbf{u} is a position unit vector, in this case pointing from the transmitter towards the receiver.

Doppler shift is influenced most strongly by the relative velocity terms. They are directly proportional, so a doubling in the relative velocity results in a doubling of the observed Doppler shift. This is what makes Doppler well-suited for estimating velocity. The idea here is to leverage Doppler measurements to estimate the *position* of a receiver (in addition to its velocity). To that end, attention must be paid to the unit vector term that projects the relative velocity onto the line of sight. For the usual purpose of estimating velocity, this unit vector is little more than a nuisance term, but in the absence of any other positioning information, it offers a foothold.

Unlike its relationship with velocity, Doppler shift is only a *weak* function of position. Consider an observer standing a few meters from a railroad track as a train passes by. Were the train to alter its speed, either speeding up or slowing down, the change in apparent frequency would be readily noticed by the observer, but if instead the observer moved closer to or further from the railroad tracks, the change would be much more subtle, requiring the observer to move tens or potentially hundreds of meters before noticing much of a change in observed frequency shifts. The mathematical manifestation of this idea is examined in Section 3.3.

Due to the tenuous connection between Doppler shift and position, determination of initial position is especially difficult with Doppler measurements. As such, this work focuses

on the use of Doppler navigation under *hot start* conditions, where the navigator starts with some inexact estimate of its position and velocity at time $t = 0$. Such a situation could arise if Doppler navigation is being used as a fallback to provide navigation information in the event that GNSS becomes unavailable.

3.2 Historical Doppler Navigation

Doppler navigation is certainly not a new concept. Doppler shift has been exploited for navigation purposes in coordinated navigation systems for decades. A notable example is the Transit satellite navigation system which, being introduced in 1963, predates the modern idea of a GNSS by a substantial margin. Transit relied on measurements taken across an entire satellite pass (hence its name) to determine a location estimate, and was targeted towards stationary and slow-moving receivers, such as ships and submarines [35, 36].

A more recent example of a deployed Doppler navigation system is the Sarsat distress beacon system, which in a sense turns the typical satellite navigation problem inside out. Instead of a receiver on the ground estimating its own position by measuring signals broadcast by satellites, the satellite system attempts to locate a distress beacon on the ground by examining the profile of the Doppler shift during a LEO satellite pass [37]. Despite the transposition of the transmitter and receiver roles, the overall working principle is generally the same for Sarsat as that used by Transit.

3.3 Application to an EKF Navigator

While several different navigator configurations are examined in this work, each remains, in essence, some variation on the well-known extended Kalman filter (EKF). As such, a rudimentary knowledge of EKFs as they pertain to pseudorange-based GNSS-INS navigation is assumed. A thorough grounding in the workings of GNSS-INS EKFs can be found in [38, 39]. This section focuses on how a GNSS-INS EKF may be modified to work with LEO Doppler

measurements, and the additional implementation and usage considerations introduced by such modifications.

The models discussed here suppose a tightly coupled GNSS-INS EKF as a starting point with an error state vector given by

$$\delta \mathbf{x} = \begin{bmatrix} \delta \mathbf{p} \\ \delta \mathbf{v} \\ \delta \Psi \\ \delta c\Delta t \\ \delta c\dot{\Delta} t \end{bmatrix} = \begin{bmatrix} \delta p_x \\ \delta p_y \\ \delta p_z \\ \delta v_x \\ \delta v_y \\ \delta v_z \\ \delta \varphi \\ \delta \theta \\ \delta \psi \\ \delta c\Delta t \\ \delta c\dot{\Delta} t \end{bmatrix} \quad (3.3)$$

referenced to the WGS84 earth-centered earth-fixed (ECEF) reference frame with position error $\delta \mathbf{p}$, velocity error $\delta \mathbf{v}$, attitude error $\delta \Psi$, clock bias error $\delta c\Delta t$, and clock drift error $\delta c\dot{\Delta} t$. The most substantial difference between using pseudoranges and using pseudorange rates is in the measurement model. Consider the measurement model for a pseudorange observable ρ in terms of transmitter (TX) and receiver (RX) positions \mathbf{p} , receiver clock bias $c\Delta t$, and unmodeled effects η_ρ :

$$\rho = |\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}| + c\Delta t + \eta_\rho. \quad (3.4)$$

In this case, η_ρ includes deterministic effects such as tropospheric and ionospheric delay, as well as stochastic effects like thermal noise. The model can be linearized about some nominal

state $\hat{\mathbf{x}}$ to obtain an EKF measurement matrix H , yielding

$$H_{\text{TOA}} = \left[\left(\frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \right)_{1 \times 3}^T \quad \mathbf{0}_{1 \times 6} \quad 1 \quad 0 \right] \quad (3.5)$$

which indicates that errors in pseudorange result in a combination of clock bias error and position error in the direction of the measurement line of sight. Note that the magnitudes of the position terms, being components of a unit vector, have an overall magnitude of 1, so the order of magnitude of pseudorange error and the corresponding position error are similar.

The same strategy can be applied to pseudorange rate observables. Beginning again with a full nonlinear measurement model results in

$$\dot{\rho} = (\mathbf{v}_{\text{RX}} - \mathbf{v}_{\text{TX}}) \cdot \frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} + c\dot{\Delta}t + \eta_{\dot{\rho}} \quad (3.6)$$

with $\eta_{\dot{\rho}}$ representing all unmodeled effects, the exact set of which differs between the pseudorange and pseudorange rate models. Deriving the linearized measurement model from this expression is quite cumbersome, and is covered in more detail in [12, 40, 41]. The resulting EKF measurement matrix is given by

$$H_{\text{Doppler}} = \left[\left(\frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \times \left(\frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \times \frac{\mathbf{v}_{\text{RX}} - \mathbf{v}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \right) \right)_{1 \times 3}^T \quad \left(\frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \right)_{1 \times 3}^T \quad \mathbf{0}_{1 \times 3} \quad 0 \quad 1 \right]. \quad (3.7)$$

The relationship between pseudorange rate and velocity is identical to that between pseudorange and position (i.e. the mapping onto the error state is the line-of-sight unit vector). This is unsurprising, given that the pseudorange rate and velocity are identically the derivatives of the pseudorange and position, respectively. The terms mapping pseudorange rate error onto position error are far more interesting.

Unlike the prior examples, the linearized mapping between measurement error and position error here is not easily lent to intuitive understanding. So, falling back to a more literal

mathematical understanding of the terms mapping position error to pseudorange rate error

$$H_{\text{Doppler},\mathbf{p}} = \frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \times \left(\frac{\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \times \frac{\mathbf{v}_{\text{RX}} - \mathbf{v}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \right) \quad (3.8)$$

it is important to note that the strength of the relationship between an estimated state and a measured quantity in the context of an EKF is generally proportional to the magnitude of the terms in the linearized measurement model H . Put another way, seeing larger values in the H matrix implies that the corresponding measurement provides more information about a given state than seeing smaller values. In the case of mapping pseudorange onto position, the magnitudes in question correspond to unit vector components, and thus tend to be on the order of 10^{-1} . So, even lacking a straightforward mental model for how pseudorange rate error maps onto position error, examining the expected magnitudes produced by (3.8) for various emitter types should be illuminating.

Closer inspection of (3.8) reveals that the first two fractional terms are identically the unit vector pointing from the receiver to the transmitter. Given that these unit vectors are applied to the remaining terms by a vector cross product, they do not have an effect on the overall magnitude of the result. Thus the remaining term, arbitrarily denoted ν ,

$$\nu = \frac{\mathbf{v}_{\text{RX}} - \mathbf{v}_{\text{TX}}}{|\mathbf{p}_{\text{RX}} - \mathbf{p}_{\text{TX}}|} \quad (3.9)$$

drives the magnitude. The term is the relative velocity between the transmitter and receiver divided by the scalar distance between the two. While still not being intuitively related to Doppler positioning, (3.9) is certainly easier to reason about than (3.8) involving nested vector cross products. For any given receiver implementing Doppler positioning, (3.9) shows that **fast-moving emitters are better than slow-moving ones** and that **nearby emitters are better than faraway ones**, or, equivalently, the suitability of an emitter for Doppler positioning is roughly proportional to the sweep rate of the line-of-sight unit vector.

To show this effect quantitatively, three emitter classes are considered for a receiver which is near the ground and traveling at 10 m/s: a nearby terrestrial emitter, a satellite in LEO, and a satellite in medium earth orbit (MEO). Consider a navigator located at 32.6° N, 85.5° W, 200 m altitude (height above ellipsoid) and traveling due east at 10 m/s. The terrestrial emitter is located 600 m north and 800 m east (1 km away) relative to the transmitter. Each of the two satellites is considered to be rising at 180° azimuth and 45° elevation. The LEO satellite is similar to a nominal Iridium satellite (780 km altitude, 90° inclination, 7456 m/s orbital speed) and the MEO satellite is similar to a nominal GPS satellite (20,200 km altitude, 55° inclination, 3874 m/s orbital speed). The resulting H matrix position terms for each scenario, given in Table 3.1, support the qualitative observations made earlier in this section. The terrestrial emitter, which is stationary but nearby, and the LEO satellite, which is moving quickly, offer much more position sensitivity as compared to the MEO satellite, which is extremely far away and moving substantially more slowly than the LEO satellite. The results shown also support the claim that even for well-suited emitters, Doppler shift is less strongly related to position than TOA. The largest values present in Table 3.1 are on the order of 10^{-3} , while values for TOA measurements tend to be larger, on the order of 10^{-1} .

Table 3.1: Nominal H matrix positioning values for three emitter classes (scaled for easy comparison).

(All values $\times 10^{-5}$)	H Position Terms	Norm
Terrestrial	[-379, 230., 404]	600.
LEO	[8.36, -118, 540.]	552
MEO	[10.4, -2.29, 14.1]	17.7

This method of comparing magnitudes of values in the measurement model can be useful for comparing the relative effectiveness of various emitter and measurement types, but the resulting values do not necessarily translate to a quantitative measurement of *how much better* one emitter or measurement type performs as compared to another. For example,

given that the H matrix values for TOA measurements from a LEO satellite are generally two orders of magnitude smaller than those for Doppler measurements, it is reasonable to expect that using the TOA measurements results in smaller position estimate errors than using the Doppler measurements. However, *it does not necessarily follow* that the positioning errors are two orders of magnitude smaller for TOA than for Doppler. This example only considers the affect of geometry on position estimation quality, but the nominal measurement error associated with each signal source and measurement type must also be considered. The simulated pseudorange rate observables used in this work are dirtied with white Gaussian noise with standard deviation values given in Table 3.2. The values are derived from a frequency locked loop (FLL) tracking noise model given in [42].

Table 3.2: SOOP Doppler measurement noise parameters.

Constellation	$\dot{\rho}$ Noise (m/s 1σ)
Iridium	0.04438
Globalstar	0.2736
ORBCOMM	0.1277

3.4 Navigator Implementation

This section summarizes the assumptions and implementation details for the remainder of the SOOP navigator. The navigator as discussed in this chapter does not include any mechanism for integrating collaborative range measurements. Modifications to the navigator to support collaborative navigation are presented in Chapter 6.

The EKF time updated is conducted using a 6 degree of freedom (DOF) inertial measurement unit (IMU) mechanization in the ECEF reference frame [38]. The corresponding

continuous process noise matrix Q was determined by hand tuning and is given by

$$Q = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 1} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{3 \times 3} & 0.01I_3 & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 1} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & 10 \times 10^{-9}I_3 & \mathbf{0}_{3 \times 1} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} & 0 & 0 \\ \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} & \mathbf{0}_{1 \times 3} & 0 & 10 \times 10^{-9} \end{bmatrix}. \quad (3.10)$$

Measurement noise model R_{Doppler} at each measurement epoch is constructed using the same noise standard deviations from Table 3.2 used to simulate the pseudorange rate observables. If multiple measurements are available at a single epoch (as would be the case if multiple satellites are simultaneously overhead), the measurement errors are assumed to be mutually uncorrelated (i.e. R_{Doppler} is always a diagonal matrix).

This chapter has shown that Doppler-based positioning schemes can be fitted into the widely used navigator format consisting of an inertial measurement unit (IMU) fused to RF measurements in an EKF without any structural accommodation other than a somewhat convoluted measurement model. Special care must be taken, however, in selecting well-suited signals and in adjusting positioning performance expectations as compared to TOA schemes such as GNSSs.

Chapter 4

LEO Performance Heuristics

The particular orbits used for satellites in a GNSS constellation are carefully selected to yield acceptable performance at any location on earth at any point in time. The differences in GNSS performance between two different places on earth due to satellite geometry tend to be small compared to other error sources (local obstructions such as urban canyons notwithstanding) [43]. This is not the case for the LEO constellations of interest in this work. Iridium, Globalstar, and ORBCOMM are all part of *communications* networks, not navigation systems, and their respective satellite orbits were selected with a different set of tradeoffs in mind compared to a GNSS. With Iridium, for example, there is always at least one satellite in view anywhere on the planet, but times and locations with two concurrent visible satellites are rare, and those with more than two concurrent visible satellites are rarer still. For navigation, having multiple concurrently visible emitters is an asset. For a communications network, it is a waste of resources at best and actively detrimental at worst due to crosstalk.

This chapter first examines how the expected number of visible satellites of interest changes from one location to the next. A more Doppler-focused measure is then developed, which is closely related to the dilution of precision (DOP) measure widely used with GNSSs. It should be noted that similarly to the discussion of Doppler H matrices in Section 3.3, the heuristic measures in this chapter are intended to compare the relative suitability of one

scenario to another, and are ill-suited to the determination of precise performance expectations.

4.1 Overhead Satellite Counts

A simple and straightforward proxy for the suitability of a particular satellite geometry to navigation is to count the number of visible emitters at a given location and time. In general, receiving signals from a greater number of satellites imparts a greater quantity of information to the receiver, thus resulting in better navigation performance.

Figures 4.1, 4.2, and 4.3 show the mean number of visible satellites above a 5° elevation mask for Iridium, Globalstar, and ORBCOMM, respectively. The number of overhead satellites is counted at each map point at 1-minute intervals over a 24-hour period. The resulting 1440 instantaneous satellite counts for each point are averaged to approximate the expected value. The values are calculated for a grid of points evenly spaced across 1810 latitude values and 3620 longitude values. Note that evenly spacing points across latitudes and longitudes does not result in an even spacing of points on the earth's surface, and the map resulting from such a sampling method is an equirectangular projection [44]. A similar map describing the total number of visible satellites across all three constellations is shown in Figure 4.4. The most important takeaway from these maps is that the mean number of overhead satellites does not change substantially as a function of longitude (equivalently, all color gradients in each map are exclusively in the latitude direction). This stands to reason given that the earth's rotation makes moving through time effectively the same as changing longitude from the perspective of a non-geosynchronous satellite. This observation regarding variation across latitude and longitude permits the same data shown in Figures 4.1–4.4 to be shown in a more easily interpreted form by examining a single line of longitude without loss of generality.




Figure 4.1: Expected overhead satellite map for Iridium.




Figure 4.2: Expected overhead satellite map for Globalstar.




Figure 4.3: Expected overhead satellite map for ORBCOMM.




Figure 4.4: Expected total overhead satellite map for the three LEO constellations.

Knowing that longitude is not a critical parameter, the same information shown in the color maps is presented in chart form for the line of longitude at 85.5° W in Figures 4.5, 4.6, and 4.7 — again for Iridium, Globalstar, and ORBCOMM, respectively. The reduction in spatial dimension permits easier visualization without the need for color gradients, and allows for the addition of minima and maxima. The result for Iridium corroborates the claim that the constellation does indeed have continuous visibility of at least one emitter at all locations. Further, the number of visible satellites increases monotonically with latitude, with locations outside of $\pm 60^\circ$ latitude having two or more satellites above the horizon at minimum, and more than three on average. The other two constellations, with their non-polar-inclined orbits, exhibit somewhat different behavior. They each have the highest average number of visible satellites in lower-mid latitudes — roughly $\pm 30^\circ$ for both Globalstar and ORBCOMM — and have *no* visible satellites at any time near the poles. Additionally, ORBCOMM has the curious property that no location has guaranteed satellite visibility, but many latitudes occasionally have four or five ORBCOMM satellites visible at one time. The data from these three plots are combined in Figure 4.8 to show the total number of overhead satellites as a function of latitude. The overall shape is generally similar to the behavior of Iridium, but the dropoff in visible satellites in both the mean and minimum cases is exaggerated near the equator compared to Iridium alone.

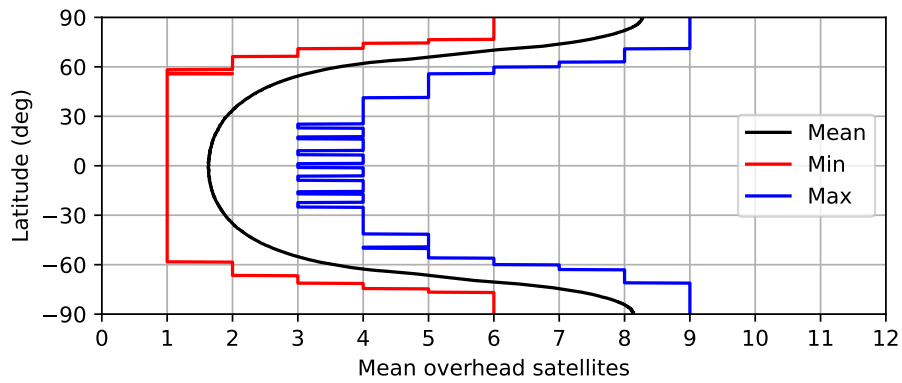


Figure 4.5: Expected number of overhead Iridium satellites at varying latitude.

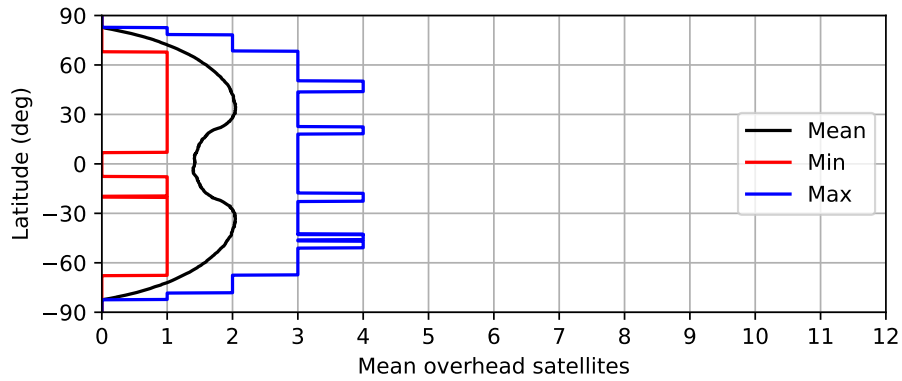


Figure 4.6: Expected number of overhead Globalstar satellites at varying latitude.




Figure 4.7: Expected number of overhead ORBCOMM satellites at varying latitude.




Figure 4.8: Expected total number of overhead satellites at varying latitude.

4.2 Doppler Geometry Index

While visible satellite counts are useful and easy to understand, they neglect to account for the type of measurement being utilized from each satellite and the uniqueness of the information imparted by each satellite. For TOA-based navigation schemes like GNSSs, the DOP serves this purpose, assigning a numeric value to a given satellite geometry which describes the geometry’s quality from a TOA navigator’s perspective. Given an EKF measurement model $H_{m \times 4}$ mapping m pseudoranges to a state vector comprised only of position and clock bias, the corresponding DOP is given by

$$\text{DOP} = \sqrt{\text{tr}((H^T H)^{-1})} \quad (4.1)$$

under the condition that $H^T H$ is invertible, which is typically (though not *necessarily*) satisfied when $m \geq 4$ measurements are present.

DOP is not linked only to TOA measurements, though there are factors which make the adaptation of DOP to Doppler navigation difficult. The first issue is that of dimension. The TOA-based conception of DOP shown in (4.1) uses an H matrix in which all values are of like dimension — or, rather, all values are dimensionless, mapping length to length for each state. This is not so with a Doppler measurement model, which includes an observable

that is of dimension $\frac{\text{length}}{\text{time}}$ and states which are a mix of length and $\frac{\text{length}}{\text{time}}$, meaning that H contains values which may be dimensionless or have dimension $\frac{1}{\text{time}}$. In [45, 46] it is shown that this issue can be mitigated with sufficient mathematical machinery. Even with proper compensation for the dimensionality issue, a more fundamental problem remains, which is that a true DOP calculation requires that $H^T H$ be invertible. For Doppler-based positioning, $m \geq 8$ measurements are necessary (though not necessarily sufficient) to meet this invertibility condition, which rarely occurs with the combination of signal sources used in this work.

Since DOP cannot serve its usual role under the conditions of interest, another geometry indicator is required. To achieve this, the usual definition of DOP is modified to fit the needs of this work. In modifying the DOP to such an end, the formal meaning of DOP is lost, and so the resulting indicator is referred to as the *Doppler geometry index* (DGI) to emphasize that it is not, in general, a direct substitute for DOP, despite it having similar properties. The DGI is defined as

$$\text{DGI} = \sqrt{\text{tr} \left((H_p^T H_p)^{-1} \right)} \quad (4.2)$$

where H_p corresponds to the columns of the measurement model relating the observable to the position only, as shown in (3.8). This modification eliminates the dimensionality mismatch present in the naïve Doppler DOP and reduces the minimum number of measurements required to achieve invertibility from eight to three. As with DOP, a lower DGI suggests that a particular set of measurements provide better navigation results than a set of measurements with a higher DGI. Unlike a true DOP metric, however, the DGI is not dimensionless; it has a dimension of time. This work measures all velocities in units of meters per second, so all DGI values presented in this work have units of seconds.

Even with the DGI's lower requirement than DOP, some locations may not have enough concurrent measurements to make $H_p^T H_p$ invertible. It is not uncommon to have only a single

satellite overhead for minutes at a time. Under such conditions, the navigator’s position is not instantaneously observable given the measurements available, but an EKF can still utilize the measurements to improve its state estimates. To adapt DGI to such conditions, the measurements can be *batched*, lumping together measurements from a series of epochs. By selecting a sufficiently long batch time, a DGI can be obtained for any location of interest. In this way, locations can be compared to one another by selecting a common batch time that is sufficient to calculate a DGI. Locations can also be compared by leveraging the property that DGI (like DOP) tends to decrease as more measurements are added. In this way, a target DGI can be selected and the batching time taken to reach the target at each location can be compared, with a lower time-to-target-DGI indicating better geometry.

Figure 4.9 illustrates how the batched DGI evolves over time across a selection of 12 different satellite geometries. The scenarios model a stationary navigator at three different latitudes and four different starting times-of-day on 13 February 2021. All the navigators are located along 85.5° W longitude, and are each obtaining measurements at a rate of 1 Hz. The first two batching epochs have been truncated to ensure that all scenarios have sufficient data to compute a DGI. This result clearly shows the DGI decreasing over time, and that the equatorial navigators — which consistently have few satellites in view — perform more poorly than their higher-latitude counterparts, as expected. Of note is the equatorial navigator starting at 12:00 UTC, which has only a single satellite in view for the first roughly 30 seconds of its run. This causes the 0° 12:00 UTC navigator to have a DGI several orders of magnitude larger than the other navigators at the same latitude, which each have a minimum of two visible satellites throughout the time window shown. When additional satellites do rise above the horizon for the 12:00 navigator, its DGI drops quickly during the following few epochs, indicating that the additional information is particularly useful for positioning.

Figures 4.10 and 4.11 offer a more general summary of the performance expectations at varying latitudes. Each figure shows the same parameters, but with differing target values and integration times. Figure 4.10 shows the necessary batching time (again obtaining

measurements at 1 Hz) necessary to achieve $\text{DGI} \leq 100$ s in addition to the DGI after batching for 100 seconds at each latitude. Both the time-to-DGI and DGI-at-100 s results lead to similar conclusions.

- Doppler geometry does, as anticipated, generally improve as the navigator is moved towards the poles.
- The quality of the geometry can vary greatly depending on start time.
- General trends as a function of latitude do not substantially differ between the northern and southern hemispheres.
- Geometry is *especially* poor in the immediate vicinity of the equator.

The results shown in Figure 4.11 show the same data with longer batching periods. These data use a target value of $\text{DGI} \leq 20$ s and a constant batching period of 600 s (10 minutes). These longer batching periods smooth out some of the moment-to-moment variation in geometry at a given location, thereby offering a better sense of the nominal differences between latitudes. The conclusions drawn from these results do not differ from those previously discussed for the moderate batching time case.

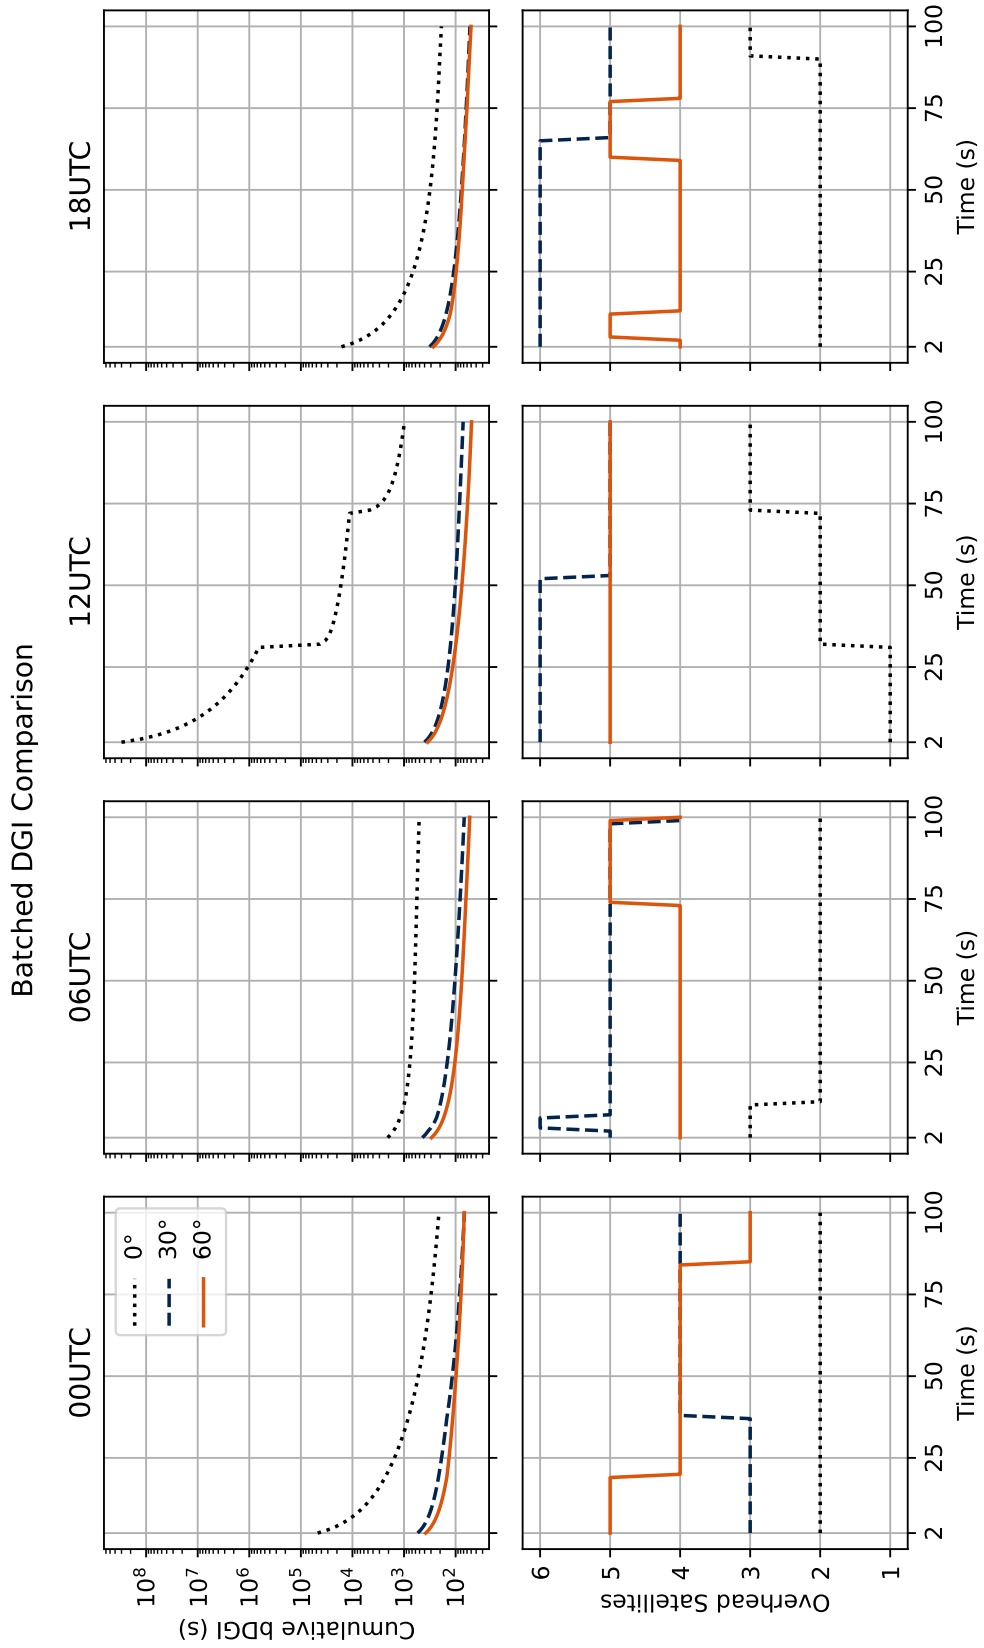


Figure 4.9: Evolution of batched DGI over time for various satellite geometries.

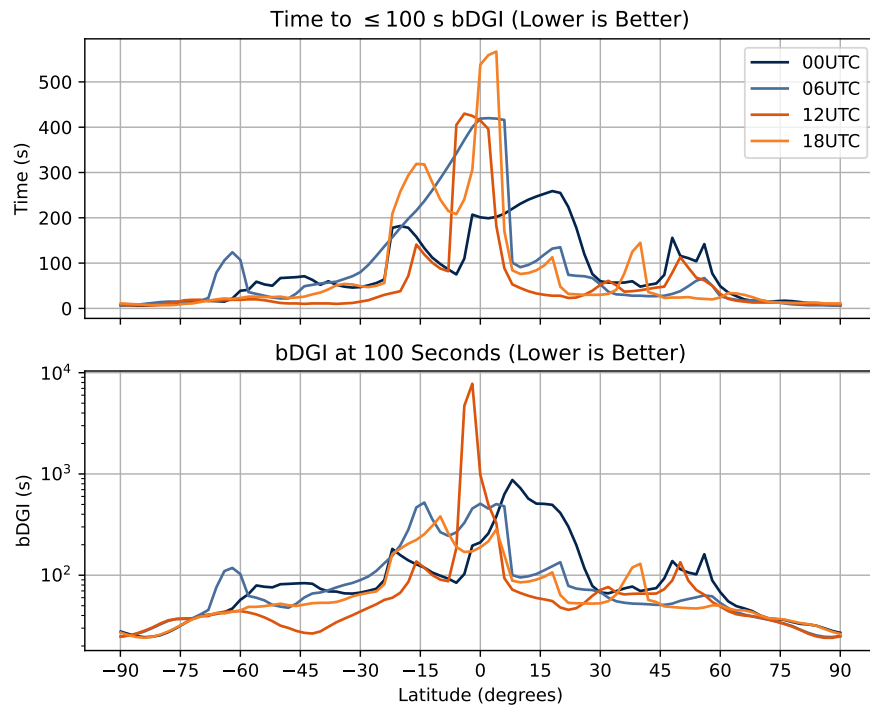


Figure 4.10: Latitude-dependent DGI characteristics with moderate batching times.

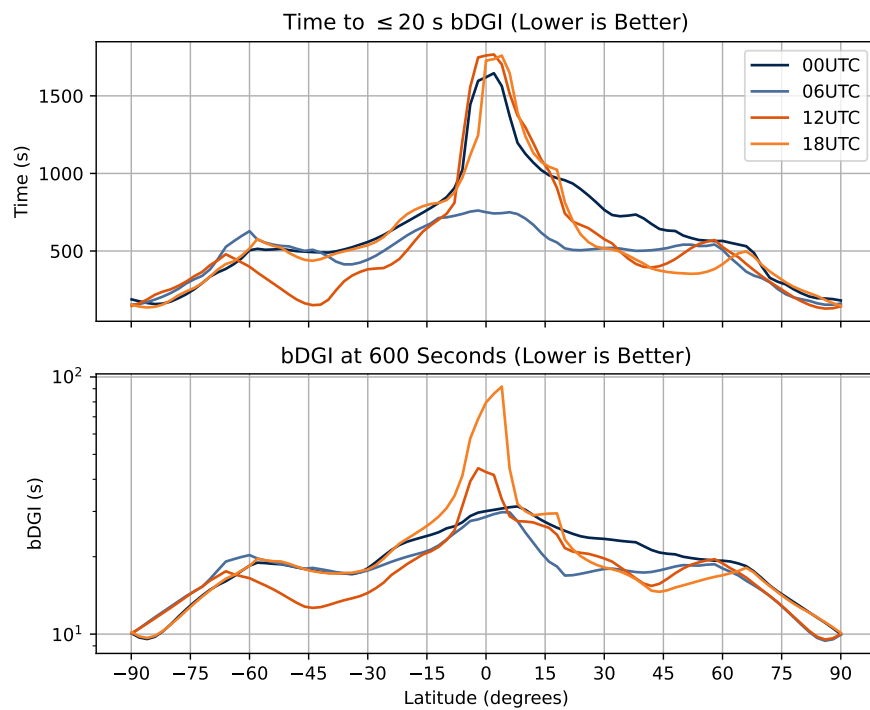


Figure 4.11: Latitude-dependent DGI characteristics with long batching times.

Chapter 5

Solo Navigator Performance Study

To demonstrate the viability of the Doppler navigator described in Chapter 3, a simulation environment is developed which models a series of UAVs navigating using the tightly-coupled Doppler-INS EKF previously discussed. This chapter describes the implementation of the simulation environment in more detail, and outlines a set of test scenarios which are used to conduct Monte Carlo performance tests of the navigator. The navigator performance characteristics resulting from the Monte Carlo tests are then presented and discussed.

5.1 Simulation Framework

To study the relative performance characteristics of the various navigators developed in this work, a measurement-level simulation environment is developed such that a series of non-collaborating navigators may be simulated. Further development of the simulation environment to support collaborative navigation is presented in Chapter 7. The simulation environment is further modified to support experimentally collected ground and aerial vehicle data in Chapter 8.

The generation of the simulated observables, it should be noted, is done *geometrically*; that is, pseudorange rate observables are generated directly from the simulated true relative motion between the receiver and each emitter. Observables are dirtied by Gaussian noise with standard deviations specified in Table 3.2.

Three functional components supporting the navigator constitute the simulation environment: inertial measurement generation, satellite position propagation, and measurement generation. Figure 5.1 shows the notional relationships between these high-level components. True trajectory generation and inertial measurement simulation are accomplished by a modified fork of Aceinna’s `gnss-ins-sim` library [47]. The trajectory generator takes as input a piecewise trajectory description and IMU error model terms to produce a time series of position, velocity, and attitude values which are used as truth, as well as specific force and angular velocity measurements from a simulated IMU. For this work, an IMU model for a generic low-end tactical grade IMU generating measurements at a rate of 100 Hz is used [38]. The parameters of the IMU error model are given in Table 5.1. Modeling satellite positions over time is essential to obtaining realistic results in a system that depends on measurements made from LEO satellites. The popular astronomy library Skyfield [48] is used to translate basic orbital parameters specified in the ubiquitous two-line element (TLE) format to a time series of position and velocity values. Though the library has many more powerful functions for astronomical computation, here it serves the straightforward role of a convenient wrapper around the well-known SGP4 algorithm [49] developed by NASA’s Jet Propulsion Laboratory in the 1980s, which remains in broad use. As a final step, measurements are generated on-demand while the navigator is running. Measurement generation creates a simulated observable given true emitter and receiver state information.

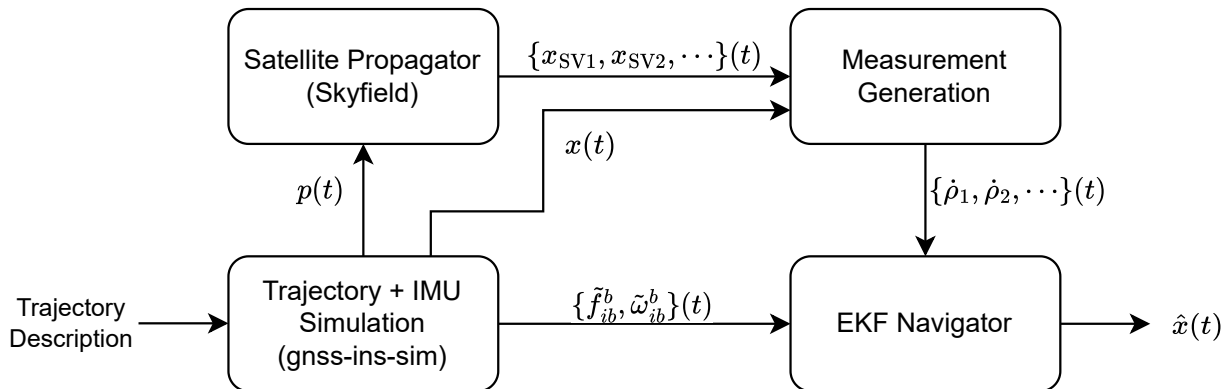


Figure 5.1: High-level simulation components and associated data flow.

Table 5.1: IMU model parameters.

Parameter	Value	Unit
Accelerometer Bias Drift	3.0×10^{-3}	m s^{-2}
Accelerometer Velocity Random Walk	10.0×10^{-3}	$\text{m s}^{-2}\text{Hz}^{-\frac{1}{2}}$
Gyroscope Bias Drift	50.0×10^{-6}	rad s^{-1}
Gyroscope Angle Random Walk	100.0×10^{-6}	$\text{rad s}^{-1}\text{Hz}^{-\frac{1}{2}}$

Given these basic building blocks of a simulation environment, the problem of *ephemeris error* should be addressed. Ephemeris error describes error in a navigation solution caused by uncorrected errors in the receiver’s knowledge of a satellite emitter’s position and/or velocity. GNSSs reduce the effects of ephemeris error by employing a sophisticated network of ground stations which can produce high-quality estimates of satellite orbital parameters communicated to the receiver through the GNSS signal’s data message. In general, opportunistic satellite navigation systems do not have the luxury of precise *a priori* satellite orbit information. Some satellite networks, such as ORBCOMM, do indeed broadcast unencrypted live ephemeris data which can be decoded by any receiver [10, 50, 51], but most do not. For satellites for which live ephemerides are not available, orbital parameters are available out-of-band from NORAD (or occasionally from the satellite network operator itself) in a variety of formats, which are updated daily can be propagated using SGP4. Satellite positions propagated in this way can accumulate errors on the order of a few kilometers over the course of a day [49]. Many prior works developing SOOP navigators employ methods such as simultaneous tracking and navigation (STAN) to mitigate the performance impacts caused by SGP4 propagation error. STAN is a technique in which the navigator attempts to simultaneously estimate its own position, velocity, and time (PVT) states in addition to those of the emitters [31, 52]. STAN is closely related to simultaneous localization and mapping (SLAM), which is often employed in vision navigation applications. The integration of STAN techniques to compensate for ephemeris errors are not in the scope of this work, so ephemeris errors are neglected.

5.2 Monte Carlo Testing

Monte Carlo simulation is used in this work as the primary mechanism for studying the navigation accuracy of the different navigators presented. In general, Monte Carlo simulation is concerned with analyzing the behavior of a process that is influenced by stochastic effects. It accomplishes this by repeating the process of interest many times, each time with freshly generated stochastic inputs and identical deterministic inputs. The set of outputs produced can then be used to draw conclusions about the general behavior of the process.

The results of Monte Carlo simulation are used to study two primary factors: the presence of estimate bias and the accuracy of the navigator's estimated covariance. Bias in the state estimate is often the result of unmodeled deterministic effects or imperfectly modeled stochastic effects, and manifests in Monte Carlo results as a bias in the sample mean errors observed across the entire test set. The second factor, the accuracy of the estimated covariance, is particularly pertinent when testing uncertainty-based processes like EKF's. Recall that the navigator itself produces not only an estimate of the state, but also a covariance matrix which describes the EKF's uncertainty in its own state estimate. For a Kalman filter that has a linear time update model (as opposed to the linear approximations used in an EKF), linear measurement models, and perfectly-modeled white Gaussian measurement and process noise, not only is its state estimate optimal, but its covariance estimate is correct. When the Kalman filter is extended to nonlinear systems, optimality of the state estimate is not the only guarantee that is lost. The covariance estimate may not be correct for an EKF, even with properly-tuned noise models. The results of a Monte Carlo test can produce a sample covariance for the entire state vector which can be compared to the EKF's estimated covariance to examine the estimator's ability to properly model uncertainty.

5.3 Test Scenario

The test results shown in this chapter utilize the sample trajectories shown in Figure 5.2 which include four aerial navigators. These trajectories are reused for the collaborative test results in Chapter 7 to aid in comparison. The scenario run time is 100 seconds and includes speeds varying from 0–10 m/s. The trajectories are defined in a generic north-east-down (NED) coordinate frame with an origin that may be arbitrarily fixed to any point on the earth. As part of the hot start assumption discussed in Chapter 3, the navigator is initialized with an imperfect estimate of its states. The errors applied to the initial position and velocity estimates are independent for each ECEF axis and are drawn from a Gaussian distribution with $\sigma = 3$ m for the position terms and $\sigma = 0.1$ m/s for the velocity terms. Initial attitude error is added by selecting a random axis of rotation from a uniform distribution on the unit sphere and selecting a random angle of rotation from a Gaussian distribution with $\sigma = 0.01^\circ$. The receiver clock drift is modeled as a constant 10 m/s with the navigator’s estimate initialized to 0 m/s.

Given that the navigators in this scenario are using LEO Doppler measurements as their only source of position-fixing navigation information, the navigators’ performance is substantially impacted by the number and locations of the LEO satellites of interest. This idea is discussed in detail in Chapter 4. So that the navigator is exposed to a range of satellite geometries in this testing, scenarios are defined with the trajectories located on the equator, at 30° N, and at 60° N. Due to symmetries induced by the rotation of the earth, the test locations’ longitude may be selected arbitrarily, so all test scenarios are located along 85.5° W longitude, which runs through the city of Auburn, Alabama. To account for changes in satellite geometry over time, tests are run with start times at 00:00 UTC, 06:00 UTC, 12:00 UTC, and 18:00 UTC. All start times use orbital data for 13 February 2021. The three latitudes and four times-of-day combine to produce 12 unique scenarios, each with different LEO satellite geometry. Each scenario is Monte Carlo tested with 100 iterations.

NED True Ground Track




Figure 5.2: The test trajectory consisting of four UAVs.

5.4 Test Results

The position estimate errors for the 100 Monte Carlo iterations starting at 00:00 UTC are shown in Figure 5.3. The figure includes the position estimate error for each individual Monte Carlo iteration as a collection of grey lines. The mean error across all iterations is shown as a black line. An important takeaway from this result is that the observed errors are zero-mean, or equivalently, that the estimator is unbiased. In addition are two sets of $\pm 3\sigma$ bounds. The navigator's estimated uncertainty derived from the covariance estimate is shown in red — note that there are 100 copies of each red line on top of each other, one from each iteration. A sample standard deviation calculated across Monte Carlo iterations is shown in blue. In this instance, the uncertainties produced by the navigator and the empirical uncertainty derived from the Monte Carlo test agree well, which indicates that

the navigator has a good sense of its own uncertainty. Appendix A includes counterparts of Figure 5.3 for a broader range of scenarios.

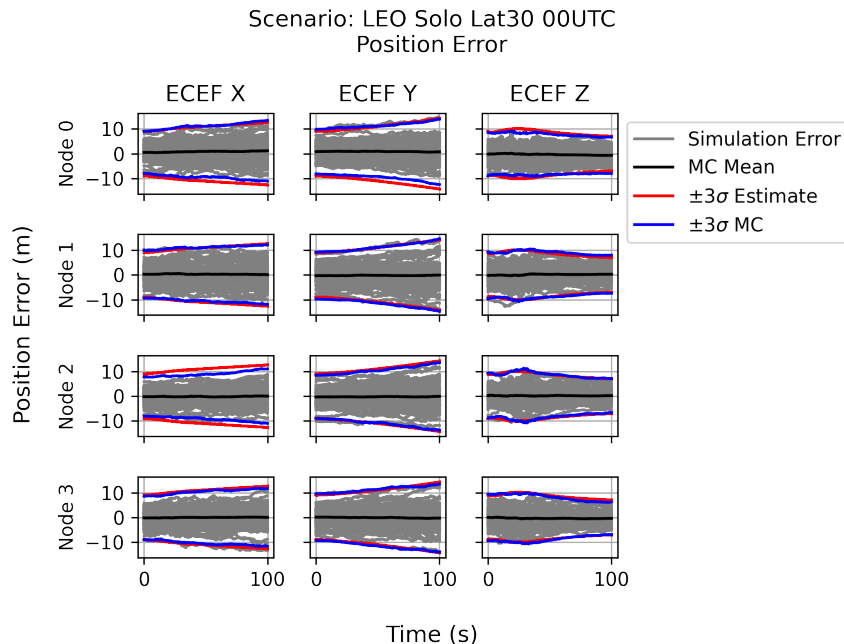


Figure 5.3: Position error for four non-collaborating navigators at 30° N.

Examining position error directly is useful for ensuring that the navigator is unbiased, but using it to compare performance across scenarios quickly becomes difficult. Error *magnitudes* are more tractable. Figures 5.4, 5.5, and 5.6 show the Monte Carlo position error magnitudes for the navigator at 0° , 30° N, and 60° N, respectively. Additionally, similar data with SOOP observables disabled entirely are shown in Figure 5.7. Positioning performance is quite poor at the equator, but improves markedly at higher latitudes with the 30° N and 60° N cases achieving similar results.

To more directly compare across latitudes, Figure 5.8 shows the estimated and Monte Carlo 3σ position error values for the three latitudes studied. Each bar consists of the mean Monte Carlo 3σ value at the end of the 100 s run, averaged across the four starting times-of-day and the four navigators in the group (so each bar includes data from a total of 1600 runs). Additionally, the mean estimated 3σ value is presented, as well as whisker bars

showing how much the filter’s estimated uncertainty varies from run to run. Such variation is caused by two effects. The first and most prominent effect is that of varying satellite geometry. The differing sets of satellites visible at the four tested times-of-day contribute to different covariance estimates produced by the navigator. The second effect arises from errors in the position or velocity estimates causing errors in measurement model H , which is a minor contributor to the variability in the results discussed in this chapter. The issue of measurement model errors becomes a more noticeable factor once collaborative methods are employed. This summary again shows the poor performance at the equator and improved performance at mid and upper latitudes. Additionally the results shown in the bar chart show that the navigator tends to be under-confident in its position estimate when it operates under poor geometric conditions. The data in Figure 5.8 are given numerically (along with results for other scenarios not yet introduced) in Table 7.1.

A comparison of the Monte Carlo end-of-run position error statistics (Figure 5.8) to the DGI metric in Section 4.2 is shown in Figure 5.9. Comparing these sets of results shows that while the DGI can serve as a general indicator of expected Doppler navigation performance, it is not perfect. With appropriate scaling, the batching times required to achieve $\text{DGI} \leq 100$ s align well with the observed positioning errors at 30° N and 60° N. Applying the same scaling to the positioning errors at the equator does not yield the same level of agreement between DGI and the Monte Carlo analysis. Both the DGI and Monte Carlo simulation agree that positioning performance is poorer at the equator than at the other latitudes, but the DGI batch time metric cannot be used to reliably determine the precise extent to which performance degrades under such conditions.

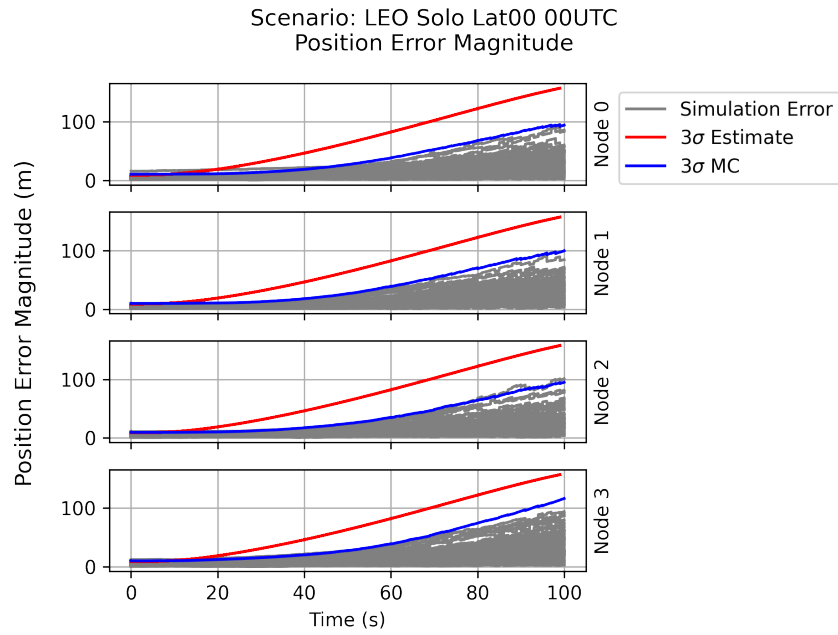


Figure 5.4: Position error magnitude for four non-collaborating navigators at the equator.

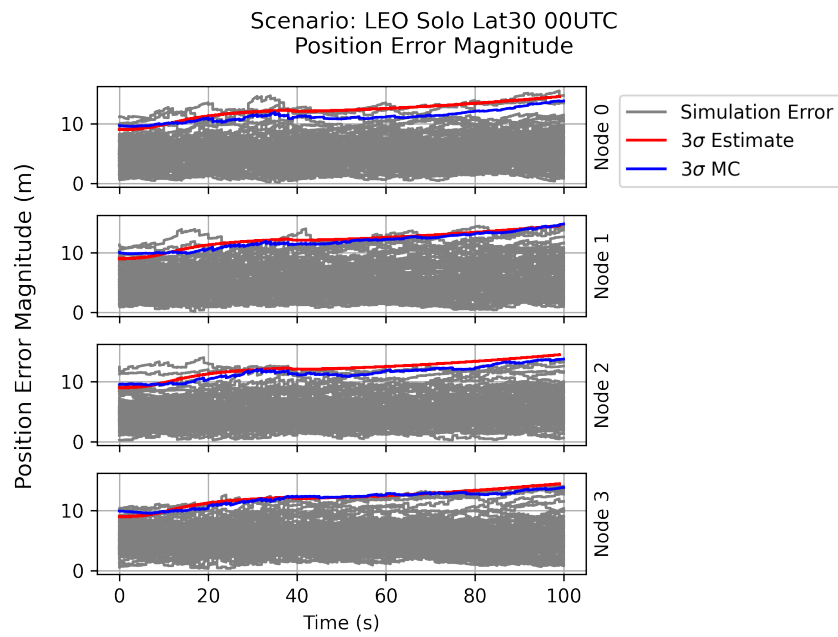


Figure 5.5: Position error magnitude for four non-collaborating navigators at 30° N.




Figure 5.6: Position error magnitude for four non-collaborating navigators at 60° N.




Figure 5.7: Position error magnitude for four non-collaborating navigators without SOOP measurements.

LEO Performance without Collaboration
100 Monte Carlo Iterations

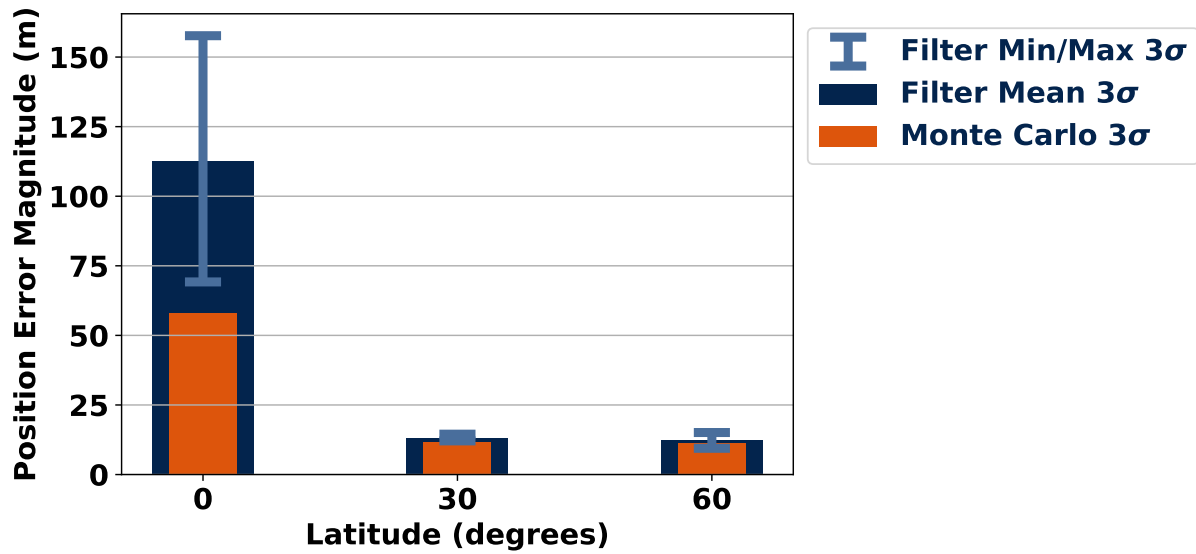


Figure 5.8: Summarized navigator positioning error magnitude.

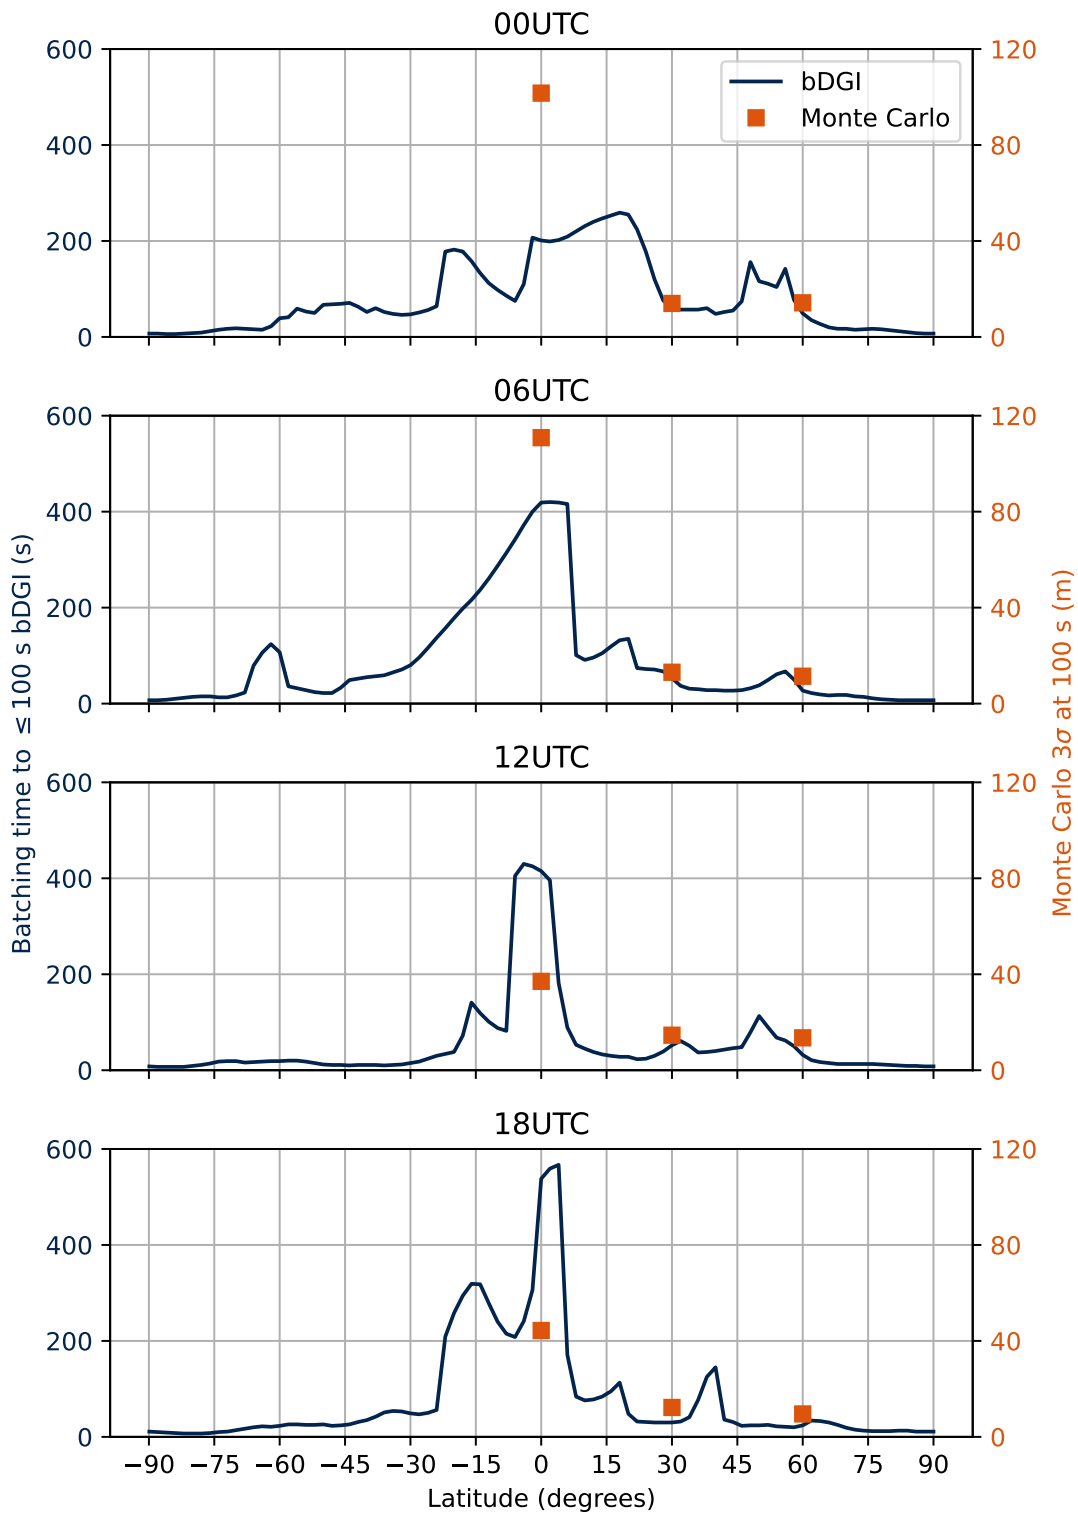


Figure 5.9: Comparison of time-to-DGI values and Monte Carlo 3σ position errors.

Chapter 6

Collaborative Navigation

Doppler-based opportunistic navigation is useful for providing navigation information in challenging environments, but the positioning accuracy discussed in Chapter 5 indicates that there is room for improvement. GNSSs, for instance, routinely achieve accuracies on the order of a few meters. One way that navigation performance can be improved is by combining information from multiple navigators in a *collaborative navigation* scheme. Collaborative navigation (or, equivalently, *cooperative navigation*) refers to any arrangement where information from multiple navigating entities is aggregated in some fashion to improve the performance of the overall system.

6.1 Peer-to-Peer Ranging

The specific collaborative environment studied in this work consists of a set of four identical aerial vehicles which are able to exchange arbitrary information with each other. Each navigator also has at its disposal high-quality range measurements to each of its peers, along with the same opportunistic Doppler measurements previously discussed. The simulation environment generates the peer-to-peer range measurements with a Gaussian noise error model with $\sigma = 5$ cm. These peer-to-peer range measurements could be obtained using a two-way time of flight (TW-ToF) system such as those utilized by ultrawideband (UWB) radio transceivers. The range between two navigators α and β as measured by α using a

TW-ToF measurement is given by

$$\begin{aligned}
 r &= \frac{1}{2}c(t_{\alpha\text{RX}} - t_{\alpha\text{TX}} - (t_{\beta\text{TX}} - t_{\beta\text{RX}})) \\
 &= \frac{1}{2}c(t_{\alpha\text{RX}} - t_{\alpha\text{TX}} - \Delta t_{\beta\text{Response}}).
 \end{aligned}
 \tag{6.1}$$

The need for four time terms comes from the *two-way* nature of TW-ToF measurements. More specifically, the situation modeled by (6.1) is that α transmits a signal at time $t_{\alpha\text{TX}}$, which is received by β at time $t_{\beta\text{RX}}$. After receiving a signal, β transmits a response at time $t_{\beta\text{TX}}$, which is received by α at time $t_{\alpha\text{RX}}$. The response time delay of β , denoted simply as $\Delta t_{\beta\text{Response}}$, is usually a well-known constant value determined by the manufacturer of the ranging equipment in use.

The principal advantage of such a range measurement system is that any dependence on absolute clock biases in either of the two participating navigators is effectively eliminated. The measurement model depends only on measured time *differences*, rather than the absolute time measurements necessary for obtaining range observables from TOA. This is not to say that oscillator quality has no effect on measurement error, but only that its importance is substantially reduced. Because of this, commercially available UWB transceivers boast range measurement accuracies on the order of 1 cm without needing to resort to using bulky and expensive atomic oscillators [53]. UWB transceivers can sometimes produce nominal error magnitudes small enough that quantization error becomes the dominant source of error, which is examined in [54]. This is in some sense a good problem to have, in that expected measurement errors are small, but quantization error is often poorly approximated by a Gaussian distribution, which can introduce modeling difficulties when using measurements obtained from a UWB transceiver with an EKF-based navigation environment.

6.2 Data Fusion

Obtaining high-quality range measurements between peer navigators is generally straightforward, as discussed in Section 6.1. Usefully leveraging such measurements to improve overall navigation performance is a more difficult task. Each of the collaborating navigators of interest is using Doppler measurements to obtain a PVT estimate as described in Chapter 3. An obvious technique for integrating peer-to-peer ranges would be to treat them like any other observable and apply them as an EKF measurement update. This naïve data fusion technique suffers from two main issues: *peer state uncertainty* and *fusion inconsistency*.

The problem of peer state uncertainty is analogous to that of ephemeris error in a GNSS navigator. The local navigator’s PVT state is not the only PVT state being estimated in the overall network. The peer navigator has its own state estimate and, critically, its own state uncertainty. Neglecting to account for the peer’s uncertainty causes the EKF to assume that the peer’s location is known perfectly, which would cause the estimator to treat the range measurement as if it imparts a greater degree of certainty than it actually does. In turn, this causes the local PVT uncertainty to be overly optimistic. Given that the peer’s state uncertainty is known (at least to the peer) and can be communicated to the local navigator, it can be used to properly model the uncertainty in the peer-to-peer range by either coalescing both uncertainties into a common covariance matrix (Section 6.2.1) or by inflating the measurement uncertainty itself (Section 6.2.2).

Fusion inconsistency is a more insidious issue. This problem arises due to improperly accounting for cross-correlations between the local and peer state estimate errors. More concretely, this means that even if peer state uncertainty is properly considered, applying a peer-to-peer range measurement as a regular EKF measurement update carries with it the implicit assumption that error in the local state estimate is uncorrelated with error in the peer state estimate. This is unlikely to be the case, especially given that the local and peer navigators are generally making use of the same SOOP measurements. In essence,

breaking the EKF’s assumptions in this way leads to a *double-counting* of information. The information held by each navigator in isolation is, to some extent, redundant when compared to its peers, yet the EKF assumes that no such redundancy exists. The various collaborative data fusion methods presented in this section have as a common trait some way to achieve *consistency* even in the presence of correlated error among the collaborators. The mechanism by which each method achieves fusion consistency varies considerably, with some attempting to estimate the error correlation itself, while others opt to employ a conservative fusion approach to become robust to unknown correlations.

This section introduces three collaborative data fusion methods. The first is a centralized navigator, which serves as a best case point of comparison. The second is a decentralized collaboration scheme which employs *covariance intersection* (CI) to fuse measurements into each navigator’s PVT solution. The third uses *inverse covariance intersection* (ICI) for decentralized data fusion in conjunction with augmented state vectors to estimate some of the relevant correlation values. The third method, consisting both of the application of ICI to SOOP navigation and the accompanying augmented state vector approach, comprise the principal contribution of this work.

6.2.1 Centralized Collaboration

Mathematically speaking, the easiest way to eliminate issues stemming from peer state uncertainty and fusion inconsistency is to consolidate all navigation information in a single, centralized estimator. Thus for n navigators each with a state vector as defined in (3.3), the central navigator has a single state vector in which the individual states for navigators 1.. n are concatenated, resulting in

$$\delta \mathbf{x}_{\text{Central}} = \left[\delta \mathbf{p}_1^T \quad \delta \mathbf{v}_1^T \quad \delta \Psi_1^T \quad \delta c \Delta t_1 \quad \delta c \dot{\Delta} t_1 \quad \cdots \quad \delta \mathbf{p}_n^T \quad \delta \mathbf{v}_n^T \quad \delta \Psi_n^T \quad \delta c \Delta t_n \quad \delta c \dot{\Delta} t_n \right]^T. \quad (6.2)$$

The centralized navigator takes in all IMU measurements, SOOP Doppler measurements, and peer-to-peer range measurements directly. This allows the EKF to account for both peer state uncertainties and error correlations by having the state uncertainties and all associated cross-correlations present in a single covariance matrix:

$$P_{\text{Central}} = \begin{bmatrix} P_{1,1} & \cdots & P_{1,n} \\ \vdots & \ddots & \vdots \\ P_{n,1} & \cdots & P_{n,n} \end{bmatrix} \quad (6.3)$$

where $P_{i,i}$ describes the state covariance of the i^{th} navigator and $P_{i,j}$ gives the state covariance cross-correlation between the i^{th} and j^{th} navigators. It is the presence of these cross-terms that permits the centralized navigator to properly integrate peer-to-peer ranges directly as an EKF measurement update. For example, a range observable between the i^{th} and j^{th} navigators, with $i, j \in \{1..n\}$ and $i < j$, would produce an EKF measurement model of the form

$$H_{i \rightarrow j, \text{Central}} = \begin{bmatrix} \mathbf{0}_{1 \times 11(i-1)} & \left(\frac{\mathbf{p}_i - \mathbf{p}_j}{|\mathbf{p}_i - \mathbf{p}_j|} \right)^T & \mathbf{0}_{1 \times 8} & \mathbf{0}_{1 \times 11(j-i-1)} & - \left(\frac{\mathbf{p}_i - \mathbf{p}_j}{|\mathbf{p}_i - \mathbf{p}_j|} \right)^T & \mathbf{0}_{1 \times 8} & \mathbf{0}_{1 \times 11(n-j)} \end{bmatrix}. \quad (6.4)$$

While mathematically advantageous, this method presents a number of challenges hindering its implementation in any real system. Communicating all measurement information to a single navigator running the centralized EKF is impractical, as it would require reliable, high-bandwidth, low-latency communication channels to every peer navigator, which is especially challenging for IMU measurements, since they are often reported at frequencies on the order of 100 Hz. Further, such an arrangement creates a single point of failure; the loss of the central navigator, or even the disruption of communications with the central navigator, constitutes a catastrophic failure of the entire network.

Issues of practicality aside, the centralized navigator remains useful as a point of comparison for the decentralized methods. In simulation, issues of latency and communications reliability can be readily ignored. As such, the centralized navigator is presented as a best case scenario. For any given scenario, it is reasonable to expect the application of decentralized collaboration to improve navigation performance relative to an equivalent set of non-collaborating navigators. Simultaneously, it would be *unreasonable* to expect any decentralized arrangement to perform better than the centralized navigator described in this section.

6.2.2 Decentralized Collaboration with CI

Decentralized approaches, wherein navigators share information with each other in an ad hoc fashion, are generally better suited for realistic implementation than their centralized counterparts, but they present their own set of challenges. As previously discussed, the issue of consistency of data fusion (or, equivalently, avoidance of double-counting redundant information) is a prominent concern for a navigator wishing to leverage peer-to-peer range measurements to improve position estimation performance. In the absence of *a priori* information about the relevant state covariance cross-terms, alternative approaches must be considered to adapt the EKF navigator.

The CI algorithm is a well-understood and widely used method of fusing two state estimates together in the presence of unknown cross-correlations. CI achieves this by eschewing optimality (even for linear systems) in favor of *consistency*. CI may not produce an optimal fusion, but it never produces a fused covariance that is more certain than an optimal fusion would be [55, 56].

The general form of the CI algorithm takes as input two different estimates of the same state vector, $\hat{\mathbf{x}}_a$ and $\hat{\mathbf{x}}_b$, along with each estimate's respective covariance, P_a and P_b . CI combines these four values to produce a fused state estimate and covariance given by $\hat{\mathbf{x}}_{\text{CI}}$

and P_{CI} , respectively. The fused values are computed by

$$P_{\text{CI}} = (\omega P_a^{-1} + (1 - \omega) P_b^{-1})^{-1} \quad (6.5)$$

$$\hat{\mathbf{x}}_{\text{CI}} = P_{\text{CI}} (\omega P_a^{-1} \hat{\mathbf{x}}_a + (1 - \omega) P_b^{-1} \hat{\mathbf{x}}_b) \quad (6.6)$$

for some value of weighting factor $\omega \in [0, 1]$. The value of ω is selected to minimize some cost function. This work uses the trace of the fused covariance $\text{tr}(P_{\text{CI}})$ as the cost function, but this is not the only valid option.

CI can be understood spacially as overbounding the intersection of the two input covariances (hence *covariance intersection*). Figure 6.1 illustrates this point in two dimensions by showing two hypothetical input uncertainty ellipses P_a and P_b . The uncertainty ellipse resulting from an optimal fusion of the two is unknowable in the absence of more cross-correlation information, so three plausible optimal fusion ellipses are shown in purple. The uncertainty ellipse produced by CI fusion of P_a and P_b is shown in teal. The uncertainty P_{CI} overbounds all possible optimal fused P ellipses.

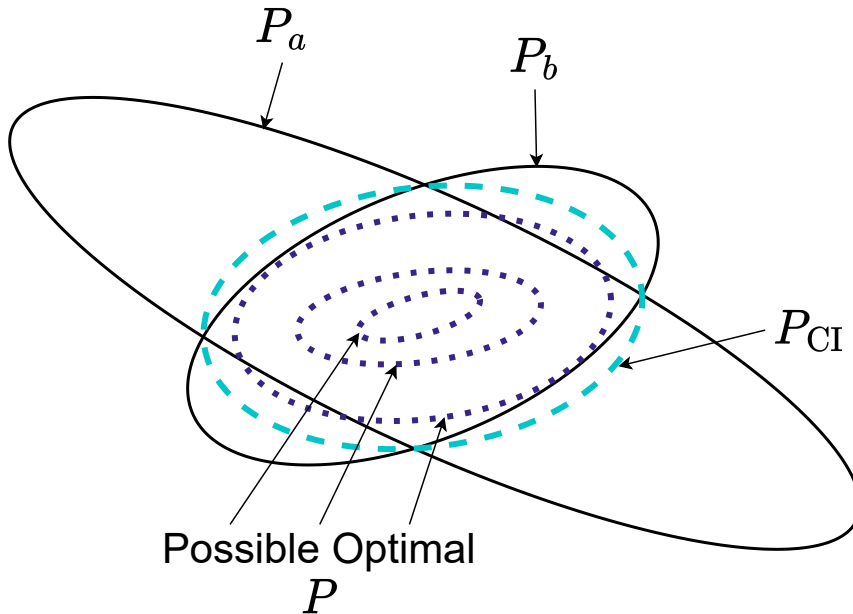


Figure 6.1: An informal illustration of the CI algorithm.

While CI addresses the issue of fusion consistency, it is not directly applicable to the problem at hand (decentralized navigation using peer-to-peer ranges) in its basic form. The CI equations given in (6.5) and (6.6) take as input two estimates of the same state vector. For the navigators described here, each has only its own estimate of its state vector. The peer-to-peer range measurements are related to its state vector through some measurement model H . Fortunately, there is a modified version of the CI algorithm which takes the form of a modified EKF measurement update [57]. The modified CI equations are

$$P_{\text{CI}} = (\omega P_a^{-1} + (1 - \omega) H^T R^{*-1} H)^{-1} \quad (6.7)$$

$$\hat{\mathbf{x}}_{\text{CI}} = \hat{\mathbf{x}}_a + (1 - \omega) P_{\text{CI}} H^T R^{*-1} (y - h(\hat{\mathbf{x}}_a)) \quad (6.8)$$

with weighting parameter ω being selected as before. New parameters not present in the original CI equations include measurement covariance R , measurand y , and nonlinear measurement model h . Equations (6.7) and (6.8) permit the integration of peer-to-peer range measurements while avoiding consistency issues, but the problem of peer state uncertainty remains. To compensate for this, the measurement covariance R is inflated to create a compensated measurement covariance R^* by

$$R_{\text{CI}}^* = (R_{\text{Collab}} + (-H) P_{\text{Peer}} (-H)^T) W_{\text{CI}} \quad (6.9)$$

for peer state uncertainty P_{Peer} and tunable weight factor $W_{\text{CI}} = 2.0$.

With the necessary modifications in place, this method can be usefully employed to integrate peer-to-peer measurements into a series of EKF-based navigators. Further information on the development of this particular navigation method, and special considerations regarding observability analysis, are examined in [58].

6.2.3 Decentralized Collaboration with ICI

As with its more well-known sibling CI, the ICI algorithm allows for consistent fusion of two state estimates in the presence of unknown cross-correlation. Also like CI, ICI is not (and *cannot* be) optimal, but ICI's advantage over CI is that it *more tightly bounds* the set of possible optimal fusion results; that is, it is less conservative than CI, and so (in general) produces superior results [19].

The general ICI equations for input state estimates and covariances $\hat{\mathbf{x}}_a$, $\hat{\mathbf{x}}_b$, P_a , and P_b are given by

$$\Theta = (\omega P_a + (1 - \omega)P_b)^{-1} \quad (6.10)$$

$$P_{\text{ICI}} = (P_a^{-1} + P_b^{-1} - \Theta)^{-1} \quad (6.11)$$

$$K = P_{\text{ICI}} (P_a^{-1} + \omega\Theta) \quad (6.12)$$

$$L = P_{\text{ICI}} (P_b^{-1} + (1 - \omega)\Theta) \quad (6.13)$$

$$\hat{\mathbf{x}}_{\text{ICI}} = K\hat{\mathbf{x}}_a + L\hat{\mathbf{x}}_b \quad (6.14)$$

with weighting parameter ω selected as before (such that $\text{tr}(P_{\text{ICI}})$ is minimized), and Θ , K , and L being intermediate terms which serve only to improve equation readability.

Returning to the two-dimensional spacial analogy first presented in Figure 6.1, if CI is an overbounding of intersecting uncertainty ellipses, then ICI can be thought about as an *under*bounding of intersecting *certainty* ellipses. Figure 6.2 shows an illustration of this idea with two input inverse covariances P_a^{-1} and P_b^{-1} , three plausible optimal fused inverse covariances P^{-1} , and the fused inverse covariance produced by ICI, P_{ICI}^{-1} . It is important to note that since these ellipses describe inverse covariances rather than covariances, a larger ellipse indicates a smaller amount of uncertainty. This way of modeling uncertainty can also be referred to as the *information form*, with the inverse covariance values describing an amount of information provided by a given state estimate.




Figure 6.2: An informal illustration of the ICI algorithm.

Thus far, ICI is in every way a dual algorithm to CI. It solves the same basic problem, it operates on the same inputs, and it produces outputs which, while not identical, are of the same form and function. The primary challenge which prevents ICI from being useful as a drop-in replacement for CI for the application of interest is that unlike CI, ICI does not have a modified form which allows it to function in place of an EKF measurement update. ICI equivalents to (6.7) and (6.8) do not exist. So to apply ICI to the problem of collaborative navigation with peer-to-peer ranges, more invasive changes to the navigator itself are required.

6.2.4 Adapting the Navigator for ICI

In its general form, ICI takes as input two estimates of the same state. Without any modification, such inputs do not exist anywhere in a network of independent EKF-based navigators.

Each navigator estimates its own state, and that estimate is the only one of its kind in the entire network. So, with the use of ICI in a collaborative navigation environment as a starting point, this section constructs a hybrid navigation architecture which borrows some ideas from the centralized navigator, while retaining most of the advantages of a decentralized network.

For clarity, this section shows the construction of the hybrid navigator with only two collaborators, called “Node 1” and “Node 2”. The concepts presented are trivially extensible to an arbitrary number of collaborators (implementation scalability concerns notwithstanding). The most fundamental modification that is made to the existing EKF navigator is the introduction of an *augmented state vector*. Each node’s state vector is extended to include estimates of the positions and velocities of each of its peers, resulting in

$$\delta \hat{\mathbf{x}}_1 = \begin{bmatrix} \delta \mathbf{p}_{1a} \\ \delta \mathbf{v}_{1a} \\ \delta \Psi_1 \\ \delta c \Delta t_1 \\ \delta c \dot{\Delta} t_1 \\ \delta \mathbf{p}_{2b} \\ \delta \mathbf{v}_{2b} \end{bmatrix} \quad \delta \hat{\mathbf{x}}_2 = \begin{bmatrix} \delta \mathbf{p}_{2a} \\ \delta \mathbf{v}_{2a} \\ \delta \Psi_2 \\ \delta c \Delta t_2 \\ \delta c \dot{\Delta} t_2 \\ \delta \mathbf{p}_{1b} \\ \delta \mathbf{v}_{1b} \end{bmatrix} \quad (6.15)$$

with numerical subscripts ($\delta \mathbf{p}_1$, $\delta \mathbf{p}_2$, etc.) indicating the node which the state describes and alphabetic subscripts ($\delta \mathbf{p}_a$, $\delta \mathbf{p}_b$, etc.) indicating whether the state estimate is (a) *local* to the node which it describes, or (b) a *remote* estimate maintained by some peer. The covariance

matrices which correspond to the state vectors given in (6.15) are

$$\begin{aligned}
P_1 &= \begin{bmatrix} P_{1a} \text{ (11}\times\text{11)} & P_{1a,2b}^* \text{ (11}\times\text{6)} \\ P_{2b,1a}^* \text{ (6}\times\text{11)} & P_{2b} \text{ (6}\times\text{6)} \end{bmatrix} = \begin{bmatrix} \mathbf{p}_{1a}\mathbf{p}_{1a} & \mathbf{p}_{1a}\mathbf{v}_{1a} & \mathbf{p}_{1a}\boldsymbol{\beta}_1 & \mathbf{p}_{1a}\mathbf{p}_{2b} & \mathbf{p}_{1a}\mathbf{v}_{2b} \\ \mathbf{v}_{1a}\mathbf{p}_{1a} & \mathbf{v}_{1a}\mathbf{v}_{1a} & \mathbf{v}_{1a}\boldsymbol{\beta}_1 & \mathbf{v}_{1a}\mathbf{p}_{2b} & \mathbf{v}_{1a}\mathbf{v}_{2b} \\ \boldsymbol{\beta}_1\mathbf{p}_{1a} & \boldsymbol{\beta}_1\mathbf{v}_{1a} & \boldsymbol{\beta}_1\boldsymbol{\beta}_1 & \boldsymbol{\beta}_1\mathbf{p}_{2b} & \boldsymbol{\beta}_1\mathbf{v}_{2b} \\ \mathbf{p}_{2b}\mathbf{p}_{1a} & \mathbf{p}_{2b}\mathbf{v}_{1a} & \mathbf{p}_{2b}\boldsymbol{\beta}_1 & \mathbf{p}_{2b}\mathbf{p}_{2b} & \mathbf{p}_{2b}\mathbf{v}_{2b} \\ \mathbf{v}_{2b}\mathbf{p}_{1a} & \mathbf{v}_{2b}\mathbf{v}_{1a} & \mathbf{v}_{2b}\boldsymbol{\beta}_1 & \mathbf{v}_{2b}\mathbf{p}_{2b} & \mathbf{v}_{2b}\mathbf{v}_{2b} \end{bmatrix} \\
P_2 &= \begin{bmatrix} P_{2a} \text{ (11}\times\text{11)} & P_{2a1b}^* \text{ (11}\times\text{6)} \\ P_{1b,2a}^* \text{ (6}\times\text{11)} & P_{1b} \text{ (6}\times\text{6)} \end{bmatrix} = \begin{bmatrix} \mathbf{p}_{2a}\mathbf{p}_{2a} & \mathbf{p}_{2a}\mathbf{v}_{2a} & \mathbf{p}_{2a}\boldsymbol{\beta}_2 & \mathbf{p}_{2a}\mathbf{p}_{1b} & \mathbf{p}_{2a}\mathbf{v}_{1b} \\ \mathbf{v}_{2a}\mathbf{p}_{2a} & \mathbf{v}_{2a}\mathbf{v}_{2a} & \mathbf{v}_{2a}\boldsymbol{\beta}_2 & \mathbf{v}_{2a}\mathbf{p}_{1b} & \mathbf{v}_{2a}\mathbf{v}_{1b} \\ \boldsymbol{\beta}_2\mathbf{p}_{2a} & \boldsymbol{\beta}_2\mathbf{v}_{2a} & \boldsymbol{\beta}_2\boldsymbol{\beta}_2 & \boldsymbol{\beta}_2\mathbf{p}_{1b} & \boldsymbol{\beta}_2\mathbf{v}_{1b} \\ \mathbf{p}_{1b}\mathbf{p}_{2a} & \mathbf{p}_{1b}\mathbf{v}_{2a} & \mathbf{p}_{1b}\boldsymbol{\beta}_2 & \mathbf{p}_{1b}\mathbf{p}_{1b} & \mathbf{p}_{1b}\mathbf{v}_{1b} \\ \mathbf{v}_{1b}\mathbf{p}_{2a} & \mathbf{v}_{1b}\mathbf{v}_{2a} & \mathbf{v}_{1b}\boldsymbol{\beta}_2 & \mathbf{v}_{1b}\mathbf{p}_{1b} & \mathbf{v}_{1b}\mathbf{v}_{1b} \end{bmatrix}
\end{aligned} \tag{6.16}$$

with

$$\boldsymbol{\beta}_1 = \begin{bmatrix} \delta\Psi_1 \\ \delta c\Delta t_1 \\ \delta c\dot{\Delta}t_1 \end{bmatrix} \quad \boldsymbol{\beta}_2 = \begin{bmatrix} \delta\Psi_2 \\ \delta c\Delta t_2 \\ \delta c\dot{\Delta}t_2 \end{bmatrix} \tag{6.17}$$

so that $\boldsymbol{\beta}$ encompasses the attitude and clock terms which are only present in the respective local state vectors. Augmenting the state vectors and covariances in this way, with each navigator keeping an estimate of the positions and velocities of each of its peers, creates two estimates of the same state, at least for some subset of the original EKF state vector. This fact is leveraged to apply ICI. During the EKF time update, these auxiliary states are propagated assuming constant velocity. A process noise model given by

$$Q_{\text{Aux}} = \begin{bmatrix} \mathbf{0}_{3\times 3} & \mathbf{0}_{3\times 3} \\ \mathbf{0}_{3\times 3} & 1000.0I_3 \end{bmatrix} \tag{6.18}$$

is concatenated along a block diagonal for each set of auxiliary states added onto the navigator.

Augmenting the state vectors presents an additional advantage. The presence of the P^* terms in (6.16) constitute a subset of the inter-navigator cross-correlation values present in the centralized navigator. The presence of these terms allows a range measurement taken by Node 1 to be applied between states \mathbf{p}_{1a} and \mathbf{p}_{2b} directly as an EKF measurement update with no special machinery needed to ensure fusion consistency. The corresponding measurement model would be

$$H_{1 \rightarrow 2} = \left[\left(\frac{\mathbf{p}_1 - \mathbf{p}_2}{|\mathbf{p}_1 - \mathbf{p}_2|} \right)^T \quad \mathbf{0}_{1 \times 8} \quad - \left(\frac{\mathbf{p}_1 - \mathbf{p}_2}{|\mathbf{p}_1 - \mathbf{p}_2|} \right)^T \right] \quad (6.19)$$

with

$$R_{\text{ICI}}^* = R_{\text{Collab}} W_{\text{ICI}} \quad (6.20)$$

where $W_{\text{ICI}} = 4.0$ is an adjustable tuning parameter.

This does not mean that this hybrid navigator system effectively reduces to a centralized navigator, nor does it imply that performance must match that of a centralized navigator. It is important to note that in the case shown in (6.19), the measurement is applied between \mathbf{p}_{1a} and \mathbf{p}_{2b} , *not* \mathbf{p}_{2a} . Node 1 still does not have access to all the information that Node 2 might have, including Node 2's IMU measurements and SOOP observables. All Node 1 is able to do to update its peer (*b*) state estimates is to employ constant velocity time updates and peer-to-peer range updates. If Node 1's navigator were left alone to operate in this way, updating its own local (*a*) states as normal, the quality of the peer state estimates would become poorer over time. Node 1 would never benefit from information that Node 2 has in its peer state estimates. This is where the ICI algorithm itself becomes important: it allows Node 1 and Node 2 to periodically swap state estimates, sharing any beneficial information

that each has about the other’s position and velocity, while also effectively re-initializing each’s peer state estimates.

Two logistical hurdles remain to applying an ICI fusion update to fuse the Node 1 and Node 2 state vectors: each node has some states which are unique and the nodes’ states are not in the same order. The second problem is quite easy to solve. Each of the two collaborators must permute its respective state vector and covariance into some commonly accepted *canonical order*. Solving the first issue, that of unique states, is somewhat more involved. Consider the clock bias term in Node 1’s state vector, $c\Delta t_1$. This term is not found in Node 2’s state vector. In a Kalman filter sense, Node 2 has an infinite and uncorrelated uncertainty of this state (i.e. it has no information whatsoever). Leveraging this, an arbitrarily-selected state value can be placed in Node 2’s state vector, accompanied by the aforementioned infinite and uncorrelated covariance value. In the interest of avoiding floating-point numerical issues when implementing this idea, an extremely-large-but-finite value (denoted Ω) is used in place of infinity (this work uses $\Omega = 10^{100}$). These dummy state estimates can be appended as needed to create a common state vector composition and ordering between the two navigators, while ensuring that the ICI algorithm knows that some nodes have no useful information to contribute on some of the states. Carrying out this process for Node 1 and Node 2 results in

$$\hat{\mathbf{x}}'_1 = \begin{bmatrix} \mathbf{p}_{1a} \\ \mathbf{v}_{1a} \\ \mathbf{p}_{2b} \\ \mathbf{v}_{2b} \\ \beta_1 \\ \mathbf{0}_{5 \times 1} \end{bmatrix} \qquad \hat{\mathbf{x}}'_2 = \begin{bmatrix} \mathbf{p}_{1b} \\ \mathbf{v}_{1b} \\ \mathbf{p}_{2a} \\ \mathbf{v}_{2a} \\ \mathbf{0}_{5 \times 1} \\ \beta_2 \end{bmatrix} \qquad (6.21)$$

$$\begin{aligned}
P'_1 &= \begin{bmatrix} P_{1a} & P_{1a,2b} & P_{1a,\beta_1} & \mathbf{0}_{6 \times 5} \\ P_{2b,1a} & P_{2b} & P_{2b,\beta_1} & \mathbf{0}_{6 \times 5} \\ P_{\beta_1,1a} & P_{\beta_1,2b} & P_{\beta_1} & \mathbf{0}_{5 \times 5} \\ \mathbf{0}_{5 \times 6} & \mathbf{0}_{5 \times 6} & \mathbf{0}_{5 \times 5} & \Omega I_{5 \times 5} \end{bmatrix} = \begin{bmatrix} p_{1a} p_{1a} & p_{1a} v_{1a} & p_{1a} p_{2b} & p_{1a} v_{2b} & p_{1a} \beta_1 & \mathbf{0}_{3 \times 5} \\ v_{1a} p_{1a} & v_{1a} v_{1a} & v_{1a} p_{2a} & v_{1a} v_{2b} & v_{1a} \beta_1 & \mathbf{0}_{3 \times 5} \\ p_{2b} p_{1a} & p_{2b} v_{1a} & p_{2b} p_{2b} & p_{2b} v_{2b} & p_{2b} \beta_1 & \mathbf{0}_{3 \times 5} \\ v_{2b} p_{1a} & v_{2b} v_{1a} & v_{2b} p_{2b} & v_{2b} v_{2b} & v_{2b} \beta_1 & \mathbf{0}_{3 \times 5} \\ \beta_1 p_{1a} & \beta_1 v_{1a} & \beta_1 p_{2b} & \beta_1 v_{2b} & \beta_1 \beta_1 & \mathbf{0}_{5 \times 5} \\ \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 5} & \Omega I_{5 \times 5} \end{bmatrix} \\
P'_2 &= \begin{bmatrix} P_{1b} & P_{1b,2a} & P_{1b,\beta_2} & P_{1b,\beta_2} & P_{1b,\beta_2} \\ P_{2a,1b} & P_{2a} & P_{2a,\beta_2} & P_{2a,\beta_2} & P_{2a,\beta_2} \\ \mathbf{0}_{5 \times 6} & \mathbf{0}_{5 \times 6} & \Omega I_{5 \times 5} & \mathbf{0}_{5 \times 5} & \mathbf{0}_{5 \times 5} \\ P_{\beta_2,1b} & P_{\beta_2,2a} & P_{\beta_2} & P_{\beta_2} & P_{\beta_2} \end{bmatrix} = \begin{bmatrix} p_{1b} p_{1b} & p_{1b} v_{1b} & p_{1b} p_{2a} & p_{1b} v_{2a} & p_{1b} \beta_2 \\ v_{1b} p_{1b} & v_{1b} v_{1b} & v_{1b} p_{2a} & v_{1b} v_{2a} & v_{1b} \beta_2 \\ p_{2a} p_{1b} & p_{2a} v_{1b} & p_{2a} p_{2a} & p_{2a} v_{2a} & p_{2a} \beta_2 \\ v_{2a} p_{1b} & v_{2a} v_{1b} & v_{2a} p_{2a} & v_{2a} v_{2a} & v_{2a} \beta_2 \\ \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 3} & \mathbf{0}_{5 \times 3} & \Omega I_{5 \times 5} \\ \beta_2 p_{1b} & \beta_2 v_{1b} & \beta_2 p_{2a} & \beta_2 v_{2a} & \beta_2 \beta_2 \end{bmatrix}
\end{aligned} \tag{6.22}$$

which, when taken together, can be used as-is to perform an ICI fusion update to obtain fused state estimate and covariance:

$$\hat{\mathbf{x}}'_{\text{ICI}} = \begin{bmatrix} \mathbf{p}_1 \\ \mathbf{v}_1 \\ \mathbf{p}_2 \\ \mathbf{v}_2 \\ \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{bmatrix} \quad (6.23)$$

$$P'_{\text{ICI}} = \begin{bmatrix} P_1 \text{ (6}\times\text{6)} & P_{1,2} \text{ (6}\times\text{6)} & P_{1,\beta_1} \text{ (6}\times\text{5)} & P_{1,\beta_2} \text{ (6}\times\text{5)} \\ P_{2,1} \text{ (6}\times\text{6)} & P_2 \text{ (6}\times\text{6)} & P_{2,\beta_1} \text{ (6}\times\text{5)} & P_{2,\beta_2} \text{ (6}\times\text{5)} \\ P_{\beta_1,1} \text{ (5}\times\text{6)} & P_{\beta_1,2} \text{ (5}\times\text{6)} & P_{\beta_1} \text{ (5}\times\text{5)} & \mathbf{0}_{5\times 5} \\ P_{\beta_2,1} \text{ (5}\times\text{6)} & P_{\beta_2,2} \text{ (5}\times\text{6)} & \mathbf{0}_{5\times 5} & P_{\beta_2} \text{ (5}\times\text{5)} \end{bmatrix} \quad (6.24)$$

$$= \begin{bmatrix} \mathbf{p}_1\mathbf{p}_1 & \mathbf{p}_1\mathbf{v}_1 & \mathbf{p}_1\mathbf{p}_2 & \mathbf{p}_1\mathbf{v}_2 & \mathbf{p}_1\boldsymbol{\beta}_1 & \mathbf{p}_1\boldsymbol{\beta}_2 \\ \mathbf{v}_1\mathbf{p}_1 & \mathbf{v}_1\mathbf{v}_1 & \mathbf{v}_1\mathbf{p}_2 & \mathbf{v}_1\mathbf{v}_2 & \mathbf{v}_1\boldsymbol{\beta}_1 & \mathbf{v}_1\boldsymbol{\beta}_2 \\ \mathbf{p}_2\mathbf{p}_1 & \mathbf{p}_2\mathbf{v}_1 & \mathbf{p}_2\mathbf{p}_2 & \mathbf{p}_2\mathbf{v}_2 & \mathbf{p}_2\boldsymbol{\beta}_1 & \mathbf{p}_2\boldsymbol{\beta}_2 \\ \mathbf{v}_2\mathbf{p}_1 & \mathbf{v}_2\mathbf{v}_1 & \mathbf{v}_2\mathbf{p}_2 & \mathbf{v}_2\mathbf{v}_2 & \mathbf{v}_2\boldsymbol{\beta}_1 & \mathbf{v}_2\boldsymbol{\beta}_2 \\ \boldsymbol{\beta}_1\mathbf{p}_1 & \boldsymbol{\beta}_1\mathbf{v}_1 & \boldsymbol{\beta}_1\mathbf{p}_2 & \boldsymbol{\beta}_1\mathbf{v}_2 & \boldsymbol{\beta}_1\boldsymbol{\beta}_1 & \mathbf{0}_{5\times 5} \\ \boldsymbol{\beta}_2\mathbf{p}_1 & \boldsymbol{\beta}_2\mathbf{v}_1 & \boldsymbol{\beta}_2\mathbf{p}_2 & \boldsymbol{\beta}_2\mathbf{v}_2 & \mathbf{0}_{5\times 5} & \boldsymbol{\beta}_2\boldsymbol{\beta}_2 \end{bmatrix}.$$

The resulting canonical state vector and covariance can then be rearranged back into the preferred state ordering for each navigator, at which point each carries on as normal until another ICI fusion update is conducted. All Monte Carlo simulation results involving ICI in this work carry out an ICI fusion update every 10 seconds.

Chapter 7

Collaborative Navigator Performance Study

To study the effectiveness of each collaborative technique presented in Chapter 6, each may be Monte Carlo tested using the same trajectories used for the solo navigator results (see Figure 5.2). Each navigator in this chapter is tested at the same three latitudes and four times of day described in Section 5.3, so the observed satellite geometry, and all other deterministic effects, are identical across collaborative methods.

Figures 7.1, 7.2, and 7.3 show summarized performance characteristics computed similarly to those in Figure 5.8 for 0° , 30° N, and 60° N, respectively. Note the drastically different y -axis scaling for the equatorial scenario compared to the other two. Instead of comparing performance across latitudes for non-collaborating navigators, each of these plots examines the relative performance of each collaborative technique one latitude at a time.

As anticipated, the scenarios without collaboration generally perform the worst and those with a centralized navigator perform the best across latitude classes, with the two decentralized methods usually falling somewhere in between. One exception to this generalization is the performance of the CI navigator at 60° N (Figure 7.3), which produced slightly larger position errors than the solo navigators, but overall the degradation in that case is small. Further, for the two higher latitudes, CI offered little benefit as a collaborative method. This could be caused by the high degree of correlation among the navigators' covariances, which would cause the fused uncertainty produced by CI to be no better than the input uncertainties. In this case, the performance gains realized by the ICI navigator arise

not so much from the substitution of ICI for CI, but from the hybrid navigator architecture that is required to permit the use of ICI in the first place. The extended state vectors allow each node to estimate parts of the cross-correlations between navigators — even if those cross-correlations are not as accurate as those produced by the centralized navigator. Thus the ICI navigator can achieve accuracies closer to that of the centralized navigator than that of the solo navigators.

The results at the equator (Figure 7.1), however, tell a different story. Here, it is the CI navigator that comes closer to replicating the performance of the centralized case. It is noteworthy, though, that despite the CI navigator offering improvement, this is not reflected in the estimated covariance, indicating that the EKF has no idea that it has improved its position estimate. CI does not always perform so well in environments with poor preexisting performance, as exhibited in Figure 7.4, which shows performance characteristics for the various collaborative methods with no LEO positioning information whatsoever. The relative performance of each technique in this IMU-only case more closely resembles the scenarios at higher latitudes than it does the scenario at the equator.

The performance figures shown in Figures 7.1–7.4 are given numerically in Table 7.1.




Figure 7.1: Comparison of collaborative positioning techniques at the equator.

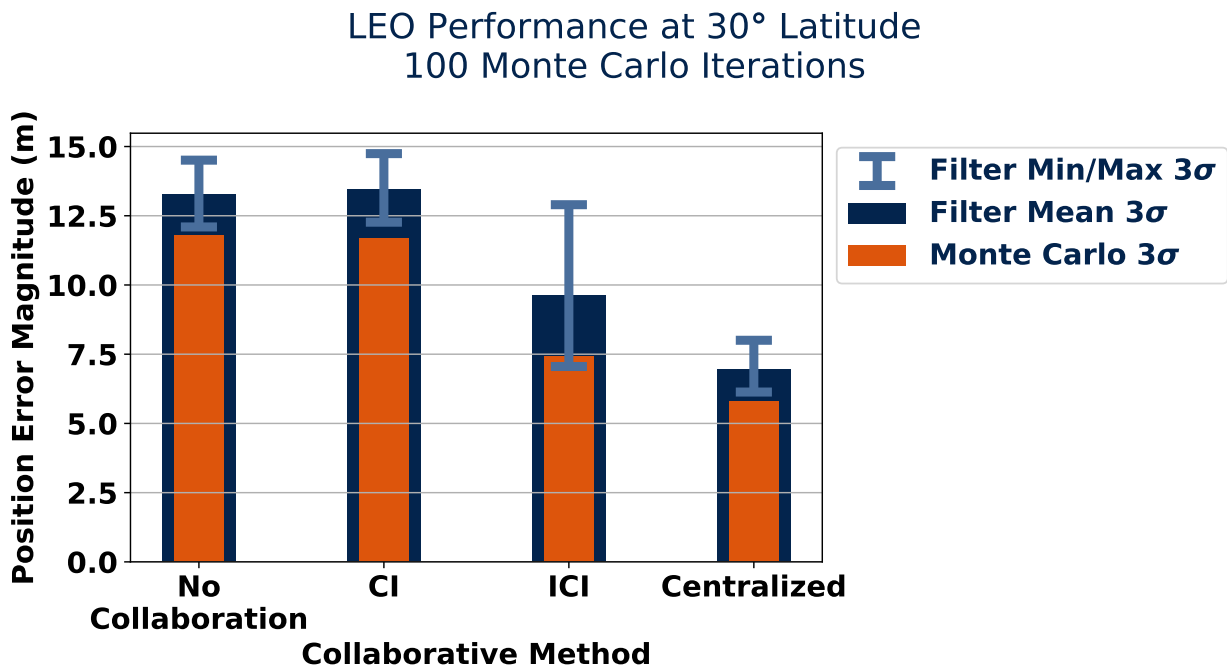


Figure 7.2: Comparison of collaborative positioning techniques at 30° N.

LEO Performance at 60° Latitude
100 Monte Carlo Iterations




Figure 7.3: Comparison of collaborative positioning techniques at 60° N.

IMU Performance at 00° Latitude
100 Monte Carlo Iterations

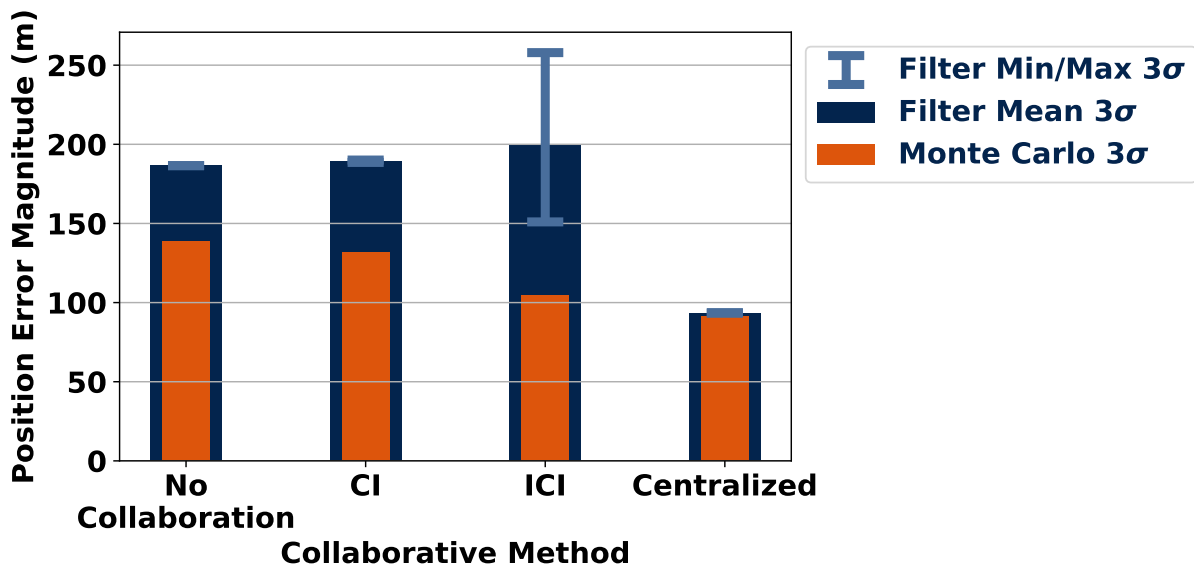


Figure 7.4: Comparison of collaborative positioning techniques with no LEO measurements.

Table 7.1: Comparison of collaborative positioning accuracies.

<i>(All values in meters)</i>	No Collaboration	CI	ICI	Centralized
<i>IMU Only</i>				
Monte Carlo 3σ	138.6	131.6	104.3	91.6
Min Filter 3σ	186.5	188.6	151.0	93.3
Mean Filter 3σ	186.5	189.3	199.6	93.4
Max Filter 3σ	186.6	190.1	257.9	93.6
<i>LEO 0° N Latitude</i>				
Monte Carlo 3σ	57.9	34.0	47.5	31.3
Min Filter 3σ	69.2	74.1	57.8	35.1
Mean Filter 3σ	112.5	118.5	103.7	56.3
Max Filter 3σ	157.7	177.0	156.5	78.7
<i>LEO 30° N Latitude</i>				
Monte Carlo 3σ	11.8	11.7	7.4	5.8
Min Filter 3σ	11.8	12.3	7.1	6.1
Mean Filter 3σ	13.3	13.5	9.6	7.0
Max Filter 3σ	14.5	14.7	12.9	8.0
<i>LEO 60° N Latitude</i>				
Monte Carlo 3σ	11.2	11.6	7.2	5.7
Min Filter 3σ	9.4	9.6	4.1	4.7
Mean Filter 3σ	12.4	12.6	7.1	6.3
Max Filter 3σ	15.1	15.3	12.2	7.7

Chapter 8

Test Track Experiments

The central focus of this work is the development of the distributed ICI architecture and the comparison of the ICI navigator against other collaborative methods. This end is well-covered in Chapter 7. Nonetheless, it is useful to take one more step in the direction of implementation on real systems. This chapter adapts the solo SOOP navigator described in Section 3.3 and the decentralized collaborative navigators described in Sections 6.2.2 and 6.2.4 to work with experimentally collected data.

Tests were conducted at Auburn University’s NCAT test track, which is located in Lee County, Alabama, roughly 9 km (5.6 mi) southeast of Opelika at 32.596° N, 85.302° W. The test track’s primary purpose is the testing of highway paving materials and techniques [59], but the site also hosts facilities for vehicle autonomy or navigation research such as this. In particular, the tests discussed in this chapter were conducted on the asphalt skid pad located near the southeast corner of the test track property (32.5955° N, 85.2953° W).

8.1 Modifications to the Navigators

The navigators developed in Chapters 3 and 6 cannot be used as-is with the test track data for two reasons: (1) the navigators assume that the dead reckoning timestep and the spacing of measurement epochs is perfectly uniform and known *a priori* and (2) the two cars use an entirely different time update mechanism than previously developed (wheel speed instead

of 6-DOF IMU). The former issue is more or less a matter of implementation, and is easily dealt with, while overcoming the latter issue requires more invasive changes.

To better accommodate ground vehicles, the state vector for all three vehicles is converted from an ECEF reference frame to an NED reference frame fixed to the ground near the center of the skid pad. Other than the change in reference frame, the resulting state vector,

$$\delta \mathbf{x} = \begin{bmatrix} \delta \mathbf{p} \\ \delta \mathbf{v} \\ \delta \Psi \\ \delta c\Delta t \\ \delta c\dot{\Delta} t \end{bmatrix} = \begin{bmatrix} \delta p_N \\ \delta p_E \\ \delta p_D \\ \delta v_N \\ \delta v_E \\ \delta v_D \\ \delta \varphi \\ \delta \theta \\ \delta \psi \\ \delta c\Delta t \\ \delta c\dot{\Delta} t \end{bmatrix} \quad (8.1)$$

is unchanged from the one described in (3.3). The navigators for the cars are further modified to track attitude as yaw angle only (assuming pitch and roll to be zero). While the dimension of the estimated attitude was reduced, the car navigators are *not* entirely reduced to two dimensions. They each still estimate a vertical component of position and velocity, which allows the cars to better account for the two-meter elevation difference from one end of the skid pad to the other. This difference in elevation would negatively impact Doppler and collaborative navigation performance if neglected.

The EKF time update for the cars consists of a constant-velocity update of the estimated position, and a dead reckoned update of the yaw from the single-axis gyro. The corresponding empirically determined continuous process noise model is

$$Q_{\text{Car}} = \begin{bmatrix} 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.16 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.04 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 2.25 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2.25 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.5625 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 40 \times 10^{-9} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.01 \end{bmatrix}. \quad (8.2)$$

Wheel speed measurements are applied as a separate EKF measurement update. Integrating wheel speed measurements in a measurement update allows for a more straightforward modeling of the uncertainty of the wheel speed measurements (i.e. the errors introduced by wheel speed measurements are more readily modeled as a measurement noise source than a process noise source). Additionally, this measurement update allows for a more fine-tuned modeling of the implicit no-side-slip assumption that is often a part of wheel speed odometry. As such, each odometry measurement is given by

$$y_{\text{Odometer}} = \mathbf{v}^b = \begin{bmatrix} v_F \\ v_R \\ v_D \end{bmatrix} = \begin{bmatrix} v_F \\ 0 \\ 0 \end{bmatrix} \quad (8.3)$$

where the body frame velocity \mathbf{v}^b is expressed as the measured wheel speed in the forward direction v_F and implicit 0 measurements of the velocity in the right and down directions, v_R and v_D , respectively. The measurement model associated with such a measurement, neglecting the effects of heading estimate error, is

$$H_{\text{Odometer}} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & C_n^b & \mathbf{0}_{3 \times 5} \end{bmatrix} \quad (8.4)$$

where C_n^b is the rotation matrix from the navigation (NED) frame to the body forward-right-down (FRD) frame. The rotation matrix is computed by assuming that pitch and roll are zero and applying an elementary rotation about the vertical (down) axis of the NED frame, or

$$C_n^b = \begin{bmatrix} \cos \hat{\psi} & \sin \hat{\psi} & 0 \\ -\sin \hat{\psi} & \cos \hat{\psi} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (8.5)$$

for estimated heading $\hat{\psi}$. Using this wheel speed mechanization technique requires making assumptions about lateral and vertical velocities, but utilizing an EKF measurement update in this way allows the EKF to account for the uncertainty in these assumptions. The associated measurement covariance for this odometry update consists of hand-tuned parameters given by

$$R_{\text{Odometer}} = \begin{bmatrix} 0.04 & 0 & 0 \\ 0 & 0.04 & 0 \\ 0 & 0 & 0.01 \end{bmatrix}. \quad (8.6)$$

The navigator design used for the UAV in these tests is unchanged from the design used in simulation, save for a re-tuned continuous process noise model,

$$Q_{\text{Tarrot}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1.21 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1.21 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.0025 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 30.5 \times 10^{-9} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 30.5 \times 10^{-9} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 305 \times 10^{-6} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.01 \end{bmatrix}. \quad (8.7)$$

The process noise model for auxiliary states used in the ICI navigator is also changed relative to that used in the simulation environment. The process noise model used depends on the peer vehicle's type (car or UAV), not the local vehicle's type. The models are given by

$$Q_{\text{Aux,Car}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.01 \end{bmatrix} \quad (8.8)$$

and

$$Q_{\text{Aux,UAV}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.1 \end{bmatrix} \quad (8.9)$$

for a car and UAV peer vehicle, respectively.

For the measurement updates integrating UWB range measurements, the UWB unit reports an estimated 1σ uncertainty value. This reported uncertainty is combined with residual monitoring and a tunable constant scale factor to obtain the collaborative measurement covariance for the ICI navigator,

$$R_{\text{UWB,ICI}} = \max\left(\sigma_{\text{UWB}}^2, \left(\frac{r}{2}\right)^2\right) * W_{\text{ICI,NCAT}} \quad (8.10)$$

for range residual r computed during the measurement update and constant weight factor $W_{\text{ICI,NCAT}} = 4.0$. This method of computing R_{UWB} involving a $\max()$ function is used to de-weight faulty measurements obtained from the UWB unit. Applying the same residual-based technique to the CI-based navigator did not have a noticeable impact on the observed position solution errors, so the form of the collaborative measurement covariance used with CI is unchanged from (6.9),

$$R_{\text{UWB,CI}}^* = (\sigma_{\text{UWB}}^2 + (-H)P_{\text{Peer}}(-H)^T) W_{\text{CI,NCAT}} \quad (8.11)$$

with re-tuned weight factor $W_{\text{CI,NCAT}} = 1.0$.

8.2 Test Scenario

Testing with four UAVs in a similar manner to the simulated scenarios previously discussed was not feasible due a lack of both four UAVs as well as — more importantly — four UAV operators. A test scenario was developed using the vehicles that were available. Two cars and one octo-rotor UAV were outfitted with Time Domain P440 UWB transceivers which measured the range between each pair of vehicles. Note that at the time of writing, the company which produced the UWB units, Time Domain, is defunct, but new company called TDSR sells UWB transceivers with claim similar performance characteristics [60].

The cars consisted of a Lincoln MKZ and a Kia Optima — henceforth referred to as “MKZ” and “Kia”, respectively. Both cars are equipped with high-quality GPS receivers which produce position estimates that include real-time kinematic (RTK) corrections. The position estimates produced by these receivers are nominally accurate to within a few centimeters, suitable for use as truth values for the purposes of this work. RTK observables are not provided to the navigators under test. The UWB units and GPS antennae were attached to each car’s roof as shown in Figures 8.1 and 8.2.

The UAV is a Tarrot X8 octo-rotor unit. Like the two cars, it includes an RTK-enabled GPS receiver which produces high-quality position estimates used as truth values, which again are not provided to the navigators under test. Dead reckoning was performed using the UAV’s onboard three-axis IMU. The equipment carried onboard the UAV is shown in Figure 8.3, and all three test vehicles are shown together in Figure 8.4.

The test route consisted of the two cars beginning at opposite ends of the skid pad, driving towards each other, circling each other for a turn and a half, then driving away from each other such that each car ends up at the opposite end of the skid pad compared to where it started. Meanwhile the UAV hovered roughly 20 m above the ground and either held stationary or circled. The view from the driver’s seat of the Kia at the beginning of one of the test runs is shown in Figure 8.5. RTK truth trajectories for the stationary and circling



Figure 8.1: MKZ roof-mounted test equipment.




Figure 8.2: Kia roof-mounted test equipment.




Figure 8.3: UAV-mounted test equipment.



Figure 8.4: Test vehicles during final checkouts before testing.

UAV scenarios are shown in Figures 8.6 and 8.7, respectively. For sake of clarity, Section 8.3 only includes results from the trajectory with the stationary UAV. Corresponding results for the trajectory with the circling UAV are given in Appendix B. The additional results do not lead to substantially different conclusions.



Figure 8.5: View from the driver’s seat of the Kia before starting a test run.

During testing, each vehicle logged its respective RTK GPS solution, UWB range measurements, and dead reckoning measurements (6-DOF IMU or yaw-gyro and wheel speed, as appropriate). The integration of real-time SOOP measurement generation or a software defined radio (SDR) SOOP receiver are beyond the scope of this work, so the data collected are combined with the same LEO Doppler measurement simulation techniques used throughout this work. All position estimation, including the integration of recorded UWB range measurements, was completed in post-processing. Results are presented for both trajectories (stationary and circling UAV), three collaborative modes (CI, ICI, and no collaboration), and two SOOP modes (on and off), giving a total of twelve scenarios. For all ICI scenarios, ICI fusion updates are conducted at 5 second intervals (0.2 Hz). Recall that in the ICI navigator, the ICI fusion update is its own distinct operation separate from the peer-to-peer range measurement update (see Section 6.2.3).

Stationary UAV Ground Tracks




Figure 8.6: Ground track of test run with a stationary hovering UAV.

Circling UAV Ground Tracks




Figure 8.7: Ground track of test run with a circling UAV.

8.3 Results

8.3.1 With SOOPs

The position error magnitudes for each vehicle with LEO Doppler measurements and no collaborative ranging are shown in Figure 8.8. Since the error magnitudes do not substantially differ from one vehicle to the next in this case, the RMS error across the three vehicles can be examined to more easily compare collaborative methods. Such a comparison is shown in Figure 8.9. Note that in this and similar plots the term “swarm” is used to denote that the error describes an overall value for the group of three vehicles. More detailed results can be found in Appendix B.

Interestingly, the addition of collaborative range measurements in this instance results in no substantial gain in positioning performance. In fact, according to RMS error values, the addition of collaborative ranging slightly *degrades* performance in both of the collaborative scenarios in Figure 8.9. The reason this has happened is that the range measurements provided by the UWB transceivers were *markedly* worse than the UWBs’ own reported standard deviation estimates would indicate. To demonstrate this, error of the range measurements reported by the UWB units (using range computed from RTK GPS positions as truth) are shown in Figure 8.10. These results show that the range measurement error is often on the order of tens of meters, while the reported confidence estimate is much smaller. Not only were the UWB transceivers performing worse than their specsheet would suggest, but the units failed to properly report the degradation in performance in all but a few measurement epochs. To avoid drastic negative performance impacts from these measurements, the ICI navigator employs the residual-based fault de-weighting scheme described in (8.10). The same approach did not improve the performance of the CI navigator, and so was not used.




Figure 8.8: Position error magnitude with no collaboration.

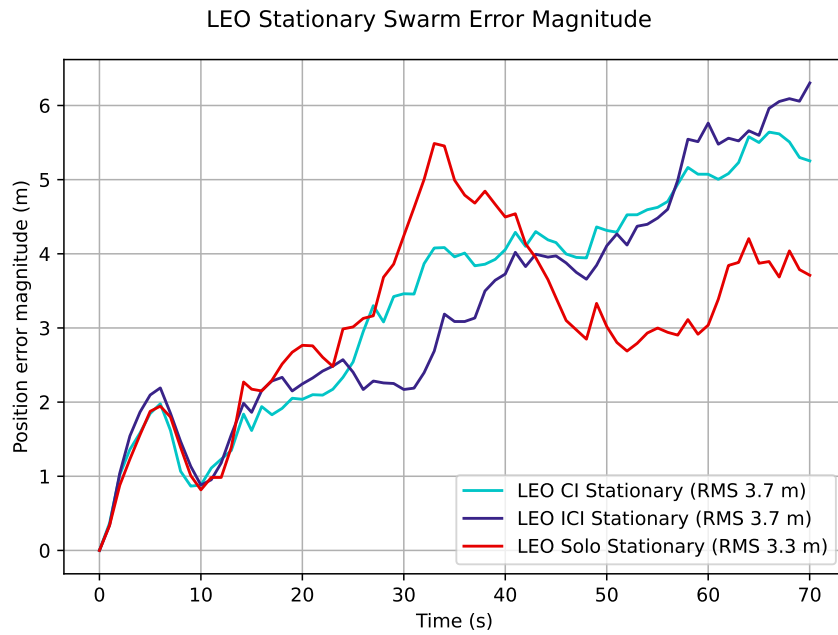


Figure 8.9: Overall position error magnitude for various collaborative methods.




Figure 8.10: UWB measurement error.

8.3.2 Without SOOPs

One important difference between the results in this chapter and those of the simulated UAVs is that ground vehicle dead reckoning using wheel speed and a single-axis gyro performs *much* better than a full 6-DOF IMU mechanization (see Figure 7.4). Figures 8.11 and 8.12 show the true and estimated ground tracks for the MKZ and Kia (respectively) without any positioning information other than wheel speed and yaw rate. (Note that the truth data in Figure 8.11 appears as a series of discrete points. This is not an error, and is merely a byproduct of the MKZ’s onboard GPS receiver having a slower reporting rate compared to that of the Kia.) Each car completes the run with an RMS position error magnitude of 2.3 m. The ground track for the UAV under the same conditions is shown in Figure 8.13 (note the wider zoom level compared to Figures 8.11 and 8.12). The position errors for the three vehicles in dead reckoned conditions are shown in Figure 8.14, which indicates an RMS position error magnitude of 115.7 m for the UAV. Given the excellent positioning

performance of the cars even with *no* RF measurements, tests were conducted with CI and ICI on the test track datasets with LEO Doppler measurements disabled.

Position error magnitudes without SOOP navigation are shown in Figure 8.15. Both collaborative methods are able to substantially reduce the position error of the UAV. The per-vehicle position errors for the stationary trajectory are shown in Figures 8.14, 8.16, and 8.17 for the solo, CI, and ICI cases, respectively. These results show that the improvement to the UAV’s position solution does not come at the expense of adversely affecting the positioning performance of the cars. Both collaborative techniques are able to effectively reduce the UAV’s RMS position error compared to the dead reckoned case, with reductions of 84% for CI and 83% for ICI. Even with peer-to-peer ranges that have abnormally high measurement error, collaborative navigation can improve position estimation quality. Similarly to the previously discussed test track results with LEO measurements, a more comprehensive set of results can be found in Appendix B.

8.3.3 Simulated UWB Measurements

Given the poor quality of the range measurements obtained from the UWB hardware, a selection of the previously discussed tests were re-run with the range measurements simulated based on the RTK true position data. The results of doing so are shown in Figure 8.18 for the scenarios with LEO SOOPs and in Figure 8.19 for the scenarios without. Comparing Figure 8.18 to Figure 8.9 shows that replacing the faulty UWB measurements with simulated ones allows for both the CI and ICI methods to improve the overall positioning performance in the presence of SOOP measurements. Comparing Figure 8.19 to Figure 8.15 results in a similar improvement over the UWB-derived ranges for the cases with SOOP measurements disabled. The performance gains caused by switching to simulated UWB measurements remain modest, however. Achieving greater performance improvements is hampered by two effects.

- The RTK GPS position measurement epochs are not mutually synchronized across the three vehicles. Consequently, the simulated UWB measurements are in some cases using out-of-date position information, which, for the speeds and measurement frequencies in this dataset, can introduce errors on the order of a few meters.
- Having only three vehicles, two of which are always on the ground, limits the geometric diversity available to the collaborative navigation algorithms. The effect of relative geometry on collaborative navigation performance is discussed in detail in [58].




Figure 8.11: MKZ dead reckoned ground track.




Figure 8.12: Kia dead reckoned ground track.




Figure 8.13: UAV dead reckoned ground track.




Figure 8.14: Position error magnitude with no collaboration and no SOOP measurements.




Figure 8.15: Overall position error magnitude with no SOOP measurements.




Figure 8.16: Position error magnitude with CI and no SOOP measurements.




Figure 8.17: Position error magnitude with ICI and no SOOP measurements.




Figure 8.18: Overall position error magnitude with simulated UWB ranges.




Figure 8.19: Overall position error magnitude with simulated UWB ranges and no SOOP measurements.

Chapter 9

Conclusions and Future Work

9.1 Summary

This thesis introduced the concept of opportunistic radionavigation using signals available from LEO satellite communications networks (Chapter 2). In particular, the use of opportunistic Doppler measurements obtained from said LEO signals was introduced as a workable alternative to the more common TOA measurements for use in a tightly-coupled RF-INS EKF navigator (Chapter 3). Performance heuristics were developed in Chapter 4 to study how the expected positioning performance of LEO Doppler measurements changes for different times and locations on the earth. In Chapter 5, the efficacy of a Doppler-based navigator was examined more directly using Monte Carlo analysis of simulated navigators across a selection of time and location conditions for the signals of interest. Collaborative peer-to-peer ranging was introduced in Chapter 6 as a way to further aid navigation with a selection of well-known and novel approaches presented. These collaborative navigation methods were simulated in a similar manner to the non-cooperating Doppler navigators to study the effectiveness of each method, again leveraging Monte Carlo analysis techniques (Chapter 7). The decentralized collaborative navigation schemes were tested experimentally using a set of two cars and one UAV obtaining range measurements using UWB transceivers in Chapter 8.

This chapter presents conclusions for three broad segments of this work. Section 9.2 discusses SOOP Doppler navigation in the absence of collaborative techniques, Section 9.3 examines the simulated experiments involving collaborative navigation, and Section 9.4 discusses the experimental results. Section 9.5 presents a number of potential avenues for future study.

9.2 Opportunistic Doppler Navigation

This work introduces the Doppler geometry index (DGI) as a way of examining the suitability of a given satellite geometry to Doppler navigation in the absence of a straightforward equivalent to the DOP parameter used for GNSSs. The performance results in Chapter 5 generally agree with the performance expectations set in Chapter 4, which indicates that the DGI is a useful indicator of expected Doppler navigation performance, at least for LEO emitters.

The simulation results in Chapter 5 also demonstrate that opportunistic Doppler navigation is an effective means of constraining error growth in the absence of GNSS signals. Even at the equator, where heuristics suggest that LEO Doppler navigation quality will be poor, the addition of SOOP observables results in a modest reduction in position error magnitude. Away from the equator the effect is more pronounced, with little error growth whatsoever during the course of the 100-second run. The position error growth is *not* bounded, however, and the observed error would continue to grow larger for longer scenarios. This means that SOOP Doppler navigation using the constellations of interest is not by itself a workable navigation strategy in the general case.

9.3 Collaborative Navigation Simulation Results

The collaborative navigation simulation study in Chapter 7 shows that the introduction of peer-to-peer range measurements is effective for reducing position estimate error. In the best

case — i.e. using a centralized navigator — the mean position error magnitude is reduced by roughly half compared to an identical scenario without collaborative ranging. For the scenarios away from the equator, the proposed ICI architecture demonstrated a reduction in position error magnitude nearly as good as that of the centralized navigator, albeit with the navigators producing less accurate covariance estimates than the centralized approach. At the equator, the ICI navigator was outperformed by the simpler and more well-studied CI navigator, though both decentralized methods reduced position error magnitude to some extent. These results show that the hybridized architecture of the ICI navigators can be a superior alternative to direct application of CI in some but not all cases. Each of the two methods carries a similar communication load, but the ICI navigator is more computationally expensive.

9.4 Collaborative Navigation Experimental Results

The experiments conducted at the NCAT test track allowed the two proposed decentralized collaborative techniques to be studied in the presence of range measurements from real hardware. The presence of measurement errors from the UWB units was substantially larger than expected. This degraded the performance of both techniques, requiring additional deweighting to be applied to nearly all range measurements. Despite this challenge, both collaborative techniques delivered usable navigation solutions and avoided producing substantially degraded navigation solutions in the presence of bad measurements.

The collaborative navigators *were*, however, still able to reduce the position estimate error for the cases without LEO Doppler measurements. Both the CI and ICI navigators reduced overall position error magnitudes by roughly 80%, relying on the excellent dead reckoning error characteristics of the cars to improve the position estimate quality of the UAV even in the presence of worse-than-expected range measurements.

9.5 Future Work

A selection of opportunities for further research is presented here. Some of these items are extensions of the concepts discussed in this work, while others regard alternative approaches to the problems of interest.

- Test the proposed SOOP navigators with a higher-fidelity measurement model which more closely approximates the measurement errors that would be expected from real measurements.
- Integrate the proposed SOOP navigators with a software defined receiver (SDR) to test with measurements collected from live sky.
- Examine the effects of adding terrestrial SOOP sources such as cellular signals in various geometric configurations.
- Improve the data collection architecture to decrease the effects of time synchronization offsets.
- Implement the proposed navigators in a real-time environment.
- Study the impacts of auxiliary state propagation methods other than constant velocity. The addition of yaw rate or steer angle sharing among ground vehicles is particularly promising as an area of improvement.
- Investigate the scalability of the collaborative navigation schemes to larger numbers of collaborators.
- Employ heuristic measures such as DGI to dynamically adjust tuning, or to dynamically select the best collaborative navigation method on a per-scenario basis.
- Adapt the proposed navigators to functioning in non-fully connected, sparsely connected, or dynamically connected graph topologies.

- Investigate the effects of different trajectory shapes on collaborative navigation performance as studied in [58].
- Combine the hybridized navigator architecture proposed for use with ICI with the CI algorithm to more directly compare the performance differences between the two data fusion algorithms.

References

- [1] M. Jones, “Signals of opportunity: Holy Grail or a waste of time?” *GPS World*, Feb. 2018. [Online]. Available: <https://www.gpsworld.com/signals-of-opportunity-holy-grail-or-a-waste-of-time/>
- [2] T. D. Hall, “Radiolocation using AM broadcast signals,” Ph.D., Massachusetts Institute of Technology, Cambridge, MA, Sep. 2002.
- [3] M. Joerger, J. Neale, and B. Pervan, “Iridium/GPS carrier phase positioning and fault detection over wide areas,” in *Proceedings of the 22nd International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2009)*, 2009, pp. 1371–1385.
- [4] J. Khalife and Z. M. Kassas, “Characterization of sector clock biases in cellular CDMA systems,” in *Proceedings of the 29th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2016)*, 2016, pp. 2281–2285.
- [5] J. Khalife, K. Shamaei, and Z. M. Kassas, “A software-defined receiver architecture for cellular CDMA-based navigation,” in *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. Savannah, GA: IEEE, Apr. 2016, pp. 816–826. [Online]. Available: <http://ieeexplore.ieee.org/document/7479777/>

- [6] K. Shamaei, J. Khalife, and Z. M. Kassas, “Performance characterization of positioning in LTE systems,” in *Proceedings of the 29th international technical meeting of the satellite division of the institute of navigation (ION GNSS+ 2016)*, 2016, pp. 2262–2270.
- [7] —, “Ranging precision analysis of LTE signals,” in *2017 25th European Signal Processing Conference (EUSIPCO)*. Kos, Greece: IEEE, Aug. 2017, pp. 2719–2723. [Online]. Available: <http://ieeexplore.ieee.org/document/8081705/>
- [8] —, “Pseudorange and multipath analysis of positioning with LTE secondary synchronization signals,” in *2018 IEEE Wireless Communications and Networking Conference (WCNC)*. Barcelona: IEEE, Apr. 2018, pp. 1–6. [Online]. Available: <https://ieeexplore.ieee.org/document/8377438/>
- [9] C. T. Ardito, J. J. Morales, J. Khalife, A. Abdallah, Z. M. Kassas *et al.*, “Performance evaluation of navigation using LEO satellite signals with periodically transmitted satellite positions,” in *Proceedings of the 2019 International Technical Meeting of The Institute of Navigation*, 2019, pp. 306–318.
- [10] J. J. Morales, J. Khalife, U. S. Cruz, and Z. M. Kassas, “Orbit modeling for simultaneous tracking and navigation using LEO satellite signals,” in *Proceedings of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*, 2019, pp. 2090–2099.
- [11] M. Orabi, J. Khalife, and Z. M. Kassas, “Opportunistic navigation with Doppler measurements from Iridium Next and Orbcomm LEO satellites,” in *2021 IEEE Aerospace Conference (50100)*. Big Sky, MT, USA: IEEE, Mar. 2021, pp. 1–9. [Online]. Available: <https://ieeexplore.ieee.org/document/9438454/>
- [12] Z. Tan, H. Qin, L. Cong, and C. Zhao, “New method for positioning using Iridium satellite signals of opportunity,” *IEEE Access*, vol. 7, pp. 83 412–83 423, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8744228/>

- [13] J. Khalife, M. Neinavaie, and Z. M. Kassas, “Universal receiver architecture for blind navigation with partially known terrestrial and extraterrestrial signals of opportunity,” in *Proceedings of the 34th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2021)*, 2021, pp. 2201–2211.
- [14] A. Lehtinen, “Doppler positioning with GPS,” Master’s thesis, Tampere University of Technology, Tampere, Finland, Oct. 2001.
- [15] R. Sharma and C. Taylor, “Cooperative navigation of MAVs in GPS denied areas,” in *2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*. Seoul: IEEE, Aug. 2008, pp. 481–486. [Online]. Available: <http://ieeexplore.ieee.org/document/4648041/>
- [16] A. Chakraborty, C. N. Taylor, R. Sharma, and K. M. Brink, “Cooperative localization for fixed wing unmanned aerial vehicles,” in *2016 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. Savannah, GA: IEEE, Apr. 2016, pp. 106–117. [Online]. Available: <http://ieeexplore.ieee.org/document/7479689/>
- [17] H. Mokhtarzadeh, “Correlated-data fusion and cooperative aiding in GNSS-stressed or denied environments,” Ph.D., The University of Minnesota, Sep. 2014.
- [18] H. Mokhtarzadeh and D. Gebre-Egziabher, “Performance of networked dead reckoning navigation system,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 5, pp. 2539–2553, Oct. 2016, number: 5. [Online]. Available: <http://ieeexplore.ieee.org/document/7812894/>
- [19] B. Noack, J. Sijs, M. Reinhardt, and U. D. Hanebeck, “Decentralized data fusion with inverse covariance intersection,” *Automatica*, vol. 79, pp. 35–41, May 2017. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0005109817300298>

- [20] B. Noack, J. Sijs, and U. D. Hanebeck, “Inverse covariance intersection: new insights and properties,” in *2017 20th International Conference on Information Fusion (Fusion)*. Xi’an, China: IEEE, Jul. 2017, pp. 1–8. [Online]. Available: <http://ieeexplore.ieee.org/document/8009694/>
- [21] S. Y. Kim, C. H. Kang, and J. W. Song, “1-point RANSAC UKF with inverse covariance intersection for fault tolerance,” *Sensors*, vol. 20, no. 2, p. 353, Jan. 2020. [Online]. Available: <https://www.mdpi.com/1424-8220/20/2/353>
- [22] T. Sun and M. Xin, “Distributed estimation with iterative inverse covariance intersection,” *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 70, no. 4, pp. 1645–1649, Apr. 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/9964039/>
- [23] —, “Inverse-covariance-intersection-based distributed estimation and application in wireless sensor network,” *IEEE Transactions on Industrial Informatics*, pp. 1–12, 2023. [Online]. Available: <https://ieeexplore.ieee.org/document/10018527/>
- [24] W. J. Park and C. G. Park, “Distributed GM-CPHD filter based on generalized inverse covariance intersection,” *IEEE Access*, vol. 9, pp. 94 078–94 086, 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9468680/>
- [25] I. Catalano, H. Sier, X. Yu, J. P. Queralta, and T. Westerlund, “UAV tracking with solid-state Lidars: dynamic multi-frequency scan integration,” Apr. 2023, arXiv:2304.12125 [cs]. [Online]. Available: <http://arxiv.org/abs/2304.12125>
- [26] C. D. Moomaw and S. M. Martin, “Collaborative signal of opportunity Doppler navigation with inverse covariance intersection,” in *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*, 2022, pp. 2883–2894.

- [27] J.-H. Huh and K. Seo, “An indoor location-based control system using Bluetooth beacons for IoT systems,” *Sensors*, vol. 17, no. 12, p. 2917, Dec. 2017. [Online]. Available: <http://www.mdpi.com/1424-8220/17/12/2917>
- [28] C. Vicek, P. McLain, and M. Murphy, “GPS/dead reckoning for vehicle tracking in the “urban canyon” environment,” in *Proceedings of VNIS '93 - Vehicle Navigation and Information Systems Conference*. Ottawa, Ont., Canada: IEEE, 1993, pp. 461–34. [Online]. Available: <http://ieeexplore.ieee.org/document/585671/>
- [29] S. E. Preston and D. M. Bevly, “CSAC-aided GPS multipath mitigation,” in *Proceedings of the 46th Annual Precise Time and Time Interval Systems and Applications Meeting*, 2014, pp. 228–234.
- [30] C. Matyszczyk, “Truck driver has GPS jammer, accidentally jams Newark airport,” *CNET*, Aug. 2013. [Online]. Available: <https://www.cnet.com/culture/truck-driver-has-gps-jammer-accidentally-jams-newark-airport/>
- [31] S. Thompson, “Collaborative simultaneous tracking and navigation with low earth orbit satellite signals of opportunity and inertial navigation system,” Master’s thesis, Auburn University, Auburn, Alabama, Jul. 2021.
- [32] C. Huang, H. Qin, C. Zhao, and H. Liang, “Phase-time method: accurate Doppler measurement for Iridium NEXT signals,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 58, no. 6, pp. 5954–5962, Dec. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9792171/>
- [33] J. J. Morales, J. Khalife, A. A. Abdallah, C. T. Ardito, and Z. M. Kassas, “Inertial navigation system aiding with Orbcomm LEO satellite Doppler measurements,” in *Proceedings of the 31st International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2018)*, 2018, pp. 2718–2725.

- [34] N. Wolchover, “The Webb space telescope will rewrite cosmic history. If it works.” *Quanta Magazine*, Dec. 2022. [Online]. Available: <https://www.quantamagazine.org/why-nasas-james-webb-space-telescope-matters-so-much-20211203/#>
- [35] R. J. Danchik, “An overview of Transit development,” *JOHNS HOPKINS APL TECHNICAL DIGEST*, vol. 19, no. 1, 1998.
- [36] J. Hospodka, “Doppler shift satellite navigation - NAVSAT-TRANSIT and adherents,” *MAD - Magazine of Aviation Development*, vol. 1, no. 2, p. 11, Mar. 2013. [Online]. Available: <https://ojs.cvut.cz/ojs/index.php/mad/article/view/3620>
- [37] D. S. Ilcev, “Introduction to polar earth orbits (PEO) in the function of the satellite distress and safety systems (SDSS),” *Indonesian Journal of Electrical Engineering and Computer Science*, vol. 19, no. 1, p. 285, Jul. 2020. [Online]. Available: <http://ijeecs.iaescore.com/index.php/IJEECS/article/view/21457>
- [38] P. Groves, *Principles of GNSS, inertial, and multi-sensor integrated navigation systems*, 2nd ed. Artech House, 2013.
- [39] R. F. Stengel, *Optimal control and estimation*, ser. Dover books on advanced mathematics. New York: Dover Publications, 1994.
- [40] W.-H. Hsu and S.-S. Jan, “Assessment of using Doppler shift of LEO satellites to aid GPS positioning,” in *2014 IEEE/ION Position, Location and Navigation Symposium - PLANS 2014*. Monterey, CA, USA: IEEE, May 2014, pp. 1155–1161. [Online]. Available: <http://ieeexplore.ieee.org/document/6851486/>
- [41] J. Wood, S. Thompson, S. Martin, and D. Bevly, “GPS positioning in reduced coverage environments using batched Doppler and pseudorange measurements,” in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*.

- Portland, OR, USA: IEEE, Apr. 2020, pp. 915–924. [Online]. Available: <https://ieeexplore.ieee.org/document/9110186/>
- [42] S. Martin, “Closely coupled GPS/INS relative positioning for automated vehicle convoys,” Master’s thesis, Auburn University, May 2011.
- [43] R. Langley, “Dilution of precision,” *GPS World*, May 1999.
- [44] J. P. Snyder, “Map projections: a working manual,” United States Geological Survey, Washington, D.C., Professional Paper 1395, 1987.
- [45] M. L. Psiaki, “Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids,” *NAVIGATION*, vol. 68, no. 3, pp. 621–641, Sep. 2021. [Online]. Available: <https://navi.ion.org/content/68/3/621>
- [46] B. McLemore and M. L. Psiaki, “DOP analysis for a LEO navigation constellation that relies on Doppler shift and pseudorange,” in *Proceedings of the 35th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2022)*, 2022, pp. 2503–2524.
- [47] Aceinna, “gnss-ins-sim.” [Online]. Available: <https://github.com/Aceinna/gnss-ins-sim>
- [48] B. Rhodes, “Skyfield.” [Online]. Available: <https://github.com/skyfielders/python-skyfield>
- [49] D. Vallado and P. Crawford, “SGP4 orbit determination,” in *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*. Honolulu, Hawaii: American Institute of Aeronautics and Astronautics, Aug. 2008. [Online]. Available: <http://arc.aiaa.org/doi/10.2514/6.2008-6770>
- [50] C. T. Ardito, J. J. Morales, J. Khalife, A. A. Abdallah, and Z. M. Kassas, “Performance evaluation of navigation using LEO satellite signals with periodically

- transmitted satellite positions,” Reston, Virginia, Feb. 2019, pp. 306–318. [Online]. Available: <https://www.ion.org/publications/abstract.cfm?articleID=16743>
- [51] M. Kenny, “Ever wondered what is on the Orbcomm satellite downlink?” Dec. 2002. [Online]. Available: <http://mdkenny.customer.netspace.net.au/Orbcomm.pdf>
- [52] J. J. Morales, J. Khalife, U. S. Cruz, and Z. M. Kassas, “Orbit modeling for simultaneous tracking and navigation using LEO satellite signals,” Miami, Florida, Oct. 2019, pp. 2090–2099. [Online]. Available: <https://www.ion.org/publications/abstract.cfm?articleID=17029>
- [53] TDSR, “P440 UWB module data sheet / user guide.” [Online]. Available: <https://tdsr-uwb.com/wp-content/uploads/2021/03/320-0317G-P440-Data-Sheet-User-Guide.pdf>
- [54] B. Jones, “Collaborative architectures for relative position estimation of ground vehicles with UWB ranging and vehicle dynamic models,” Master’s thesis, Auburn University, Auburn, Alabama, Dec. 2022.
- [55] S. Julier and J. Uhlmann, “A non-divergent estimation algorithm in the presence of unknown correlations,” in *Proceedings of the 1997 American Control Conference (Cat. No.97CH36041)*. Albuquerque, NM, USA: IEEE, 1997, pp. 2369–2373 vol.4. [Online]. Available: <http://ieeexplore.ieee.org/document/609105/>
- [56] H. Mokhtarzadeh and D. Gebre-Egziabher, “Cooperative inertial navigation,” *Navigation*, vol. 61, no. 2, pp. 77–94, Jun. 2014, number: 2. [Online]. Available: <http://doi.wiley.com/10.1002/navi.60>
- [57] P. Arambel, C. Rago, and R. Mehra, “Covariance intersection algorithm for distributed spacecraft state estimation,” in *Proceedings of the 2001 American Control Conference*.

(*Cat. No.01CH37148*). Arlington, VA, USA: IEEE, 2001, pp. 4398–4403 vol.6.
[Online]. Available: <http://ieeexplore.ieee.org/document/945670/>

[58] J. Pryor, “Evaluation of cooperative navigation strategies with maneuvering UAVs,”
Master’s thesis, Auburn University, Auburn, Alabama, Aug. 2021.

[59] Auburn University, “About NCAT.” [Online]. Available: <https://eng.auburn.edu/research/centers/ncat/about/index.html>

[60] TDSR, “TDSR UWB technology.” [Online]. Available: <https://tdsr-uwb.com/technology/>

Appendices

Appendix A

Additional Monte Carlo Position Error Results

This appendix gives detailed position error results for all twelve 00:00 UTC navigator scenarios (three latitudes and four collaborative modes) as well as point-of-comparison results for navigators with GPS and with no satellite measurements at all. Some plots shown here are duplicates of those discussed in Chapters 5 and 7. They are repeated here to create a complete set of aggregated results.

GPS Results with GPS enabled, shown in Figures A.1, A.2, A.3, and A.4, demonstrate how the navigators behave in the presence of abundant navigation information. The resulting position errors are substantially less than what would be obtained from a real GPS receiver, which is a byproduct of the simplified measurement error models used in this work.

IMU Monte Carlo tests were also conducted with neither GPS nor LEO Doppler measurements. These results show how the various collaborative techniques behave with only each other's dead reckoned estimates. The corresponding position error results are shown in Figures A.5, A.6, A.7, and A.8.

LEO The remainder of the chapter is dedicated to the position error results across LEO Doppler scenarios. Equatorial results are given in Figures A.9, A.10, A.11, and A.12. Results from 30° N are given in Figures A.13, A.14, A.15, and A.16. Results from 60° N are given in Figures A.17, A.18, A.19, and A.20.




Figure A.1: Position error for a navigator using GPS and no collaboration.




Figure A.2: Position error for a navigator using GPS and CI.




Figure A.3: Position error for a navigator using GPS and ICI.




Figure A.4: Position error for a navigator using GPS and centralized collaboration.




Figure A.5: Position error for a navigator using IMU and no collaboration.




Figure A.6: Position error for a navigator using IMU and CI.




Figure A.7: Position error for a navigator using IMU and ICI.

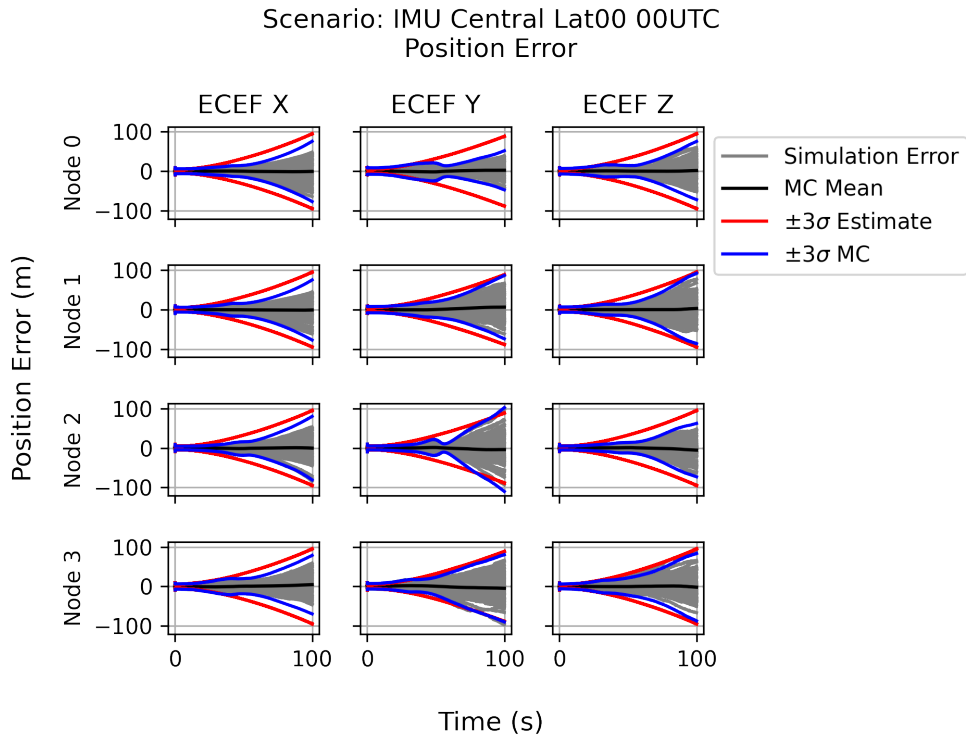


Figure A.8: Position error for a navigator using IMU and centralized collaboration.




Figure A.9: Position error for a navigator using LEO at the equator and no collaboration.

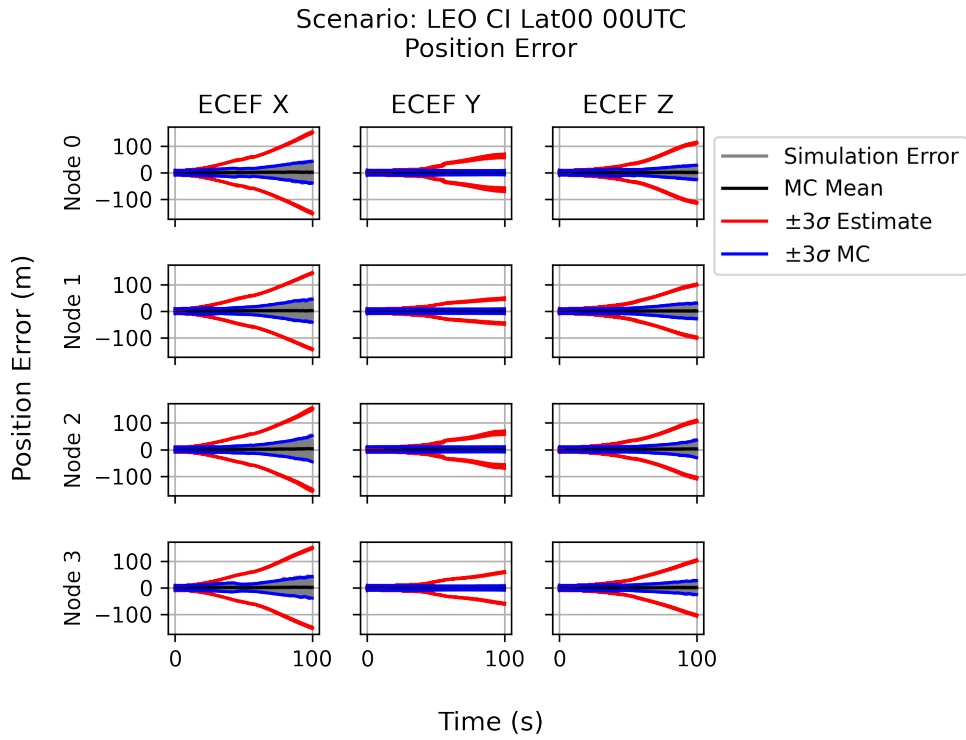


Figure A.10: Position error for a navigator using LEO at the equator and CI.




Figure A.11: Position error for a navigator using LEO at the equator and ICI.




Figure A.12: Position error for a navigator using LEO at the equator and centralized col-
laboration.




Figure A.13: Position error for a navigator using LEO at 30° N and no collaboration.




Figure A.14: Position error for a navigator using LEO at 30° N and CI.




Figure A.15: Position error for a navigator using LEO at 30° N and ICI.




Figure A.16: Position error for a navigator using LEO at 30° N and centralized collaboration.

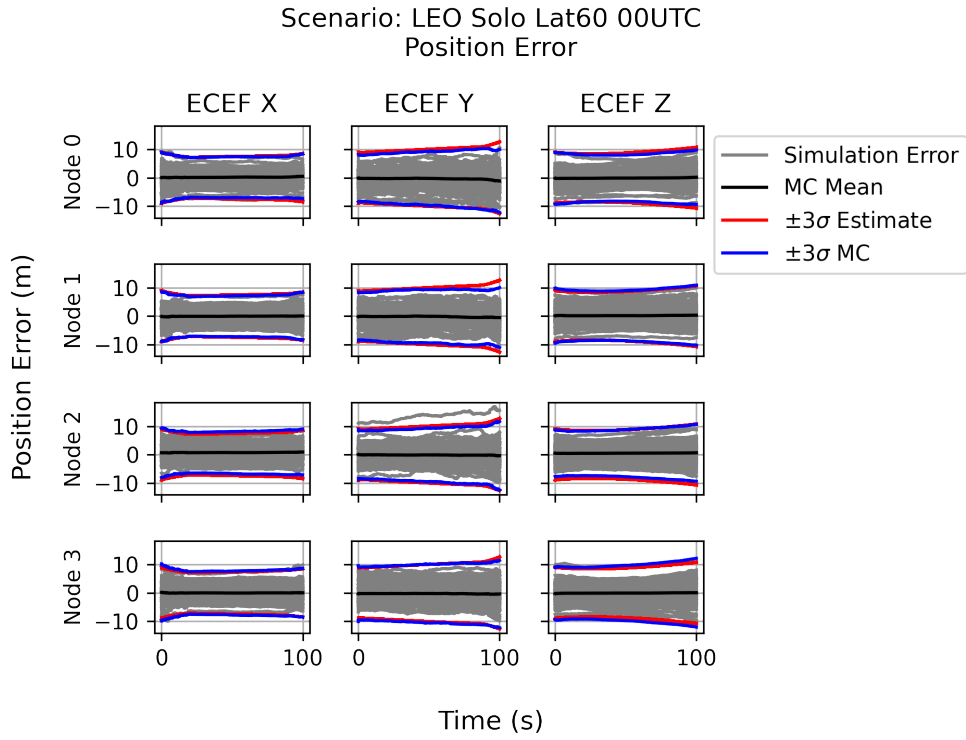


Figure A.17: Position error for a navigator using LEO at 60° N and no collaboration.




Figure A.18: Position error for a navigator using LEO at 60° N and CI.




Figure A.19: Position error for a navigator using LEO at 60° N and ICI.




Figure A.20: Position error for a navigator using LEO at 60° N and centralized collaboration.

Appendix B

Additional Test Track Results

This appendix includes a broad set of test results relating to the NCAT test runs discussed in Chapter 8. Section B.1 gives results corresponding exactly to those given in Chapter 8, but for the trajectory with the circling UAV rather than with the static UAV. Section B.2 gives additional supporting data for both trajectory types.

B.1 Circling Trajectory Results

Figures B.1 – B.10 mirror exactly those in Section 8.3, but using the circling rather than static trajectory. Corresponding plots appear in the same order, and each plot caption in this section includes a reference to its counterpart.

B.2 Additional Data

This section includes a selection of plots showing position and velocity estimate errors in each of the NED axes, along with the accompanying EKF estimated covariances. Such data demonstrates how well or how poorly the EKF is estimating its uncertainty. Note that y -axis scale varies from one plot to the next so that the maximum amount of detail may be shown for each scenario, so care must be taken when comparing across scenarios. Figures B.11 – B.16 describe scenarios with SOOP measurements, and Figures B.17 – B.22 describe those without. All results shown in this section use the stationary trajectory. Results using the

circling trajectory have been omitted for brevity and do not substantially differ. Figure B.23 shows one of the UWB units being removed from its mounting hardware; attempts to remove the unit non-destructively were unsuccessful.




Figure B.1: Position error magnitude with no collaboration (see Fig. 8.8).




Figure B.2: Overall position error magnitude for various collaborative methods (see Fig. 8.9).




Figure B.3: UWB measurement error (see Fig. 8.10).




Figure B.4: MKZ dead reckoned ground track (see Fig. 8.11).




Figure B.5: Kia dead reckoned ground track (see Fig. 8.12).




Figure B.6: UAV dead reckoned ground track (see Fig. 8.13).




Figure B.7: Position error magnitude with no collaboration and no SOOP measurements (see Fig. 8.14).




Figure B.8: Overall position error magnitude with no SOOP measurements (see Fig. 8.15).




Figure B.9: Position error magnitude with CI and no SOOP measurements (see Fig. 8.16).




Figure B.10: Position error magnitude with ICI and no SOOP measurements (see Fig. 8.17).




Figure B.11: NED position errors with LEO SOOPs and no collaboration.




Figure B.12: NED velocity errors with LEO SOOPs and no collaboration.




Figure B.13: NED position errors with LEO SOOPs and CI.




Figure B.14: NED velocity errors with LEO SOOPs and CI.




Figure B.15: NED position errors with LEO SOOPs and ICI.




Figure B.16: NED velocity errors with LEO SOOPs and ICI.




Figure B.17: NED position errors with no SOOPs and no collaboration.




Figure B.18: NED velocity errors with no SOOPs and no collaboration.




Figure B.19: NED position errors with no SOOPs and CI.




Figure B.20: NED velocity errors with no SOOPs and CI.




Figure B.21: NED position errors with no SOOPs and ICI.




Figure B.22: NED velocity errors with no SOOPs and ICI.



Figure B.23: Removing a UWB unit from its mounting hardware.

COLOPHON

This thesis was typeset with L^AT_EX 2_ε using the T_EX Live package running on Ubuntu 20.04.1. The body serif font is Computer Modern, the sans-serif font used in figures is DejaVu Sans, and the monospaced font is DejaVu Sans Mono. Writing and editing was done using GNU Emacs 28.2, with special thanks to the GNU Aspell package for finding and fixing innumreable spelling errors.

Affiliation Disclosure

The author's affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions, or viewpoints expressed by the author.