

**Autonomous Navigation for Distributed Space Systems via
Spacecraft to Spacecraft Absolute Tracking**

by

J. A. Greaves

B.A., Rensselaer Polytechnic Institute, 2018

M.S., University of Colorado Boulder, 2020

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Aerospace Engineering Sciences
2023

Committee Members:

Daniel J. Scheeres, Chair

Marcus Holzinger

Jay McMahan

Nisar Ahmed

David Woffinden

Greaves, J. A. (Ph.D., Aerospace Engineering Sciences)

Autonomous Navigation for Distributed Space Systems via Spacecraft to Spacecraft Absolute Tracking

Thesis directed by Prof. Daniel J. Scheeres

A distributed space system utilizes multiple spacecraft working in tandem to complete unified objectives. Distributed architectures are appealing because they can provide enhanced robustness for reduced costs over traditional monolithic spacecraft. Additionally, utilizing multiple vehicles allows for the capacity to fly sensor in large formations which can enable novel science. While distributed systems are promising, they require more complex planning to properly operate. Particularly, there is a need to navigate every vehicle in the system. Spacecraft navigation traditionally relies on ground-based sensor networks but scheduling observation time for each vehicle is not only costly, but potentially intractable. Thus, there is a need to examine autonomous navigation methods for distributed space systems.

In this dissertation we examine the feasibility of generating a full state estimate for every vehicle in a distributed system from relative measurements between vehicles in the system. We show that a full state estimate can be obtained due to nonlinearities in spacecraft motion, and therefore such a method could provide fully autonomous navigation. Still, even if the full state can be observed, there are regions of the state space which are slow to accrue information. To abate this, we develop guidance policies to minimize state uncertainty along a desired state projection. The policies are analytic and well suited for autonomous computation. Finally, we include methods to identify and classify unknown maneuvers on target vehicles. Having the capacity to capture and label unknown behaviors improves the robustness of the navigation solutions when communication or cooperation is not available. Overall, the combination of methods presented form a completely autonomous navigation method for distributed space systems which can improve its own state estimate and account for unmodeled events.

Dedication

To my family Amna, Paul, and Ian Greaves. Without their love and support none of this would be possible. To my friends from CU Boulder, Rensselaer Polytechnic Institute, and Andover. Without them this journey would not have been enjoyable. Lastly, to my teachers and mentors who gave me the skills and tools to complete this work. Without them I would not have made it this far.

Acknowledgements

First, I would like to thank my advisor Dr. Daniel Scheeres for taking me on as a student and guiding me throughout my time at Boulder. He afforded me the freedom to pursue my research interests and was instrumental in solving problems whenever they arose. I would also like to thank my committee Marcus Holzinger, Jay McMahon, Nisar Ahmed, and David Woffinden for their thoughtful feedback and mentorship which has helped shape this thesis.

Next, I must thank the group EG6 at Johnson Space Center and 392J at the Jet Propulsion Lab. 392J and EG6 brought me on as an intern and later as a NASA Pathways student so I could gain necessary experience and contribute to NASA's missions. At EG6 I would specifically like to thank my mentors Chris D'Souza, David Woffinden, Greg Holt, and Jack Brazzel. In 392J I want to thank Daniel Lubey and Shyam Bhaskaran. All of them have directly boosted my trajectory through life, imparted invaluable knowledge, and presented me with the greatest opportunities.

A highlight of my time at CU Boulder where the friends and connections I made along the way. From outdoors sports to cookouts and campouts, to critical research support and advice. The people I have met along the way have made this journey both feasible and fun. I know I will always look back at this period of my life with nothing but fond memories largely due to their influence on me.

Finally, much of this work was supported by the Air Force Research Lab. Thank you very much for funding my time, efforts, and travel while developing my work. They are the reason for the existence of this work.

Contents

| Chapter | |
|----------------|-----------|
| 1 | 1 |
| 1.1 | 4 |
| 1.2 | 5 |
| 1.3 | 6 |
| 1.3.1 | 6 |
| 1.3.2 | 7 |
| 2 | 9 |
| 2.1 | 10 |
| 2.1.1 | 10 |
| 2.1.2 | 13 |
| 2.1.3 | 15 |
| 2.2 | 19 |
| 2.2.1 | 19 |
| 2.2.2 | 20 |
| 2.2.3 | 21 |
| 3 | 23 |
| 3.1 | 23 |
| 3.2 | 25 |

| | | |
|----------|---|-----------|
| 3.2.1 | Spacecraft to Spacecraft Absolute Tracking Model | 25 |
| 3.2.2 | Linearized Observability Theory | 28 |
| 3.3 | Observability of Spacecraft to Spacecraft Absolute Tracking | 30 |
| 3.3.1 | Derivation of the Generic Observability Matrix | 30 |
| 3.3.2 | Generic Observability Conditions | 32 |
| 3.3.3 | Multiple Targets and Agent-Based Absolute State | 34 |
| 3.3.4 | Spaceflight Specific Modeling | 35 |
| 3.4 | Simulation and Results | 38 |
| 3.4.1 | SSAT Under Keplerian Motion | 39 |
| 3.4.2 | SSAT With J2 Perturbation | 44 |
| 3.4.3 | SSAT in the CR3BP | 48 |
| 3.4.4 | SSAT of a Multi-Target Constellation | 51 |
| 3.5 | Conclusion | 52 |
| 4 | Autonomous Guidance for Information Gathering with Optics | 54 |
| 4.1 | Introduction and Background on Information Gathering | 54 |
| 4.2 | Problem Formulation | 56 |
| 4.2.1 | Dynamical Models | 57 |
| 4.2.2 | Measurement Models | 58 |
| 4.2.3 | Uncertainty Cost Function | 60 |
| 4.2.4 | Angular Geometry and Variation | 61 |
| 4.3 | Information Gathering Guidance | 64 |
| 4.3.1 | Heuristic Guidance for Optical Sensors | 64 |
| 4.3.2 | Maximum Measurement Deviation Guidance | 68 |
| 4.3.3 | Combined Information Gathering and Station Keeping | 69 |
| 4.4 | Simulation and Results | 70 |
| 4.4.1 | Information Gathering Operation | 71 |

| | | |
|----------|---|------------|
| 4.4.2 | Informed Station Keeping Operation | 77 |
| 4.5 | Conclusion | 81 |
| 5 | Spacecraft to Spacecraft Absolute Tracking and Maneuver Classification in Cislunar Space | 83 |
| 5.1 | Introduction | 83 |
| 5.2 | Background Literature | 84 |
| 5.2.1 | Space-Based Optical Tracking in the Cislunar Environment | 84 |
| 5.2.2 | Maneuver Detection and Classification | 86 |
| 5.3 | Problem Formulation | 87 |
| 5.3.1 | Dynamical Models | 88 |
| 5.3.2 | Selected Orbits | 90 |
| 5.3.3 | Measurement Models | 92 |
| 5.4 | Optical-Only Spacecraft to Spacecraft Absolute Tracking in Cislunar Space | 93 |
| 5.4.1 | Observability Results | 93 |
| 5.4.2 | State Estimation Solutions | 95 |
| 5.5 | Maneuver Classification of Unmodeled Events | 106 |
| 5.5.1 | Classifier Input and Setup | 107 |
| 5.5.2 | Classification Algorithms | 108 |
| 5.5.3 | Maneuver Classification Results | 110 |
| 5.6 | Conclusion | 120 |
| | Bibliography | 124 |
| | Appendix | |
| A | Ballistic Linear OCBE for Spacecraft to Spacecraft Absolute Tracking | 131 |

| | | |
|----------|--|------------|
| B | Modifications for the ballistic OCBE | 135 |
| B.1 | Unscented Modification | 135 |
| B.2 | Square Root Information Modification | 137 |
| C | Space-Based Observation of Cislunar Vehicles with the Optimal Control Based Estimator | 140 |
| C.1 | Introduction | 140 |
| C.2 | Problem Formulation | 142 |
| C.2.1 | Selected Orbits | 142 |
| C.2.2 | Target Maneuvers | 143 |
| C.2.3 | Measurements | 144 |
| C.3 | Optical-Only Tracking with a Known Observer in Cislunar Space | 147 |
| C.3.1 | Observational Analysis | 148 |
| C.3.2 | Filter Tuning Strategy | 150 |
| C.4 | Maneuver Detection | 160 |
| C.4.1 | Hypothesis Testing Using Optimal Control Policies | 160 |
| C.4.2 | Control Policy Distributions | 161 |
| C.4.3 | Maneuver Detection | 166 |
| C.4.4 | Control as a Health Metric | 167 |
| C.5 | Conclusion | 167 |

Tables

Table

| | | |
|-----|---|-----|
| 3.1 | SSAT Uncertainty Ranges | 44 |
| 3.2 | Multi-target SSAT Uncertainties | 53 |
| 5.1 | Target Maneuver Statistics | 90 |
| 5.2 | Selected Orbit Parameters for Chapter 5 | 90 |
| 5.3 | Time to Observability Results | 95 |
| 5.4 | Filter Consistency Sans Maneuvers | 97 |
| 5.5 | Filter Consistency with Maneuvers | 105 |
| C.1 | Selected Target Orbit Parameters | 142 |
| C.2 | Station Keeping Statistics | 144 |
| C.3 | Measurement Statistics | 146 |
| C.4 | Idealistic Filter Uncertainty | 150 |
| C.5 | Expected Maneuver Detection from Hypothesis Test on Large Maneuvers | 162 |
| C.6 | Expected Maneuver Detection from Hypothesis Test on Small Maneuvers | 164 |
| C.7 | Maneuver Detection Results for 9:2 NRHO | 166 |
| C.8 | Maneuver Detection Results for Stable NRHO | 167 |

Figures

Figure

| | | |
|------|--|----|
| 2.1 | Two Body Problem Diagram | 10 |
| 2.2 | Circular Restricted Three Body Diagram | 13 |
| 2.3 | Circular Restricted Three Body Orbit Families | 16 |
| 2.4 | Bounded Motion of a QPO about a NRHO | 16 |
| 3.1 | STT Approximation Error of the SSAT Observability Matrix | 40 |
| 3.2 | Keplerian SSAT $\delta(e, i)$ Navigation Map | 41 |
| 3.3 | Keplerian SSAT $\delta(a, i)$ Navigation Map | 42 |
| 3.4 | Keplerian SSAT $\delta(a, e)$ Navigation Map | 43 |
| 3.5 | J2 Perturbed Keplerian SSAT $\delta(e, i)$ Navigation Map | 46 |
| 3.6 | J2 Perturbed Keplerian SSAT $\delta(a, i)$ Navigation Map | 46 |
| 3.7 | J2 Perturbed Keplerian SSAT $\delta(a, e)$ Navigation Map | 47 |
| 3.8 | CR3BP SSAT Navigation Map With Gateway to Southern L2 Orbits | 49 |
| 3.9 | CR3BP SSAT Navigation Map With Gateway to Northern L1 Orbits | 50 |
| 3.10 | CR3BP SSAT Navigation Map With Gateway to Planar Lyapunov Orbits | 50 |
| 3.11 | Multi-Target SSAT Geometry | 52 |
| 4.1 | Information Gathering Maneuver Geometry | 62 |
| 4.2 | Optimal Displacement Angle to Minimize Range Uncertainty | 65 |
| 4.3 | Information Gathering Operation Uncertainty Results | 72 |

| | | |
|------|--|-----|
| 4.4 | Information Gathering Operation Fuel Results | 74 |
| 4.5 | Information Gathering Operation Timeline | 75 |
| 4.6 | Information Gathering Operation Pareto Front | 76 |
| 4.7 | QPO Target Section for Station Keeping | 78 |
| 4.8 | Station Keeping Maneuver Surface | 79 |
| 4.9 | Informed Station Keeping Uncertainty Results | 79 |
| 4.10 | Informed Station Keeping Fuel Results | 80 |
| 5.1 | Selected Orbits for Chapter 5 | 91 |
| 5.2 | UKF Performance for NL1 Agent Sans Maneuver | 97 |
| 5.3 | UKF Performance for PL2 Agent Sans Maneuver | 98 |
| 5.4 | UKF Performance for QPO Agent Sans Maneuver | 98 |
| 5.5 | Smoothed U-OCBE Performance for NL1 Agent with Maneuver | 101 |
| 5.6 | Smoothed U-OCBE Outputs for NL1 Agent with Maneuver | 101 |
| 5.7 | Smoothed U-OCBE Performance for PL2 Agent with Maneuver | 102 |
| 5.8 | Smoothed U-OCBE Outputs for PL2 Agent with Maneuver | 103 |
| 5.9 | Smoothed U-OCBE Performance for QPO Agent with Maneuver | 104 |
| 5.10 | Smoothed U-OCBE Outputs for QPO Agent with Maneuver | 104 |
| 5.11 | True and Estimated Control Vectors for QPO Agent and Noisy Maneuvers | 111 |
| 5.12 | Estimated Control Magnitude for QPO Agent and Noisy Maneuvers | 111 |
| 5.13 | True and Estimated Control Vectors for QPO Agent and Quiet Maneuvers | 112 |
| 5.14 | Estimated Control Magnitude for QPO Agent and Quiet Maneuvers | 112 |
| 5.15 | Estimated Control Vectors for PL2 Agent | 113 |
| 5.16 | Estimated Control Magnitudes for PL2 Agent | 114 |
| 5.17 | Estimated Control Vectors for NL1 Agent | 115 |
| 5.18 | Estimated Control Magnitudes for NL1 Agent | 116 |
| 5.19 | Confusion Matrix Results for QPO Agent | 117 |

| | | |
|------|--|-----|
| 5.20 | Confusion Matrix Results for PL2 Agent | 118 |
| 5.21 | Confusion Matrix Results for NL1 Agent | 118 |
| C.1 | Selected Target Orbits | 143 |
| C.2 | Observer and Target Locations | 147 |
| C.3 | Filter Failure Sans Angle-Rate Measurements | 149 |
| C.4 | Filter Success with Angle-Rate Measurements | 151 |
| C.5 | Adaptive SRI-OCBE with a Large Maneuver | 152 |
| C.6 | Adaptive SRI-OCBE Control and Metric with a Large Maneuver | 153 |
| C.7 | Adaptive SRI-OCBE with a Small Maneuver | 154 |
| C.8 | Adaptive SRI-OCBE Control and Metric with a Small Maneuver | 155 |
| C.9 | Constant Tuning SRI-OCBE with a Small Maneuver | 156 |
| C.10 | Constant Tuning SRI-OCBE Control and Metric with a Small Maneuver | 157 |
| C.11 | Constant Tuning SRI-OCBE for Station Keeping | 158 |
| C.12 | Constant Tuning SRI-OCBE Control and Metric for Station Keeping | 159 |
| C.13 | Histogram of \mathcal{Z} for Large Maneuvers | 161 |
| C.14 | PDF approximation of \mathcal{Z} for Large Maneuvers | 162 |
| C.15 | Estimated Control Magnitudes and True Control Pointing for Large Maneuvers . . . | 163 |
| C.16 | Estimated Control Magnitudes and True Control Magnitude for Large Maneuvers . . | 164 |
| C.17 | Histogram of \mathcal{Z} for Large Maneuvers | 164 |
| C.18 | PDF approximation of \mathcal{Z} for Small Maneuvers | 165 |
| C.19 | Estimated Control Magnitudes and True Control Pointing for Small Maneuvers . . . | 165 |
| C.20 | Estimated Control Magnitudes and True Control Magnitude for Small Maneuvers . . | 165 |
| C.21 | Filter Failure with Normal Residuals | 168 |

Chapter 1

Introduction

Autonomous distributed space systems (DSS's) are currently garnering significant attention for their potential to accomplish missions with lower costs, a higher degree of reliability, and reduced loads on traditional external resources. Furthermore, and perhaps even more appealing, DSS architectures present novel capabilities which cannot be achieved with a single craft such as: global communications coverage, precise gravity field reconstruction, and solar interferometry [1, 2, 3]. While the previous examples were cooperative in scope, DSS techniques may also be applied in non-cooperative settings. For example, space debris removal or defunct satellite repair both consider operations with multiple space objects and utilize many of the same fundamentals as a DSS, excluding cooperative communication and planning of course. Moreover, equipping a spacecraft with the capacity to consider additional objects in its vicinity is a versatile tool which is applicable to many important scenarios ranging from relative navigation to collision avoidance.

Even with the great potential that DSS's offer, they have yet to be widely adopted. Their slow adoption is likely due to their complex nature and heavy resource demand. One particularly difficult requirement is the need to provide orbit determination for every vehicle in the system. Calculating an estimated orbit for a single vehicle is not technically challenging since there are many methods to obtain a solution, but there are often challenges when it comes to scheduling observational resources since ground-based observers are heavily subscribed. The scheduling problem is further exacerbated when tracking an entire DSS, let alone a single vehicle. Thus, the primary motivation for this thesis: providing autonomous navigation for distributed space systems.

To achieve the goal of autonomous navigation, three research objectives were sought. First, proof that relative sensors onboard DSS can produce a fully observable state space for all vehicles in the system is proffered. Second, guidance to improve navigation estimates is developed since natural motion and sensing may not always provide sufficient accuracy. Finally, methods to identify and classify unknown perturbation are implemented such that an autonomous DSS can be robust to unmodeled events. With the aim of autonomy in mind, the developed methods are required to only use onboard sensors and attempt to limit computational demand. These objectives are summarized by the thesis statement of this work.

Thesis Statement

Autonomous navigation methods for distributed space systems have the potential to enable and enhance spaceflight missions while reducing dependencies and requirements. To achieve further levels of autonomy there is a need for low computation cost navigation techniques that are fully observable with onboard-only sensors, incorporate guidance to improve state estimation, and have the capacity to identify and classify unmodeled events.

Achieving the three outlined goals means that a DSS could robustly and autonomously navigate itself without requiring any external resources. Such a paradigm presents numerous benefits such as alleviating sensor network loads, enabling real-time operations, and increasing navigation robustness. Moreover, a fully autonomous system could function in conditions where communications are not readily available, which could facilitate novel mission scenarios. The first step to advance navigation autonomy is to prove that a DSS can use onboard-only sensors to fully reconstruct a state estimate for every vehicle in the system.

The use of onboard-only sensors to solve for the state of all vehicles in a distributed system has a wide set of background literature. In the robotics and estimation communities this topic is commonly covered under simultaneous localization and mapping (SLAM) and distributed/networked

multi-agent localization. SLAM tends to focus on mapping an unknown environment while then localizing an agent relatively within that model. Distributed/networked multi-agent localization is primarily concerned with generating relative estimates between vehicles in a distributed system to empower formation control. This is in no way a comprehensive review of these topics, but an important feature of both solutions is that they require an absolute sensor, such as GPS, if global localization is desired. Obtaining absolute information for a space system greatly hinders system autonomy because such information generally requires external resources, and GPS is not available for cislunar or deep space activities.

Fortunately, it has been shown that relative-only measurements can produce both local and global information for a DSS by estimating the state of all vehicles in the system [4, 5, 6, 7]. This is synonymous with obtaining an estimate for the absolute state along with all relative states for the distributed system. This is possible for space systems because the dynamics of space flight are predictable and nonlinear in a way which they can be used to effectively measure the acceleration field driving the system's motion [8]. It was speculated that origin of absolute state observability is due to experiencing unique gravity forces in the two body problem, and asymmetry in the three body problem [4, 5]. Chapter 3 will show by derivation that, no matter the dynamical system, observability is obtained due to differential motion between vehicles. The resulting state estimate generated from estimating all vehicles is therefore entirely dependent on the configuration of the system, with some configurations being unobservable while others are highly accurate.

Using relative-only measurements from onboard sensors for a navigation solution is appealing because it only requires an a-priori estimate. But it is important to acknowledge that limiting data sources can lead to a limited state reconstruction. Single sensor setups often suffer from limited information, so guidance algorithms have been developed to enhance state estimates to abate this issue [9, 10, 11]. Angle-only sensors have been particularly well studied for space systems because they represent a common measurement type that is often easy to implement [12, 13, 14]. Chapter 4 builds on previous angle-only guidance techniques by studying the geometric nature of information gathering for angular measurements.

For autonomous navigation algorithms to be trusted then they must show that they can maintain a convergent state solution even under unknown perturbations. This is particularly concerning for a DSS in cislunar space because many trajectories in the region require frequent station keeping maneuvers due to the strongly nonlinear dynamics [15, 16]. These dynamics also cause filtering difficulties because the high-order moments are non-negligible in some cases and can cause linearized filters to diverge [17, 15, 18]. Therefore, cislunar navigation methods must consider maneuver detection and high-order filtering to ensure proper tracking. Numerous methods are available for maneuver detection and high-order filtering, but they often come at the cost of additional computational requirements. The optimal control based estimator (OCBE) and unscented Kalman filter (UKF) address each concern respectively and require minimal additional computation cost compared to alternative methods [19, 20]. Chapter 5 presents a maneuver classification method based on a modified OCBE with UKF updates, called the U-OCBE, in conjunction with classification algorithms.

1.1 Dissertation Overview

The next chapter lays out common background material that will appear in all the following chapters and contributions. The common material covers spacecraft dynamics and simulated measurements. Then contributions to three main research areas are made over Chapters 3-5. Each chapter covers a unique topic, so they are introduced individually with its own literature review, problem formulation, contributed solution, and simulations.

Chapter 3 contains an analytic derivation that proves relative measurements can produce absolute and relative information to fully estimate a distributed system without the need for external observation. The derivation assumes a single observing agent and observed target pair but is trivially expanded to a multi-target system. The navigation technique laid out by the derivation is called spacecraft-to-spacecraft absolute tracking (SSAT). SSAT is a solution to autonomous navigation of a DSS from onboard relative measurements, and it is used throughout the remainder of the thesis. Simulated results are computed to demonstrate potential applications of SSAT, and

the resulting analysis is useful to identify where the technique excels or falters.

Chapter 4 builds upon SSAT by introducing maneuvers to enhance the system state estimate. The maneuvers in Chapter 4 are planned and executed by the observing agent to minimize a desired cost function which represents a projection of the state uncertainty matrix. The work in this chapter focuses optical measurements and chooses to minimize uncertainty that is orthogonal to the measurement space because the orthogonal space is often the most difficult to estimate. The work begins with a geometric analysis which leads to a heuristic policy. Lessons from the heuristic are used to define a surrogate cost function which is solved analytically. The analytic policy is easy to calculate and closely approximates an optimal solution making it ideal for online guidance.

Finally, Chapter 5 incorporates unknown disturbances and uses estimation techniques in combination with machine learning to reconstruct and label the unknown events. The unknown event reconstruction is carried out by the optimal control based estimator (OCBE). To effectively apply the OCBE to the SSAT system the ballistic version of the filter is modified to take on non-linear updates. The reconstructed events are labeled by various supervised learning algorithms to compare speed, accuracy, and simplicity. The event reconstruction and classification add robustness to the autonomous navigation architecture by capturing unmodeled events so they do not unknowingly cause tracking divergence.

1.2 Contributions

The primary contribution in Chapter 3 is the proof of SSAT observability and the resulting navigation maps. The proof derives the observability matrix for a distributed system using relative measurements and shows that the absolute state observability requires differential motion between vehicles in the system. This result is a generalization of previously published works which have shown that relative measurements produce fully observable DSS's. The generalization does not assume any particular dynamical system and therefore is completely system agnostic, which leads to insights into the origins of observability for SSAT. Then simulations are executed to build navigation maps which predict navigation accuracy for a DSS using SSAT in two and three body

environments.

Chapter 4 has begins with a geometric analysis on the impact of maneuvers for optical only information gathering to enhance range information. The analysis produces a relationship between measurement uncertainty and range uncertainty given a maneuver. The resulting equation is analogous to previous works but functionally unique in that ensures a positive uncertainty and has the capacity to incorporate maneuver uncertainty if desired. The derived relationship is the basis of two novel guidance equations: a heuristic policy and an analytic solution to a surrogate cost function. The guidance policies either capture or closely approximate the true optimal solution which minimizes the projection of state uncertainty orthogonal to the measurement space. Finally, the analytic policy is combined with station keeping objectives to identify a maneuver which returns a spacecraft to a desired quasi-period orbit while minimizing range uncertainty.

Chapter 5 has two major contributions. First, the ballistic OCBE is modified with nonlinear updates to produce accurate state and control estimates in cislunar space given optical-only measurements in a SSAT setting. Cislunar optical-only SSAT requires nonlinear updates due to large associated uncertainties and strongly nonlinear dynamics. Second, the control estimates from the modified OCBE are classified with supervised learning algorithms. One classifier is a custom method specifically developed for maneuvering spacecraft, and the second classifier is a classical algorithm from the machine learning domain. The classifiers produce similar levels of results, with the custom method being faster and easier to implement.

1.3 Publications

1.3.1 Journal Papers

- J.A. Greaves, and D.J. Scheeres. "Observation and Maneuver Detection for Cislunar Vehicles: Using Optical Measurements and the Optimal Control Based Estimator", *The Journal of the Astronautical Sciences*. 68(4). pp.826-854. 2021
- J.A. Greaves, and D.J. Scheeres. "Autonomous Optical-Only Spacecraft-to-Spacecraft Ab-

solute Tracking and Maneuver Classification in Cislunar Space,” (Accepted)

- J.A. Greaves, and D.J. Scheeres. ”Spacecraft to Spacecraft Absolute Tracking for Autonomous Navigation of a Distributed Space System with Relative Optical Measurements,” (In Review)
- J.A. Greaves, and D.J. Scheeres. ”Analytic Information Gathering for a Distributed Space System to Enhance State Estimation,” (In Preparation)

1.3.2 Conference Papers

- J.A. Greaves, and D.J. Scheeres. ”Estimation of Stochastic Events for Vehicles in NRHOs.” In AIAA Scitech 2020 Forum. p. 0227. 2020.
- J.A. Greaves, and D.J. Scheeres. ”Maneuver Detection for Cislunar Vehicles Using Optical Measurements.” Advanced Maui Optical and Space Surveillance Technologies Conference. 2020.
- J.A. Greaves, and D.J. Scheeres. ”Relative estimation in the cis-lunar regime using optical sensors.” Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference. 2021.
- Greaves, Jesse A., and Daniel J. Scheeres. ”Optical Only Co-Estimation and Maneuver Classification in the Cislunar Regime.” Rocky Mountain AAS GN&C Conference. 2022.
- Greaves, Jesse A., and Daniel J. Scheeres. ”Autonomous Guidance and Navigation with Optical Measurements for Formation Flying in Cislunar Space.” International Workshop on Satellite Constellations & Formation Flying. 2022.
- Greaves, Jesse A., and Daniel J. Scheeres. ”Absolute and Autonomous Navigation for Distributed Space Systems from Relative Measurements.” 33rd AAS/AIAA Space Flight Mechanics Meeting. 2023.

- Greaves, Jesse A., and Daniel J. Scheeres. "Autonomous Information Gathering Guidance for Distributed Space Systems with Optical Sensors." 2023 American Controls Conference. 2023.
- Greaves, Jesse A., and Daniel J. Scheeres. "Combined Station Keeping and Information Gathering for a Distributed Space System." Advanced Maui Optical and Space Surveillance Technologies Conference. 2023.

Chapter 2

Preliminary Material

Nomenclature

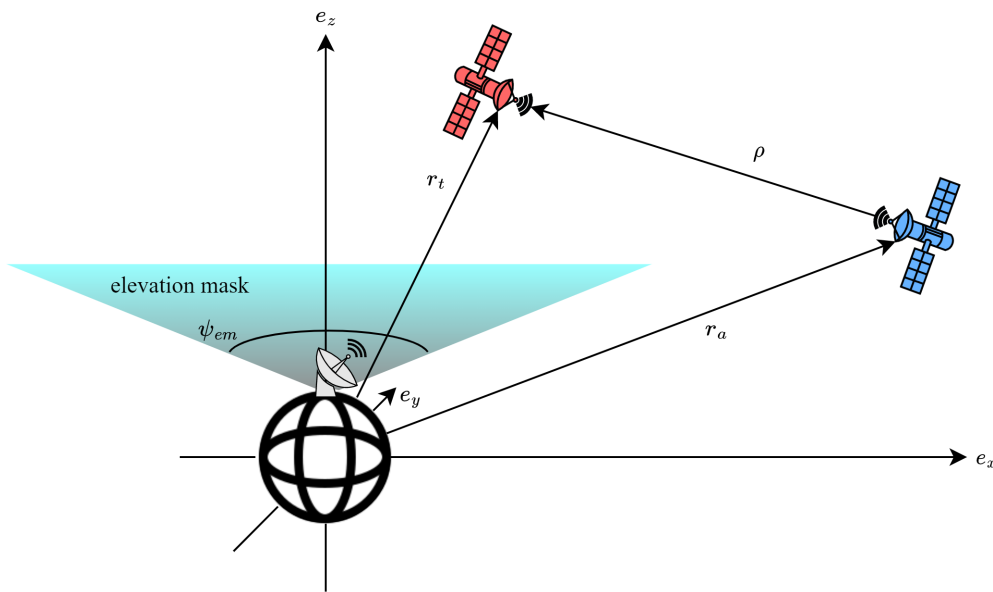
| | |
|---------------------------|--|
| x | Scalar quantity x (L2 norm if x is a vector) |
| \mathbf{x} | Vector quantity \mathbf{x} |
| \mathbf{X} | Matrix quantity \mathbf{x} |
| $\hat{\mathbf{x}}$ | Unit vector of \mathbf{x} |
| $\dot{\mathbf{x}}$ | Derivative of \mathbf{x} w.r.t. the dynamic frame denoted $\frac{D}{dt}$ |
| $\dot{\hat{\mathbf{x}}}$ | Derivative of \mathbf{x} w.r.t. the LVLH frame denoted $\frac{L}{dt}$ |
| $[\tilde{\mathbf{x}}]$ | Matrix cross product of \mathbf{x} |
| \mathbf{r} | Position vector |
| \mathbf{v} | Velocity vector |
| $\hat{\mathbf{e}}_i$ | Basis vector i |
| x_i | Dot product of \mathbf{x} and $\hat{\mathbf{e}}_i$ |
| G | Gravitational constant |
| \mathbf{I}_n | Identity matrix of size n by n |
| $\mathbf{0}_{m \times n}$ | Zero matrix of size m by n |

This chapter covers fundamental space system dynamics and measurement that reappear throughout the remainder of the thesis. The dynamical models describe spacecraft motion and are split into the two and three body problem. The measurement models represent common sensor, such as radar and optics, which are used to estimate the state of spacecraft as they navigate their environment. Since the materiel in this section is used in each subsequent chapter, models presented here and the notation established in the above nomenclature section remain constant throughout the following chapters. Each individual chapter introduces the necessary notation in the text as equations arise, but because of the considerable disparity in topics the notation between chapters may overlap.

2.1 Spacecraft Dynamics

Spacecraft motion is a Newtonian problem, it relies on gravity and other forces to integrate the position of the spacecraft as it moves through space. The models used in this work are not overly detailed or complex, they include the primary driving forces such as the two and three body problem along with major perturbations such as non-uniform gravity fields and solar radiation pressure. The position of the spacecraft is integrated as a second order ordinary differential equation so the spacecraft state \mathbf{s} includes the position vector \mathbf{r} and velocity vector \mathbf{v} such that $\mathbf{s} = [\mathbf{r}^T \ \mathbf{v}^T]^T$. The differential equation is then $\dot{\mathbf{s}} = [\mathbf{v}^T \ \mathbf{a}^T]^T$ where \mathbf{a} is the acceleration that the spacecraft experiences due to the forces acting on it.

Figure 2.1: Diagram of the two body problem.



2.1.1 Two Body Problem and Perturbations

The most fundamental representation of spacecraft motion is the two body problem (2BP). As the name implies, the problem only considers two bodies and how their mutual attraction drives motion. This problem was originally solved by Isaac Newton and has become a basis for orbital

dynamics. The 2BP is prized because it is the only orbital problem that is fully solvable due to the existence of six integrals of motion. In the presence of perturbations these integrals experience oscillation, secular drift, or both. The acceleration of a spacecraft in the two body problem can be written as

$$\mathbf{a}_{2BP} = -\frac{\mu}{r^3}\mathbf{r} + \mathbf{a}_{pert} \quad (2.1)$$

$$\mu = G(m_1 + m_2) \quad (2.2)$$

where the first term is the acceleration due to the primary body, and the second term contains perturbing accelerations. The gravitational constant of the two body problem μ contains the mass of the primary m_1 and secondary m_2 , but it is often approximated by only the primary as $\mu \approx Gm_1$ because spacecraft mass is negligible. Considered perturbations generally include non-point mass gravity fields, solar radiation pressure, atmospheric drag, and third body gravity.

Planetary bodies are not perfect point masses and have non-spherical gravity potentials which are often approximated with spherical harmonics. Spherical harmonics are akin to a Fourier series approximation of the gravity field of a planetary body which uses Legendre polynomials as a basis. The general potential for spherical harmonics is

$$U_{SH} = \frac{\mu}{r} \left(-\sum_{l=2}^{\infty} J_l \left(\frac{R_{\oplus}}{r} \right)^l P_l[\sin(\phi)] + \sum_{l=2}^{\infty} \sum_{m=1}^l \left(\frac{R_{\oplus}}{r} \right)^l P_{l,m}[\sin(\phi)] (C_{l,m}\cos(m\lambda) + S_{l,m}\sin(m\lambda)) \right) \quad (2.3)$$

where P_l are Legendre polynomials and R_{\oplus} is the planetary radius. The most important harmonic for Earth's gravity field is J2, which is nearly 1000 times larger than the next most important harmonic, and represents Earth's oblateness. The acceleration due to the J2 harmonic is

$$\mathbf{a}_{J_2} = \frac{-3\mu R_{\oplus}^2 J_2}{2r^5} \begin{bmatrix} (1 - 5 \sin^2(r_z/r))r_x \\ (1 - 5 \sin^2(r_z/r))r_y \\ (3 - 5 \sin^2(r_z/r))r_z \end{bmatrix} \quad (2.4)$$

J2 causes the orbit plane to precess about the pole via secular drift in the node, argument of periapsis, and mean anomaly. Orbit semi-major axis, eccentricity, and inclination oscillate but

their averages do not drift due to the J2 potential.

Solar radiation pressure (SRP) is the force exerted on a spacecraft by incoming photons from a nearby star. True computation of the force requires detailed shape modeling such that the lit area and its reflective index are considered. Instead of modeling the exact shape, a cannonball model is often assumed as a sufficient approximation. The force from SRP calculated via the cannonball model is

$$\mathbf{a}_{srp} = \frac{L_s}{4c\pi\rho_s^2} \frac{SA}{m} \hat{\boldsymbol{\rho}}_s \quad (2.5)$$

$$\boldsymbol{\rho}_s = \mathbf{r} - \mathbf{r}_{sun} \quad (2.6)$$

where L_s and c are the environmental parameters: luminosity of the sun and speed of light. The spacecraft parameters S , A , and m are the reflectivity, projected area, and mass of the vehicle which are assumed to be constants in the cannonball model. Because of the relatively slow-moving nature of the sun with respect to a spacecraft orbiting a planetary body, SPR generates a near constant force on spacecraft as long as it is not in shadow. Shadowing effects are discussed later in the sensor limitations section.

The two body problem neglects the mass of all other objects in the universe. To truly solve planetary motion, and spacecraft motion, requires solving the N-body problem. Solving for the effects of gravity on all bodies is intractable but including a third body, such as the sun, is often feasible and has a tangible impact on motion. Including a third body is most easily captured by taking its position from ephemerides and calculating the gravity based on the given position. Because the third body position is treated as known it acts as a perturbation to the 2BP that is being integrated. Acceleration due to a third body is


$$\mathbf{a}_{b3} = -\mu \left(\frac{\boldsymbol{\rho}_{b3}}{\rho_{b3}^3} + \frac{\mathbf{r}_{b3}}{r_{b3}^3} \right) \quad (2.7)$$

$$\boldsymbol{\rho}_{b3} = \mathbf{r} - \mathbf{r}_{b3} \quad (2.8)$$

where \mathbf{r}_{b3} is the position of the third body. Of course, there are times when the third body can be nearly as important as the primary when solving for orbital motion. In cases where the third body

is such a contributing factor then a more complex model, such as the three body problem, should be adopted.

Figure 2.2: Diagram of the circular restricted three body problem.



2.1.2 Circular Restricted Three Body Problem

The simplest form of the three body problem is the circular restricted three body problem (CR3BP). The restricted assumption is that the third body's mass is negligible compared to the two primary bodies and does not impact their motion. When the restricted assumption is true then the two massive bodies follow an analytic two body solution, and the third body has its motion dictated by two massive bodies. The circular assumption is that the two massive bodies are in co-circular orbits about their barycenter. The CR3BP is often viewed in a rotating frame, referred to as the Jacobi frame, which moves with the massive bodies such that they are fixed to the first axis. Adopting the rotating frame enables easier identification of solutions that are fixed with respect to both massive bodies. The equations of motion for a spacecraft in the CR3BP as seen in the Jacobi

rotating frame are

$$\mathbf{a}_{cr3bp} = -2[\tilde{\boldsymbol{\omega}}_{D/N}]\mathbf{v} - [\tilde{\boldsymbol{\omega}}_{D/N}][\tilde{\boldsymbol{\omega}}_{D/N}]\mathbf{r} - \frac{Gm_1}{\rho_{b1}^3}\boldsymbol{\rho}_{b1} - \frac{Gm_2}{\rho_{b2}^3}\boldsymbol{\rho}_{b2} \quad (2.9)$$

$$\boldsymbol{\omega}_{D/N} = \sqrt{\frac{G(m_1 + m_2)}{r_{12}^3}}\hat{\mathbf{e}}_z \quad (2.10)$$

In this equation r_{12} is the distance between the massive bodies. The vector $\boldsymbol{\omega}_{D/N}$ is the angular rotation rate of the CR3BP rotating frame, which is denoted as the dynamical frame D, with respect to the inertial frame N. The parameters m_1, m_2 are masses of the primary bodies, and $\boldsymbol{\rho}_{b1}$ and $\boldsymbol{\rho}_{b2}$ are the vectors from the bodies to the spacecraft \mathbf{r} .

The CR3BP is commonly non-dimensionalized because it reduces free parameters and alleviates numerical issues. The system is non-dimensionalized by normalizing position by the distance between the massive bodies and time by the rotation rate of the circular two body orbit. The conversions from dimensional to non-dimensional position, velocity, and time are

$$\mathbf{r}' = \frac{\mathbf{r}}{r_{12}} \quad (2.11)$$

$$\mathbf{v}' = \frac{\mathbf{v}}{r_{12}\boldsymbol{\omega}_{D/N}} \quad (2.12)$$

$$t' = nt \quad (2.13)$$

where the primes denote non-dimensional states. The non-dimensional equations of motion are

$$\mathbf{a}'_{cr3bp} = -2[\tilde{\boldsymbol{\zeta}}]\mathbf{v}' - [\tilde{\boldsymbol{\zeta}}][\tilde{\boldsymbol{\zeta}}]\mathbf{r}' - \frac{1-\mu}{\rho_{b1}'^3}\boldsymbol{\rho}'_{b1} - \frac{\mu}{\rho_{b2}'^3}\boldsymbol{\rho}'_{b2} \quad (2.14)$$

$$\mu = \frac{m_2}{m_1 + m_2} \quad (2.15)$$

In this equation μ is the non-dimensional gravitational parameter, and $[\tilde{\boldsymbol{\zeta}}]$ is the matrix cross product representation of the z-unit vector.

Unfortunately, unlike the two body problem, the motion of the third body cannot be generally analytically solved. Only special solutions such as equilibrium points, periodic orbits, and quasi-periodic orbits (QPOs) can be analytically obtained. Fortunately, the CR3BP is well studied and methods to obtain special solutions such as QPOs are abundant. Important algorithms utilized

within this text are shooting algorithms to calculate single trajectories, continuation to calculate orbit families, and the GMOS algorithm to compute QPO families [21]. Several periodic orbit families are examined throughout the following chapters to explore how different orbits impact navigation solutions. The included families are: northern L1 halos, planar L2 Lyapunovs, southern L2 halos, and distant retrograde orbits. The orbit families are depicted in Figure 2.3.

The southern L2 halo family warrants additional attention due to the expected deployment of NASA’s lunar gateway to the 9:2 synodic near rectilinear halo orbit (NRHO) in that family. Because of the focus on the 9:2 NRHO, QPOs about this orbit are generated and used to ensure naturally bounded relative motion for proximity operations. When generating QPOs it is useful to ensure that the base frequency, or stroboscopic period, matches the base periodic orbit so that the relative QPO has does not drift. The resulting relative trajectory produces diverse viewing geometries which provides navigational benefits. Figure 2.4 presents the 9:2 NRHO and an associated QPO about it in the CR3BP rotating frame, and the natural relative trajectory of the QPO centered about the PO over 15 orbit periods.

2.1.3 Relative Motion

Designing and flying DSS’s requires an understanding of relative motion between vehicles in the system. Relative motion is derived by defining a relative position as the difference between vehicles and then taking the appropriate derivatives. Frames, transport theorem, and relative motion are all extensively covered in Reference [22], and this book is basis for the following section. Relative motion is defined from a primary vehicle, which we call the agent, that is responsible for generating measurements and planning maneuvers. Other vehicles in the DSS are target vehicles which operate independently to achieve their own goals. By defining the system in such a manner,

Figure 2.3: Orbit families of southern L2 halos, northern L1 halos, planar L2 Lyapunov, and distant retrograde orbits from top left to bottom right respectively.

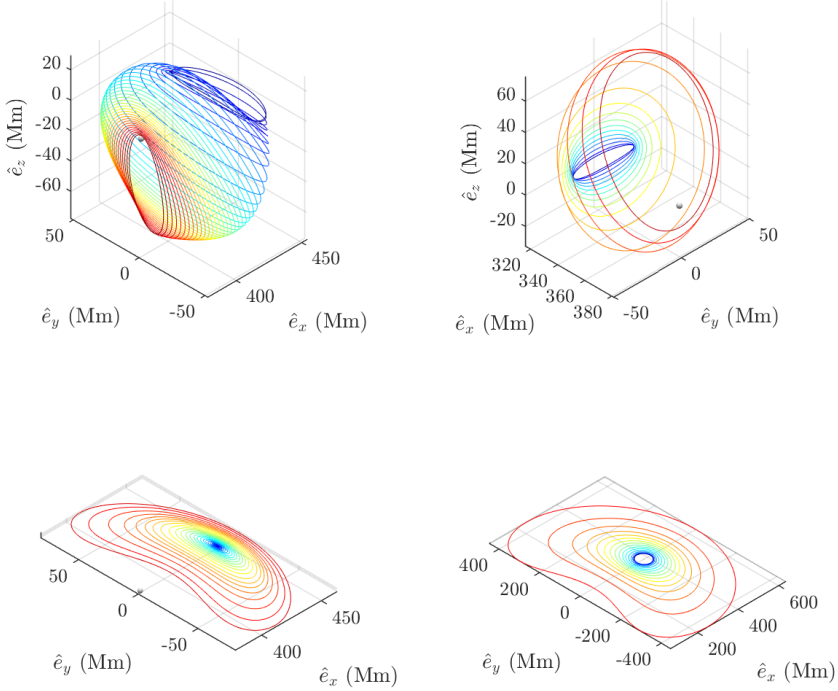
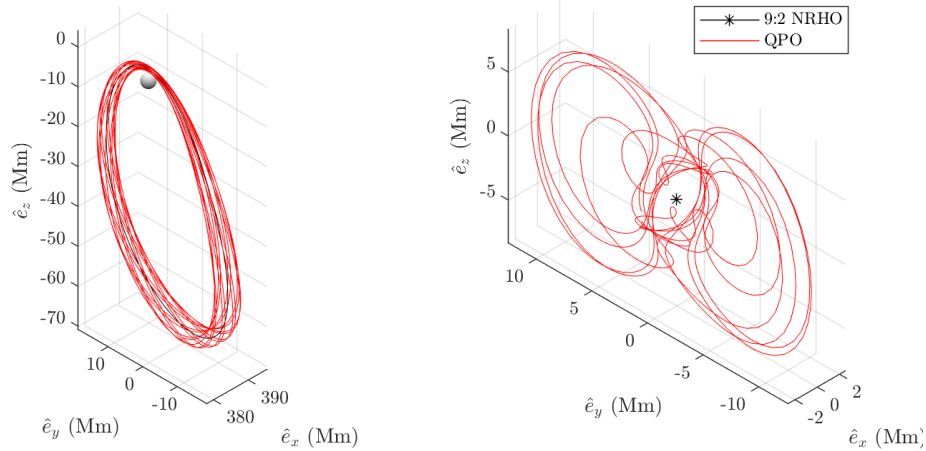


Figure 2.4: Relative motion of the 9:2 synodic QPO about the 9:2 synodic NRHO.



the relative states in the dynamic frame are

$$\boldsymbol{\rho} = \mathbf{r}_t - \mathbf{r}_a \quad (2.16)$$

$$\begin{aligned} \boldsymbol{\nu} &= \dot{\mathbf{r}}_t - \dot{\mathbf{r}}_a \\ &= \mathbf{v}_t - \mathbf{v}_a \end{aligned} \quad (2.17)$$

These states are implicitly defined with respect to the dynamic frame which the equations of motion are integrated in, which is either the inertial two body frame or rotating three body frame. The chosen dynamical model therefore defines the base frame which is referred to as the dynamic reference frame.

In addition to a relative state, a relative frame is also usually beneficial to understand rendezvous and proximity operations. The most common relative motion frame is a local-vertical local-horizontal (LVLH) frame that clearly represents in-track, out-of-track, and out-of-plane motion. The LVLH frame is defined with the first axis from the agent in the radial direction of nearby planetary body, the third axis is aligned with the angular momentum of agent, and the second axis completes the set. This LVLH basis set is written $\{\hat{\mathbf{e}}_r, \hat{\mathbf{e}}_\theta, \hat{\mathbf{e}}_h\}$. The vectors that form the LVLH basis are easily found in the dynamic reference frame and can be used to construct a rotation matrix between them via

$$[\mathbf{LD}] = \begin{bmatrix} {}^D\hat{\mathbf{e}}_r & {}^D\hat{\mathbf{e}}_\theta & {}^D\hat{\mathbf{e}}_h \end{bmatrix}^T \quad (2.18)$$

$$\hat{\mathbf{e}}_r = \hat{\mathbf{r}}_a \quad (2.19)$$

$$\hat{\mathbf{e}}_\theta = \hat{\mathbf{e}}_h \times \hat{\mathbf{e}}_r \quad (2.20)$$

$$\hat{\mathbf{e}}_h = \hat{\mathbf{h}} \quad (2.21)$$

$$\mathbf{h} = \mathbf{r}_a \times \mathbf{v}_a \quad (2.22)$$

The matrix $[\mathbf{LD}]$ is the rotation matrix from the dynamic reference frame (D) to the LVLH frame (L). The definition that is given here uses vectors written in the dynamic frame which are usually known via integration and are denoted by the pre-superscript D. With the help of the rotation

matrix and transport theorem the relative state as seen in the LVLH frame are

$${}^L\rho = [LD] {}^D\rho \quad (2.23)$$

$${}^L\dot{\rho} = [LD]({}^D\nu - {}^D\omega_{L/D} \times {}^D\rho) \quad (2.24)$$

The term ${}^L\rho$ is the relative position vector in LVLH components, and ${}^L\dot{\rho}_i$ is the derivative of the relative position with respect to the LVLH frame also in LVLH components. The only missing variable now is the rotation rate between the LVLH frame and the dynamic frame denoted by $\omega_{L/D}$. In the unperturbed two body problem this is simply the angular rate of the agent spacecraft which is

$$\omega_{L/D} = \dot{f}\hat{e}_h \quad (2.25)$$

$$\dot{f} = \frac{h}{r_a^2} \quad (2.26)$$

In the perturbed or three body problem, where orbit elements are not applicable, the process is much messier but are calculated via

$$[\tilde{\omega}_{L/D}] = [DL][LD] \quad (2.27)$$

$$[DL] = [{}^D\dot{e}_r \quad {}^D\dot{e}_\theta \quad {}^D\dot{e}_h] \quad (2.28)$$

$$\dot{e}_r = (\mathbf{I}_3 - \hat{r}_a\hat{r}_a^T) \frac{\mathbf{v}_a}{r_a} \quad (2.29)$$

$$\dot{e}_\theta = \dot{e}_h \times \hat{e}_r + \hat{e}_h \times \dot{e}_r \quad (2.30)$$

$$\dot{e}_h = (\mathbf{I}_3 - \hat{h}\hat{h}^T) \frac{\mathbf{r}_a \times \mathbf{a}_a}{h} \quad (2.31)$$

where \mathbf{a}_a is the acceleration of agent in the dynamic reference frame as given by the equations of motion. Then $[\tilde{\omega}_{L/D}]$ is the matrix cross product representation of ${}^D\omega_{L/D}$ which is the angular velocity of the LVLH frame as desired.

2.2 Measurement Models

2.2.1 Range Based Sensors

Range and range-rate measurements are typical forms of sensing data for spacecraft navigation because they are reliable sources of information obtained from radar sensors either on the ground or on spacecraft. Two examples of such sources are JPL's deep space network and Advanced Space's CAPS navigation platform. Range is the L2 norm of distance between two objects, and range-rate is the time derivative of range. These measurements usually come as a pair that are stacked in a vector denoted \mathbf{y}_ρ . The equations for range and range-rate are

$$\mathbf{y}_\rho = \begin{bmatrix} \rho & \dot{\rho} \end{bmatrix}^T \quad (2.32)$$

$$\rho = \sqrt{\boldsymbol{\rho}^T \boldsymbol{\rho}} \quad (2.33)$$

$$\dot{\rho} = \frac{(\boldsymbol{\nu}^T \boldsymbol{\rho})}{\rho} \quad (2.34)$$

The target velocity is left as the general relative velocity vector because this work considers both ground-based and space-based sensors. For ground-based networks the velocity is replaced with a function of the sensor position and planetary rotation rate.

The linear sensitivity of this measurement to the spacecraft state is an important quantity that appears in many tools later on. The first-order sensitivity of the measurement is the Jacobian of the measurement function taken with respect the state vector \mathbf{s} defined as $\mathbf{H} = \frac{\partial \mathbf{y}}{\partial \mathbf{s}}$. Because this work focuses on relative measurements between two objects in a system, it is easiest to write the sensitivity as a function of a relative state \mathbf{s}_ρ which is comprised of a relative position $\boldsymbol{\rho}$ and relative velocity $\boldsymbol{\nu}$ such that $\mathbf{s}_\rho = [\boldsymbol{\rho}^T, \boldsymbol{\nu}^T]^T$. The relative sensitivity defined $\mathbf{H}_\rho = \frac{\partial \mathbf{y}}{\partial \mathbf{s}_\rho}$ is often the main component of any Jacobian including for inertial states found via the chain rule. This fact is

easily demonstrated by taking the Jacobian with respect to inertial states as follows:

$$\begin{aligned}
 \mathbf{H} &= \frac{\partial \mathbf{y}}{\partial \mathbf{s}} \\
 &= \begin{bmatrix} \frac{\partial \mathbf{y}}{\partial \boldsymbol{\rho}} \frac{\partial \boldsymbol{\rho}}{\partial \mathbf{r}} & \frac{\partial \mathbf{y}}{\partial \boldsymbol{\nu}} \frac{\partial \boldsymbol{\nu}}{\partial \mathbf{v}} \end{bmatrix} \\
 &= \begin{bmatrix} -\frac{\partial \mathbf{y}}{\partial \boldsymbol{\rho}} & -\frac{\partial \mathbf{y}}{\partial \boldsymbol{\nu}} \end{bmatrix} \\
 &= -\mathbf{H}_\rho
 \end{aligned}$$

The sensitivity of \mathbf{y}_ρ with respect to the relative state vector $\mathbf{s}_\rho = [\boldsymbol{\rho}^T \boldsymbol{\nu}^T]^T$, denoted as \mathbf{H}_ρ for the range measurements, is:

$$\mathbf{H}_\rho = \frac{1}{\rho} \begin{bmatrix} \boldsymbol{\rho}^T & \mathbf{0}_{1 \times 3} \\ \boldsymbol{\nu}^T - \frac{\dot{\rho}}{\rho} \boldsymbol{\rho}^T & \boldsymbol{\rho}^T \end{bmatrix} \quad (2.35)$$

2.2.2 Optical Based Sensors

Optical sensors are another common measurement source for spacecraft navigation which are readily available and low cost payloads that are deployed on many missions. Optical measurements are represented in many ways, but one common method is an azimuth and elevation angle pair as given by:

$$\mathbf{y}_\theta = \begin{bmatrix} \theta_a & \theta_e \end{bmatrix}^T \quad (2.36)$$

$$\theta_a = \tan^{-1} \left(\frac{\rho_y}{\rho_x} \right) \quad (2.37)$$

$$\theta_e = \sin^{-1} \left(\frac{\rho_z}{\rho} \right) \quad (2.38)$$

By using the above representation it is assumed that there is a separate image processing system has identified the target and converted the image data into the dynamic reference frame via necessary frame rotations and transformations. Note that because of the arc-tangent function that azimuth is not continuous and will have discontinuities in it. To ensure a continuous difference between two azimuth angles, an incoming measurement θ_a can be lifted to be continuous with a reference angle

$\theta_{a,0}$ as follows

$$\theta_a = \begin{cases} \theta_a + 2\pi & \text{if } \theta_{a,0} - \theta_a > \pi \\ \theta_a - 2\pi & \text{if } \theta_{a,0} - \theta_a < \pi \\ \theta_a & \text{otherwise} \end{cases} \quad (2.39)$$

The measurement sensitivity with respect to the relative state vector for the optical case is:

$$\mathbf{H}_\theta = \begin{bmatrix} \frac{-\rho_y}{\rho_x^2 + \rho_y^2} & \frac{\rho_x}{\rho_x^2 + \rho_y^2} & 0 & \mathbf{0}_{1 \times 3} \\ \frac{-\rho_x \rho_z}{\rho^2 \sqrt{\rho_x^2 + \rho_y^2}} & \frac{-\rho_y \rho_z}{\rho^2 \sqrt{\rho_x^2 + \rho_y^2}} & \frac{\sqrt{\rho_x^2 + \rho_y^2}}{\rho^2} & \mathbf{0}_{1 \times 3} \end{bmatrix} \quad (2.40)$$

2.2.3 Sensor Limitations

2.2.3.1 Line of Sight Obstruction

All sensors are limited by line of sight which can be obstructed by planetary bodies. Line of sight is obstructed if the relative vector between objects intersects with the planetary body. Mathematically, this boils down to solving for the intersection between a line and surface. If the planetary body is approximated as a sphere, then the surface is

$$\|\mathbf{x} - \mathbf{o}\|^2 = R_\oplus^2 \quad (2.41)$$

where \mathbf{o} is the center of planet. The equation for the line of sight has its origin at the observers position, which is assigned to the agent's position \mathbf{r}_a , and is directed towards the target along the relative range unit vector $\hat{\boldsymbol{\rho}}$ as

$$\mathbf{x} = \mathbf{r}_a + d\hat{\boldsymbol{\rho}} \quad (2.42)$$

The distance d gives the intersection point between the line and the sphere, and it is solved as a quadratic equation. The solution to the quadratic tells us if there are 0, 1, or 2 intersections depending on the quadratic equation. When there is no intersection and visibility is acquired then the following is true

$$(\hat{\boldsymbol{\rho}}^T(\mathbf{r}_a - \mathbf{o}))^2 - (\|\mathbf{r}_a - \mathbf{o}\|^2 - r_a^2) < 0 \quad (2.43)$$

2.2.3.2 Angle Masks and Keep Out Cones

Angle masks are used to impose elevation angle limits for ground based sensors and keep out regions for optical sensors. Both cases are an applications of the dot product rule to find the angle between two vectors. First, there must be a mask line $\hat{\mathbf{l}}$ which defines the center of the cone which the mask angle is defined about. Second, there is the line of sight vector which is used to determine if the viewing angle is inside or outside the angle mask. Thus the viewing angle ψ is calculated as

$$\psi = \text{acos}(\hat{\mathbf{l}}^T \hat{\boldsymbol{\rho}}) \quad (2.44)$$

To implement a ground station elevation mask, the mask line $\hat{\mathbf{l}}$ is the vector from the center of the planet to the observer position, and the viewing angle ψ must be less than the elevation mask angle ψ_{em} such that $\psi < \psi_{em}$. Conversely, the optical keep out is enforced by ensuring the viewing angle is greater than the keep out mask ψ_{ko} such that $\psi > \psi_{ko}$. For solar keep out regions the mask line is the vector from the observer to the sun.

2.2.3.3 Shadowing

Optical imaging is not possible if the target is in a shadowed region. For an object to be in a shadow, it needs to be the umbra of a planet. For this work the umbra is approximated as a cylinder behind the planet that is directed along the solar incident line.

The shadowing process is similar to the angle mask constraint in that it relies on the angle between vectors, so the same notation is used. First, the solar incident line is defined as the vector from the sun to the center of the planet and is assigned $\hat{\mathbf{l}}$. Then the vector from the center of the planet to the object is becomes the relative vector $\hat{\boldsymbol{\rho}}$. The angle between these vectors displacement angle ψ and is calculated equivalently to Eq. 2.44.

The first condition for the object to be in shadow is that is must be behind the planet, which is true if $\psi < \pi/2$. Then the object must be within the cylinder which is found by using the dot product rule again. The object is in the cone if $\rho \sin(\psi) < R_{\oplus}$. If both conditions are met then the object is in shadow and cannot be measured by an optical sensor.

Chapter 3

Spacecraft to Spacecraft Absolute Tracking

3.1 Introduction and Background on Spacecraft to Spacecraft Tracking

An autonomous DSS requires guidance, navigation, and control methods that can be computed in real time and only depend on onboard resources. With regards to navigation, the state estimation algorithm must be observable from the data acquired by sensors within the DSS and computationally tractable for flight computer. Meeting the observability requirement for a general distributed system entails solving the state of all vehicles in the system, which is a challenge due to the onboard-only sensor restriction. The primary issue stems from the fact that the global, or absolute, state is not directly sensed without external ground-based sensors. Fortunately, an absolute state can be obtained from onboard relative sensors under specific scenarios if the DSS is appropriately designed to leverage nonlinearities in spacecraft motion. Moving forwards, we hereby refer to the use of relative measurements to generate an absolute state estimate for all vehicles in a DSS as spacecraft to spacecraft absolute tracking (SSAT).

The use of relative measurement to obtain an absolute state estimate for all objects in a distributed system falls under the purview of several fields of research which include but is not limited to robotics, controls, and aerospace. Within the robotics and controls communities this most closely translates to distributed/networked localization; a topic where multiple agents localize themselves with respect to each other [23, 24]. Many of the related works within the robotics community assume that agents have the capacity to pass information between each other, which unfortunately is not always applicable for space systems. Requiring space-based agents to pass

information not only reduces the level of autonomy by requiring cooperation, but it may not be possible due to the inclusion of non-cooperative targets such as space debris or derelict space objects.

While cooperation may not be guaranteed for space operations, orbital dynamics do guarantee predictable nonlinear motion which can be leveraged to treat target objects as dynamic features in the environment. Estimating features is more akin to a general multi-agent localization or simultaneous localization and mapping problems [25]. There are various papers within the robotics community that consider these problems and have shown that estimating known landmarks can lead to global localization [26, 27]. Some have even considered angle-only observability for multi-agent localization showing nonlinearity is critical for full observability [28]. Extending this framework to space operations means that the “landmarks” are now dynamic targets and gives rise to the SSAT problem. The SSAT has been studied in the past but has recently received renewed interest due to the potential to reduce reliance on costly and oversubscribed ground systems.

To the best of our knowledge, the first proof that showed that absolute navigation can be obtained from relative measurements within a DSS was by F.L. Markley in 1984 [4]. Later, M.L. Psiaki refined the analysis by likening the spacecraft formation to a large-scale accelerometer, and further demonstrated that the addition of non-spherical gravity fields broadens observable conditions [8]. Then K.A. Hill examined SSAT in cislunar space with relative-range measurements, and demonstrated by extension, the utility of SSAT for interplanetary missions [29]. Contemporary works expanded on these results by exploring specific scenarios with various measurement sensors and dynamical environments [30, 6, 7, 31, 32]. Further study into relative measurements for absolute estimation of dynamical systems is warranted to identify conditions in which this approach excels or falters.

Additionally, there is a particular interest in studying SSAT with optical-only sensors. Optical-only SSAT is appealing because optics are readily available and low cost payloads with an abundance of literature and demonstration within spaceflight [33, 34, 35, 14]. Recent studies have also demonstrated that optical sensors can overcome range/scale ambiguities if the system dynamics

are sufficiently nonlinear, and that optical-only measurements can successfully reconstruct a fully observable state solution [36, 37, 7, 38, 32, 39]. The combined advantages of optical sensors present the potential for entirely autonomous navigation of a DSS, and therefore it is prudent to assess the capabilities optical-only SSAT for a DSS. Still all of these previous works consider specific dynamical systems, generally Keplerian dynamics, and do not generalize beyond selected case studies. Additionally, only a select few consider estimating the absolute state of a DSS with optical-only measurements, leaving further room for examination.

The remainder of the chapter begins by defining the problem formulation and introduces the SSAT state space along with basic observability theory. Next, an analytic derivation of linearized observability for SSAT is derived. Then various SSAT scenarios with optical-only measurements are simulated under a linear covariance analysis to construct potential navigation solutions and maps. Finally, the conclusions are presented, and some future work is speculated upon.

3.2 Problem Formulation

3.2.1 Spacecraft to Spacecraft Absolute Tracking Model

The SSAT problem estimates the state of every object in a DSS from relative measurements between objects in the system. In the simplest form a DSS is only two vehicles; an agent which is actively managed and a target which is nominally non-cooperative. Let $\mathbf{s}_a \in \mathbb{R}^n$ be the state of the agent in a dynamic reference frame of choice where agent's motion is only dependent on its own state. The agent is responsible for generating measurements and is actively managed. Then let $\mathbf{s}_t \in \mathbb{R}^n$ be the state of the target in the same dynamic frame. The target operates independently of the agent so that cooperative planning is not required in this architecture. Because the SSAT problem seeks to estimate the state of both vehicles, the SSAT system state is comprised of the target and the agent sub-states together:

$$\mathbf{x} = \begin{bmatrix} \mathbf{s}_a^T & \mathbf{s}_t^T \end{bmatrix}^T \quad (3.1)$$

Estimating the state of both vehicles in a dynamic reference frame where vehicle motion

is uncoupled is equivalent to estimating a absolute and relate state simultaneously. In fact, an absolute and relative state can be defined as a transformation of the SSAT system state. Let the absolute state be defined as the mean of the vehicles and the relative state be their difference as given by:

$$\mathbf{s}_\alpha = \frac{\mathbf{s}_a + \mathbf{s}_t}{2} \quad (3.2)$$

$$\mathbf{s}_\rho = \mathbf{s}_t - \mathbf{s}_a \quad (3.3)$$

Then an absolute-relative system state is defined by stacking the sub-states as follows:

$$\tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{s}_\alpha^T & \mathbf{s}_\rho^T \end{bmatrix}^T \quad (3.4)$$

From the definitions above the linear transformation \mathbf{T} from the SSAT system state \mathbf{x} to the absolute-relative system state $\tilde{\mathbf{x}}$ is:

$$\tilde{\mathbf{x}} = \mathbf{T}\mathbf{x} \quad (3.5)$$

$$\mathbf{T} = \begin{bmatrix} \frac{1}{2}\mathbf{I}_{n \times n} & \frac{1}{2}\mathbf{I}_{n \times n} \\ -\mathbf{I}_{n \times n} & \mathbf{I}_{n \times n} \end{bmatrix} \quad (3.6)$$

A convenient property of the \mathbf{T} matrix is that it is orthogonal. An orthonormal transformation can also be defined by removing the half factors in the top row, and scaling the whole matrix by $1/\sqrt{2}$. The orthonormal transformation has additional useful properties but provides less intuitive insight, so it is neglected in favor of the non-normal transformation.

Defining the absolute-relative system as the average and difference of the vehicles allows us to leverage the linear transformation \mathbf{T} between the SSAT system and absolute-relative systems. The transformation leads to various simplifications which provide insight into observability properties later, and is a representation of Psiaki's insight into the problem that the center of the configuration acts as an accelerometer. In practice it may be easier to implement a different absolute sub-state, such as using the inertial state of the agent spacecraft, but the observability properties from the following work are valid under any linear transformation of these states. In fact, we provide an agent-based absolute state later to demonstrate the equivalence.

The natural dynamics model that dictates motion for a single vehicle in the dynamic reference frame is given by function $\dot{\mathbf{s}}_i = \mathbf{f}(\mathbf{s}_i)$, which assumes that inter-vehicle gravitational effects are negligible and that the frame is suitable for integrating motion based on the state. Thus, the SSAT system state evolves according to:

$$\dot{\mathbf{x}} = \begin{bmatrix} \mathbf{f}(\mathbf{s}_a)^T & \mathbf{f}(\mathbf{s}_t)^T \end{bmatrix}^T \quad (3.7)$$

The state is integrated given initial conditions \mathbf{x}_0 to time τ through the nonlinear solution flow function $\phi(\mathbf{x}_0, \tau)$. The state transition matrix (STM) $\Phi(\tau, \tau_0)$, which propagates linearized deviations from time τ_0 to τ , is the partial of the solution flow with respect to the initial conditions. The equations for the solution flow and STM are:

$$\mathbf{x} = \phi(\mathbf{x}_0, \tau) \quad (3.8)$$

$$\Phi(\tau, \tau_0) = \frac{\partial \mathbf{x}}{\partial \mathbf{x}_0} \quad (3.9)$$

For the SSAT system, the full STM is comprised of decoupled individual vehicle STMs such that:

$$\Phi(\tau, \tau_0) = \begin{bmatrix} \Phi_a(\tau, \tau_0) & \mathbf{0}_{n \times n} \\ \mathbf{0}_{n \times n} & \Phi_t(\tau, \tau_0) \end{bmatrix} \quad (3.10)$$

where vehicle i has the STM $\Phi_i = \frac{\partial \mathbf{s}_i}{\partial \mathbf{s}_{i_0}}$. The individual vehicle STMs are easy to calculate and are computed via the differential equation:

$$\Phi_i(\tau_0, \tau_0) = \mathbf{I}_n \quad (3.11)$$

$$\dot{\Phi}_i(\tau) = \mathbf{A}_i(\tau) \Phi_i(\tau, \tau_0) \quad (3.12)$$

where $\mathbf{A}_i = \frac{\partial \mathbf{f}}{\partial \mathbf{s}_i}$ is the partial of the dynamics function with respect to the vehicle sub-state. Due to the linear transformation between the SSAT system and absolute-relative system, the STM for the absolute-relative system denoted $\tilde{\Phi}(\tau_1, \tau_0)$ is also a linear transformation. The absolute-relative STM is:

$$\begin{aligned} \tilde{\Phi}(\tau, \tau_0) &= \mathbf{T} \Phi(\tau, \tau_0) \mathbf{T}^{-1} \\ &= \begin{bmatrix} \frac{1}{2}(\Phi_a(\tau, \tau_0) + \Phi_t(\tau, \tau_0)) & \frac{1}{4}(\Phi_t(\tau, \tau_0) - \Phi_a(\tau, \tau_0)) \\ \Phi_t(\tau, \tau_0) - \Phi_a(\tau, \tau_0) & \frac{1}{2}(\Phi_a(\tau, \tau_0) + \Phi_t(\tau, \tau_0)) \end{bmatrix} \end{aligned} \quad (3.13)$$

Measurements within the SSAT framework only consider relative quantities between the agent and target. Thus, the measurements are obtained from the measurement function $\mathbf{y} = \mathbf{h}(\mathbf{s}_\ell) \in \mathbb{R}^m$ which only relies on the relative sub-state. Alternatively the measurement function must be written as $\mathbf{y} = \mathbf{h}(\mathbf{s}_a, \mathbf{s}_t) = \mathbf{h}(\mathbf{x})$ when it is used in the SSAT system. Taking the partials with respect to the SSAT system and absolute-relative system leads to:

$$\mathbf{H} = \begin{bmatrix} \frac{\partial \mathbf{h}}{\partial \mathbf{s}_\ell} & \frac{\partial \mathbf{h}}{\partial \mathbf{s}_a} \\ \frac{\partial \mathbf{h}}{\partial \mathbf{s}_\ell} & \frac{\partial \mathbf{h}}{\partial \mathbf{s}_t} \end{bmatrix} = \begin{bmatrix} -\mathbf{H}_\ell & \mathbf{H}_\ell \end{bmatrix} \quad (3.14)$$

$$\tilde{\mathbf{H}} = \begin{bmatrix} \frac{\partial \mathbf{h}}{\partial \mathbf{s}_a} & \frac{\partial \mathbf{h}}{\partial \mathbf{s}_\ell} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_{m \times n} & \mathbf{H}_\ell \end{bmatrix} \quad (3.15)$$

The nature of relative measurements creates symmetry in the SSAT system as seen in Equation (3.14). The measurement sensitivity matrix taken with respect to the absolute-relative system on the other hand has a zero component which leads to convenient simplifications later. Regardless, it is clear that relative measurements do not depend on the absolute state, and the lack of sensitivity raises concerns about the absolute state being observable.

3.2.2 Linearized Observability Theory

The following observability analysis is based on a linearization approach which estimates a linear deviation from a nominal trajectory. Taking a linearization approach assumes that there is an a-priori nominal and that the state's first-order variation is being estimated. This approach is adopted for three primary reasons. First, prior works by Markley, Psiaki, and Born set a precedence for these techniques for spacecraft orbit determination because spacecraft behavior is generally well known and within a linear regime of a set nominal [4, 8, 40]. Second, linear variational observability generally coincides with local first-order nonlinear observability [37, 28]. Finally, this approach leads to insights into observability for the SSAT system. Thus, given a linear system, the state is said to be observable at time τ_k if the observability matrix \mathcal{O} is full rank. The observability matrix is:

$$\mathcal{O} = \begin{bmatrix} [\mathbf{H}\Phi]_{\tau_0} \\ \vdots \\ [\mathbf{H}\Phi]_{\tau_k} \end{bmatrix} \quad (3.16)$$

The notation $[\cdot]_{\tau_k}$ means that the expression in the bracket is evaluated at time τ_k such that $\Phi(\tau_k, \tau_0)$ and \mathbf{H} is calculated along the nominal at the corresponding time.

There are two notes to make about testing for observability with the observability matrix. First, if the system achieves observability at some time τ_k then it remains observable for all time afterwards. Second, observability is a binary condition that identifies if an estimate for every state can be made, but does not quantify how accurate the estimate is. One method to ascertain the quality of the estimate is to construct the information matrix $\mathbf{\Lambda}$:

$$\mathbf{\Lambda} = \mathcal{O}^T \mathbf{W} \mathcal{O} \quad (3.17)$$

The matrix \mathbf{W} is the measurement weight matrix for all observations, which is a diagonal matrix if the measurements are uncorrelated.

The information matrix $\mathbf{\Lambda}$ is the inverse of the state covariance matrix \mathbf{P} which describes how accurately each state is known. A drawback of the information matrix is that it can be poorly conditioned due to the fact that it is effectively a square of the observability matrix. The poor conditioning is often due to scaling discrepancies between sub-states, either agent-target or absolute-relative. This scaling difference is easily confused with a unobservable system because unobservable systems are by definition singular and poorly conditioned. Fortunately, there are methods which can avoid the scaling issues by taking a block inversion that only requires a portion of the matrix and its Schur complement to be well conditioned [41]. The formula for these inversions are:

$$\begin{aligned} \mathbf{P} = \mathbf{\Lambda}^{-1} &= \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix}^{-1} \\ &= \begin{bmatrix} \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{C}\mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{B}(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1} \\ -(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{C}\mathbf{A}^{-1} & (\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1} \end{bmatrix} \\ &= \begin{bmatrix} (\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} & -(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1}\mathbf{B}\mathbf{D}^{-1} \\ -\mathbf{D}^{-1}\mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} & \mathbf{D}^{-1} + \mathbf{D}^{-1}\mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1}\mathbf{B}\mathbf{D}^{-1} \end{bmatrix} \end{aligned} \quad (3.18)$$

In this equation \mathbf{A} and \mathbf{D} must be square blocks, and when applied to the SSAT systems represent

the agent-target or absolute-relative states respectively. Breaking up the matrix into related sub-states in this manner often alleviates numerical precision problems and allows cases to be properly reported.

3.3 Observability of Spacecraft to Spacecraft Absolute Tracking

3.3.1 Derivation of the Generic Observability Matrix

We seek to identify the conditions where SSAT observability can be obtained solely from relative measurements in a general dynamical system. We begin by constructing the observability matrix for the absolute-relative system in order to clearly delineate the origins of absolute vs relative information. Doing this results in the observability matrix:

$$\tilde{\mathcal{O}} = \begin{bmatrix} \left[\tilde{\mathbf{H}} \tilde{\Phi} \right]_{\tau_0} \\ \vdots \\ \left[\tilde{\mathbf{H}} \tilde{\Phi} \right]_{\tau_k} \end{bmatrix} \quad (3.19)$$

Equation (3.19) is now expanded as a function of uncorrelated SSAT dynamics from Eq. (3.15), and the measurement sensitivity matrix from Equation (3.13). Doing this results in:

$$\tilde{\mathcal{O}}_{SSAT} = \begin{bmatrix} \left[\mathbf{H}_\rho(\Phi_t - \Phi_a) \right]_{\tau_0} & \frac{1}{2} \left[\mathbf{H}_\rho(\Phi_t + \Phi_a) \right]_{\tau_0} \\ \vdots & \vdots \\ \left[\mathbf{H}_\rho(\Phi_t - \Phi_a) \right]_{\tau_k} & \frac{1}{2} \left[\mathbf{H}_\rho(\Phi_t + \Phi_a) \right]_{\tau_k} \end{bmatrix} \quad (3.20)$$

Equation (3.20) leads to the critical realization that absolute state information, which corresponds to the left half of $\tilde{\mathcal{O}}$, originates from the difference in motion between the two vehicles. This realization directly indicates that absolute state will not be observable under certain dynamical conditions. The simplest condition that leads to an unobservable absolute state is when both vehicles experience uniform motion, such as in linear time invariant dynamical systems. Another condition arises even if a system does not have uniform dynamics, which is when the target's state approaches the agent's state. A system is unobservable in the limit as the two vehicles approach each other because the dynamics, and hence the STMs, become equivalent. The latter condition is

expressed as:

$$\lim_{\mathbf{s}_t \rightarrow \mathbf{s}_a} \tilde{\Theta} = \begin{bmatrix} \mathbf{0}_{m \times n} & [\mathbf{H}_\rho \Phi_a]_{\tau_0} \\ \vdots & \vdots \\ \mathbf{0}_{m \times n} & [\mathbf{H}_\rho \Phi_t]_{\tau_k} \end{bmatrix} \quad (3.21)$$

We can further elucidate the source of absolute state information by approximating the target as a deviation from the agent. Relating the agent and target in such a manner allows for the target STM to be expressed as a function of the agent STM. The process starts by defining the target as a state deviation from agent [42]:

$$\begin{aligned} \mathbf{s}_t &= \mathbf{s}_a + \delta \mathbf{s} \\ &= \phi(\mathbf{s}_{a_0} + \delta \mathbf{s}_0, \tau) \end{aligned} \quad (3.22)$$

Then the solution flow is approximated with a Taylor series expansion:

$$\begin{aligned} \mathbf{s}_t &= \phi(\mathbf{s}_{a_0}, \tau) + \frac{\partial \phi(\mathbf{s}_{a_0}, \tau)}{\partial \mathbf{s}_{a_0}} \delta \mathbf{s}_0 + \frac{1}{2} \frac{\partial^2 \phi(\mathbf{s}_{a_0}, \tau)}{\partial \mathbf{s}_{a_0}^2} \delta \mathbf{s}_0^2 + H.O.T.(\delta \mathbf{s}_0^3) \\ &= \mathbf{s}_a + \Phi_a \delta \mathbf{s}_0 + \frac{1}{2} \Phi_a^{(2)} \delta \mathbf{s}_0^2 + H.O.T.(\delta \mathbf{s}_0^3) \end{aligned} \quad (3.23)$$

The notation $\phi(\mathbf{s}_{i_0}, \tau)$ is the solution flow from the initial condition \mathbf{s}_{i_0} to time τ . The term $\Phi_i^{(2)} \delta \mathbf{s}_0^2$ is short hand for the Einstein sum of the second-order state transition tensor (STT) $\Phi_i^{(2)}$ with the state deviation $\delta \mathbf{s}_0$ which omits state and time dependencies for simplicity. Finally, $H.O.T.(\delta \mathbf{s}_0^3)$ is used to capture all higher-order terms that are third power and higher with respect to $\delta \mathbf{s}_0$.

Now the expression for \mathbf{s}_t from Eq. (3.23) can be used to define the target STM. Using the definition of the STM and utilizing the following identities $\frac{\partial \mathbf{s}_{t_0}}{\partial \mathbf{s}_{t_0}} = \mathbf{I}_{6 \times 6}$, $\frac{\partial \mathbf{s}_{a_0}}{\partial \mathbf{s}_{t_0}} = \mathbf{0}_{6 \times 6}$, and $\Phi_a^{(2)} = \Phi_a^{(2)T}$ yields:

$$\begin{aligned} \Phi_t &= \frac{\partial \mathbf{s}_t}{\partial \mathbf{s}_{t_0}} \\ &= \frac{\partial \mathbf{s}_a}{\partial \mathbf{s}_{t_0}} + \Phi_a \frac{\partial (\mathbf{s}_{t_0} - \mathbf{s}_{a_0})}{\partial \mathbf{s}_{t_0}} + \frac{1}{2} \Phi_a^{(2)} \frac{\partial (\mathbf{s}_{t_0} - \mathbf{s}_{a_0})^2}{\partial \mathbf{s}_{t_0}} + \frac{\partial H.O.T.(\delta \mathbf{s}_0^3)}{\partial \mathbf{s}_{t_0}} \\ &= \Phi_a + \Phi_a^{(2)} \delta \mathbf{s}_0 + H.O.T.(\delta \mathbf{s}_0^2) \end{aligned} \quad (3.24)$$

Now Φ_t from the above equation can substituted into Equation (3.20). Making the substitution

and finally dropping the higher order terms $H.O.T.(\delta\mathbf{s}_0^2)$ results in the conclusion:

$$\tilde{\mathcal{O}}_{STT} = \begin{bmatrix} \left[\mathbf{H}_\rho \Phi_a^{(2)} \delta \mathbf{s}_0 \right]_{\tau_0} & \left[\mathbf{H}_\rho \left(\Phi_a + \frac{1}{2} \Phi_a^{(2)} \delta \mathbf{s}_0 \right) \right]_{\tau_0} \\ \vdots & \vdots \\ \left[\mathbf{H}_\rho \Phi_a^{(2)} \delta \mathbf{s}_0 \right]_{\tau_k} & \left[\mathbf{H}_\rho \left(\Phi_a + \frac{1}{2} \Phi_a^{(2)} \delta \mathbf{s}_0 \right) \right]_{\tau_k} \end{bmatrix} \quad (3.25)$$

3.3.2 Generic Observability Conditions

In order for the SSAT system to be observable the matrix \mathcal{O}_{SSAT} must be full rank. The matrix \mathcal{O}_{STT} approximates \mathcal{O}_{SSAT} , but it provides more insight into the problem by reducing free parameters which enables conditions for observability to be stated more easily. These conditions are best approached by partitioning \mathcal{O}_{STT} into absolute information contained on the left-hand side of the matrix, and relative information on the right-hand side of the matrix such that $\tilde{\mathcal{O}}_{STT} = [\tilde{\mathcal{O}}_\alpha, \tilde{\mathcal{O}}_\rho]$. Separating absolute and relative portions allows us to identify conditions that must be satisfied for observability to be met for each sub-state.

Beginning with relative information on the right-hand side, it is clear that the relative observability is dependent on full nonlinear vehicle motion:

$$\text{rank}(\tilde{\mathcal{O}}_\rho) = n \rightarrow \left[\mathbf{H}_\rho \left(\Phi_a + \frac{1}{2} \Phi_a^{(2)} \delta \mathbf{s}_0 \right) \right]_{\tau_i} \text{ are L.I. for } k \geq \frac{n}{m} \quad (3.26)$$

This is expected because the relative measurements capture the full relative motion. More importantly, the higher order terms include the initial vehicle separation through the term $\mathbf{H}_\rho \Phi_a^{(2)} \delta \mathbf{s}_0$. The dependency on \mathbf{s}_0 means that any change in initial deviation may produce a unique viewing profile as long as the initial deviation does not lie in a null-space of $\mathbf{H}_\rho \Phi_a^{(2)}$. This dependency only manifests if there are nonlinear dynamics and means that traditional relative observability issues still arise if the dynamics are sufficiently linear. This leads to the conclusion that relative observability is obtained when the observed motion is linearly independent for k measurements such that $k \geq \frac{n}{m}$. Equation (3.26) states this constraint mathematically.

Next examining the left-hand portion of Eq. (3.25) reveals that there are several necessary

conditions for the absolute state to be observable:

$$\text{rank}(\tilde{\mathcal{O}}_\alpha) = n \rightarrow \left. \begin{array}{l} \left[\Phi_a^{(2)} \right]_{\tau_i} \neq \mathbf{0}_{n \times n \times n} \\ \left[\mathbf{H}_\rho \Phi_a^{(2)} \right]_{\tau_i} \delta \mathbf{s}_0 \text{ are L.I.} \\ \delta \mathbf{s}_0 \notin \text{kernel} \left(\left[\mathbf{H}_\rho \Phi_a^{(2)} \right]_{\tau_i} \right) \\ \delta \mathbf{s}_0 \neq \mathbf{0}_{n \times 1} \end{array} \right\} \text{for } k \geq \frac{n}{m} \quad (3.27)$$

First, the second-order STT must be non-zero. Second, the observed second order motion must be linearly independent of itself and the first order STM. Finally, the initial deviation between the two vehicles must be non-zero and not in the null space of the second-order STT. These conditions are written in Eq. (3.27) and again must be met for $k \geq \frac{n}{m}$.

The above conditions on the dynamics are required for the absolute state portion of the observability matrix to be both full rank and linearly independent of the relative state. The non-zero second-order requirement is the most stringent condition and directly implies that the absolute state cannot be observed in a linear dynamical system. The remaining conditions are less demanding, and it is fair to assume that they will be met for most nonlinear systems where the vehicles are not docked/collided. These conditions can be summarized as meaning that the absolute state will be unobservable if both vehicles experience identical forces. This summary is a generalization of previous conclusions from Markley's analysis on DSS's [4] which are now applied to nonlinear systems as a whole.

A final analysis on the information quality of the state estimate can be inferred from Eq (3.25) and Eq. (3.17). The absolute state information quantity is related to the square of the absolute portion of the observability matrix because the information matrix is akin to a square of the observability matrix. Therefore, increasing the state deviation $\delta \mathbf{s}_0$, or keeping higher order terms which add to the absolute portion of matrix, effectively increase the scale of the absolute state information. Obtaining enhanced state information from larger deviation or strongly nonlinear motion coincides with intuition developed from Eq. (3.20) since both tend to increase differences in motion, which is a conclusion that matches previous work from Hill's Liaison navigation results

[43].

3.3.3 Multiple Targets and Agent-Based Absolute State

When multiple targets are considered it is much less intuitive to use an average state to represent the absolute state. Instead, the agent's state in the dynamic reference frame is a more suitable option. Setting the agent's state as the absolute and leaving relative states as differences from the agent to targets leads to a state space formulation that is:

$$\tilde{\mathbf{x}} = [\mathbf{s}_a^T, \mathbf{s}_{\rho_1}^T, \dots, \mathbf{x}_{\rho_k}^T]^T \quad (3.28)$$

where \mathbf{s}_{ρ_i} is the relative state of target i . Using Eq. (3.28) also means that the relative transformation and observability matrices have been altered, and are now:

$$\mathbf{T} = \begin{bmatrix} \mathbf{I}_{n \times n} & \mathbf{0}_{n \times n} & \dots & \mathbf{0}_{n \times n} \\ -\mathbf{I}_{n \times n} & \mathbf{I}_{n \times n} & & \\ \vdots & & \ddots & \\ -\mathbf{I}_{n \times n} & & & \mathbf{I}_{n \times n} \end{bmatrix} \quad (3.29)$$

$$\tilde{\mathcal{O}}_{AAO} = \begin{bmatrix} \sum_{i=1}^k [\mathbf{H}_{\rho_i}(\Phi_{t_i} - \Phi_a)]_{\tau_0} & [\mathbf{H}_{\rho_1} \Phi_{t_1}]_{\tau_0} & \dots & [\mathbf{H}_{\rho_k} \Phi_{t_k}]_{\tau_0} \\ \vdots & \vdots & & \vdots \\ \sum_{i=1}^k [\mathbf{H}_{\rho_i}(\Phi_{t_i} - \Phi_a)]_{\tau_k} & [\mathbf{H}_{\rho_1} \Phi_{t_1}]_{\tau_k} & \dots & [\mathbf{H}_{\rho_k} \Phi_{t_k}]_{\tau_k} \end{bmatrix} \quad (3.30)$$

Equation (3.30) assumes that all vehicles are observed simultaneously, which we refer to as 'all at once' (AAO) viewing. Removing the simultaneous assumption spreads the sensitivities out over time where only one target will be non-zero at a given epoch, which we refer to as 'one at a time' (OAT) viewing. Importantly though, examining Eq (3.30) reveals that the agent's information is summed over all targets, indicating that including more targets has the potential to produce more information. The targets on the other hand are only dependent on measurements that include itself. For completeness, we also give the observability matrix of 'one at a time' measurement method for

three vehicles:

$$\tilde{\mathcal{O}}_{OAT} = \begin{bmatrix} [\mathbf{H}_{\varrho_1}(\Phi_{t_1} - \Phi_a)]_{\tau_0} & [\mathbf{H}_{\varrho_1} \Phi_{t_1}]_{\tau_0} & \mathbf{0}_{m,n} & \mathbf{0}_{m,n} \\ [\mathbf{H}_{\varrho_2}(\Phi_{t_2} - \Phi_a)]_{\tau_1} & \mathbf{0}_{m,n} & [\mathbf{H}_{\varrho_2} \Phi_{t_2}]_{\tau_1} & \mathbf{0}_{m,n} \\ [\mathbf{H}_{\varrho_3}(\Phi_{t_3} - \Phi_a)]_{\tau_2} & \mathbf{0}_{m,n} & \mathbf{0}_{m,n} & [\mathbf{H}_{\varrho_3} \Phi_{t_3}]_{\tau_2} \\ \vdots & & \vdots & \end{bmatrix} \quad (3.31)$$

Equations (3.30) and (3.31) are nearly equivalent if the same number of measurements are taken of each vehicle, with the only difference being the timing of evaluations. Most importantly, both equations are clearly a linear variations of the matrix $\tilde{\mathcal{O}}_{SSAT}$ and thus incorporating the state transition tensor approximation leads to equivalent observability properties.

3.3.4 Spaceflight Specific Modeling

Observability of SSAT for a DSS can be further examined by substituting the generalized dynamics and measurements in the observability derivation with orbital motion and common spacecraft sensor payloads. Examining orbital motion may lead to specific conditions for observability based on the gravitational potential that drives motion. Examining the measurement sensitivity matrix identifies specific geometries where singularities may occur. In regards to measurements, we are concerned with range-only and optical-only sensors as these are standard payloads for spacecraft. The range-only sensor includes both range and range-rate data. The optical-only sensor is represented as a right ascension and declination angle pair.

Spaceflight dynamics are modeled as a second order ordinary differential equation and therefore position and velocity are required to integrate the state. Thus, let the spacecraft state be composed of position and velocity vectors be given by the vectors $\mathbf{r}_i \in \mathbb{R}^3$ and $\mathbf{v}_i \in \mathbb{R}^3$, where the subscript i is denotes the spacecraft. Then the SSAT and absolute-relative systems become:

$$\mathbf{x} = \begin{bmatrix} \mathbf{r}_a^T & \mathbf{v}_a^T & \mathbf{r}_t^T & \mathbf{v}_t^T \end{bmatrix}^T \quad (3.32)$$

$$\tilde{\mathbf{x}} = \begin{bmatrix} \boldsymbol{\rho}_\alpha^T & \boldsymbol{\nu}_\alpha^T & \boldsymbol{\rho}^T & \boldsymbol{\nu}^T \end{bmatrix}^T \quad (3.33)$$

Spaceflight dynamics primary driving force is from a gravitational potential field $U(\mathbf{r}_i)$. The potential field only relies on position and therefore the Jacobian only has a positional dependency if no perturbing effects are included and the dynamic frame is inertial. If the aforementioned cases are true then the Jacobian becomes:

$$\mathbf{f}(\mathbf{s}_i) = \begin{bmatrix} \mathbf{v}_i^T & -\nabla U(\mathbf{r}_i)^T \end{bmatrix}^T \quad (3.34)$$

$$\mathbf{A}_i = \frac{\partial \mathbf{f}}{\partial \mathbf{s}_i} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \\ -\nabla \nabla U(\mathbf{r}_i) & \mathbf{0}_{3 \times 3} \end{bmatrix} \quad (3.35)$$

It is very important to note the the bottom right hand corner will not be zero if the dynamic frame is rotating, like in the CR3BP which is simulated later. If the rotation is constant, again as in the CR3BP, then the bottom right matrix is the matrix cross product representation of the angular velocity.

From Eq. (3.34) and (3.12) it is evident that two STMs will be equivalent if the potential is the same on both spacecraft for all time, a fact which was originally derived by Markley [4]. If the STMs are equivalent then the SSAT system cannot be absolutely observable as evident from Eq.(3.20). In the two body problem this means that two orbits with matched semi-major axis, eccentricity, and phasing will not be absolute observable. When additional perturbations are included then the requirement is generalized such that the spacecraft must experience the same forces for all time. This requirement is significantly more challenging to both identify and satisfy.

3.3.4.1 Range-Only Measurements

Range-only measurements sensors, which consist of range and range-rate data, have their equations given by Eq. (2.33). Replacing the generic relative sensitivity \mathbf{H}_ρ with the range-only sensitivity \mathbf{H}_ρ from Eq. (2.35) gives rise to one obvious unobservable edge case. Given two planar orbits, the z-components of position and velocity are zero for all time and therefore the out of plane sensitivity is zero for all time as a consequence. This means the out of plane motion is unobservable given planar motion for range-based measurements. If it is known that spacecraft motion is planar

then this doesn't create any practical issues since the out of plane state is neglected. If it is not known that spacecraft motion is planar then the system appears unobservable. The resulting unobservability of the 3D system with range-only sensors was empirically demonstrated by Hill [43].

3.3.4.2 Optical-Only Measurements

Optical-only measurements are represented with a right-ascension and elevation angle pair as given by Eq. (2.36). The optical-only measurement sensitivity with respect to the relative state, denoted such that $\mathbf{H}_\rho = \mathbf{H}_\theta$, is given by Eq. 2.40. The optical-only sensitivity \mathbf{H}_θ only has zero out of plane component during the elevation singularity at $\pm\pi/2$. This single unobservable edge case for optical sensors is in contrast to the range-only measurements which is always unobservable for planar motion. In practice the singularity can be avoided with an alternative formulation. The non-zero out of plane sensitivity, along with the non-cooperative sensing, are strong arguments for equipping spacecraft with optical-only measurements over range-only measurements in a SSAT setting.

A comment must be made about obtaining a relative range information from angle measurements. While it is not immediately obvious that range may be unobservable from the sensitivity matrix, it is well known to be so in linear systems. This unobservable case includes linearization of spaceflight dynamics for nearby relative motion [35]. Fortunately, when nonlinear dynamics are used then the relative range becomes observable [37, 28]. Effectively, the higher order moments can create a unique motion profile which is locally observable if the system is properly designed. Local observability for nonlinear systems is made more evident by examining Eq. (3.25) which incorporates the initial deviation, including the range state, via the Einstein sum with the higher order moments in both sides of the matrix. Unfortunately, the inclusion of nonlinear dynamics does not guarantee observability, it only enables it, and case by case studies must be instituted to determine if the nonlinearities produce observability for particular trajectories.

3.4 Simulation and Results

With the theory developed, we move on to generating predicted navigation performances for systems using optical-only SSAT which is simulated via a linear covariance analysis. The reported results are position and velocity norms of the 1-sigma values in the absolute-relative SSAT information matrix $\tilde{\mathbf{A}}$ from Eq. (3.17). The matrix is constructed from the corresponding observability matrix $\tilde{\mathbf{O}}$ calculated from Eq. (3.26). If the observability matrix is full rank then the information matrix is inverted using Eq. (3.18) to calculate the covariance of the state estimate. If the matrix is not invertible then the simulation is said to be numerically unobservable. Additionally, any case where the uncertainty norms are greater than 10 km or 10 m/s, or less than 1e-12 km and 1e-12 km/s, then the simulation is also disregarded due to numerical instability as these cases produced inconsistent results.

We begin with the simplest, and least likely to be observable case, two vehicles operating under unperturbed Keplerian motion. In the Keplerian setting the STT approximation of the observability matrix is tested, and navigation maps are constructed by varying delta orbit elements of the two vehicles. Next, perturbing forces and cislunar dynamics are explored to examine their effects on navigation accuracies. Finally, simulations incorporating multiple target vehicles are examined.

To minimize variability between simulations there are several constant settings across simulations. First, the agent spacecraft is responsible for observing the target with an optical-only sensor for a fixed duration and number of measurements. Next, measurements are only limited by line-of-sight which can be occluded by planetary bodies. The measurements noise is set to 10 micro-rad which is representative of a high accuracy optical sensor as described by Reference [33]. Lastly, Earth's gravitational parameter is used to simulate Keplerian motion, and the Earth-Moon gravitational parameter is used for the CR3BP.

3.4.1 SSAT Under Keplerian Motion

Keplerian motion, or the two body problem (2BP), is used to approximate near-Earth missions with an agent and target pair on ballistic trajectories. The acceleration for a single vehicle under the gravitational potential of a single attracting body is given by the first term of Eq. (2.1) and sets the perturbing acceleration to zero such that $\mathbf{a}_{pert} = \mathbf{0}$. All motion is propagated in an inertial frame and transformed into absolute-relative states using the definitions earlier. The agent and target are propagated for 5 orbit periods with 10 evenly spaced measurements every orbit.

3.4.1.1 Tensor Approximation Validation

The validity of the STT approximation made in the observability derivation is examined by calculating the difference between STT observability matrix $\tilde{\mathcal{O}}_{STT}$ obtained from Equation (3.25), from the true SSAT observability matrix $\tilde{\mathcal{O}}_{SSAT}$ obtained from Equation (3.20). The normalized error is calculated using:

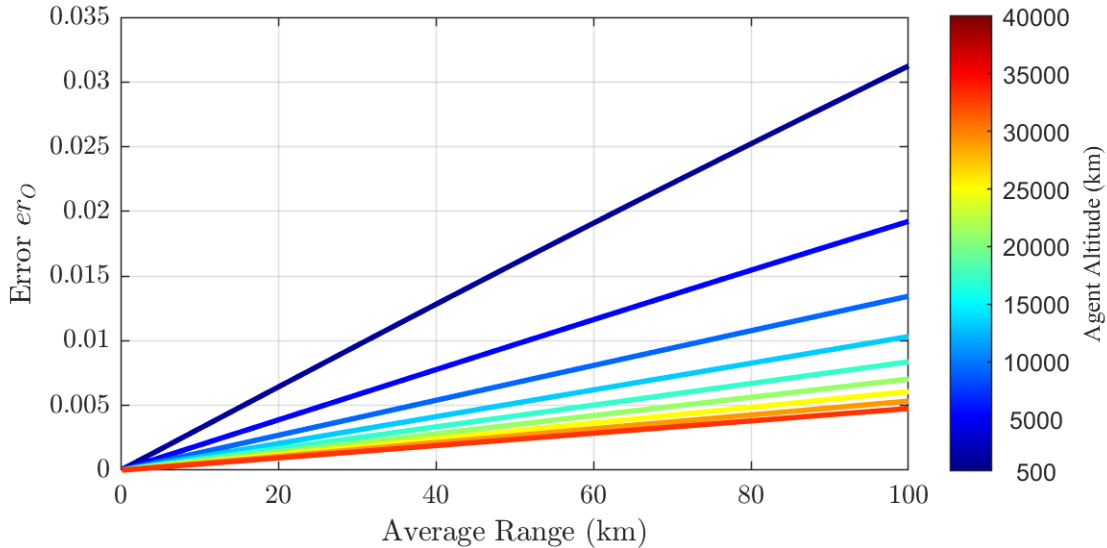
$$\text{er}_{\mathcal{O}} = \frac{\left\| \tilde{\mathcal{O}}_{STT} - \tilde{\mathcal{O}}_{SSAT} \right\|}{\left\| \tilde{\mathcal{O}}_{SSAT} \right\|} \quad (3.36)$$

This equation captures the size of the error of the STT observability matrix, and normalizes by the true SSAT observability matrix.

The agent is initialized on circular Earth orbits with an inclination of 15 degrees and an orbit altitude between 500 km to 40000 km. The target is placed on a near-by period matched orbit using the closed relative motion solutions to the Clohessy-Wiltshire (CW) equations to ensure bounded relative motion [22]. The average range of the target orbit is increased from 0.1 km to 100 km and propagated for 5 agent orbit periods. The STT error is calculated using Eq. (3.36) over all orbit combinations and is recorded in Figure 3.1.

The error is near zero for all near-by orbits under 10 km and grows linearly as the average range increases. This low error over short distances with growth as range increases is indicative of the Taylor series expansion used in the STT matrix. The error grows the fastest for the low Earth orbits because lower orbits behave more nonlinearly due to gravity being proportional to

Figure 3.1: Normalized STT observability matrix error as calculated by Eq. (3.36).



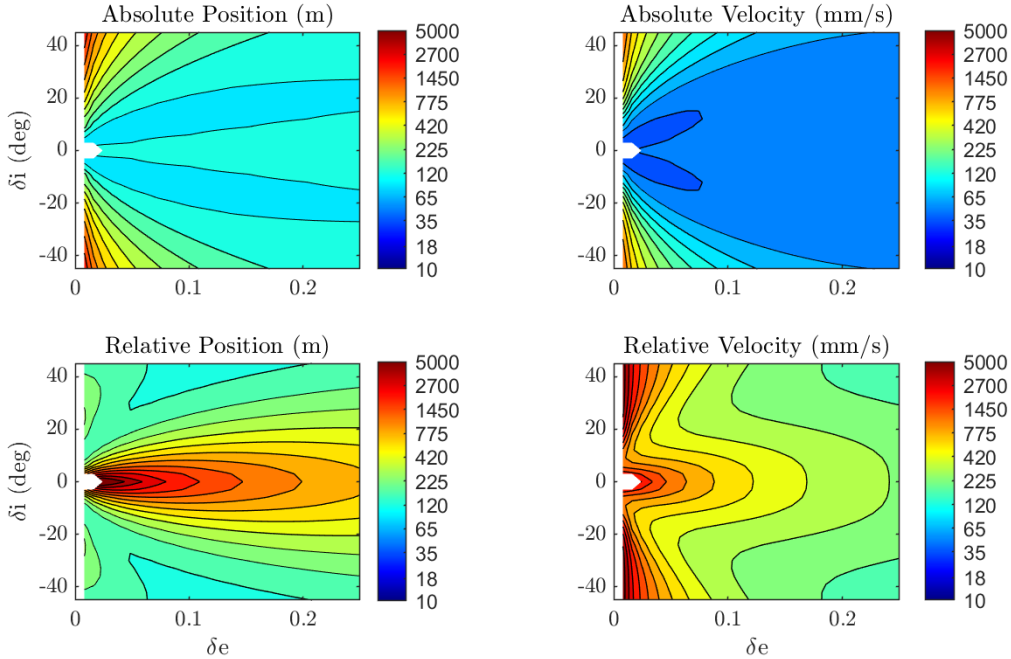
inverse of the radius squared. Thus, as expected, these low orbits have greater truncation errors and this error decays as the agent altitude increases. Most importantly these results indicate that the STT matrix is accurately approximating the SSAT matrix, and that STT approximation in the observability derivation is valid.

3.4.1.2 Linear Covariance Navigation Maps

Predicted navigation accuracies for optical-only SSAT navigation are studied by varying relative orbital motion between the agent and target crafts. The agent remains fixed in a given orbit and the target is varied using delta orbit elements such that $\alpha_t = \alpha_a + \delta\alpha$ where α_i is an array of orbit elements for a given vehicle. For the remainder of the Keplerian results the agent is placed on a circular medium Earth orbit with an altitude of 10000 km and an inclination of 0 deg. The target is given matching orbit elements to begin, and then two elements are varied to create navigation maps. The maps represent 1-sigma navigation uncertainties from the linear covariance matrix, with a minimum uncertainty in blue and maximum uncertainty in red. The colormaps are always be on a log scale due to large differences in navigation accuracies. Uncolored regions

represent conditions where the SSAT problem is poorly posed such that \mathbf{A} is not invertible and the system is numerically unobservable. The first of the delta element maps are delta eccentricity and inclination, resulting in Figure 3.2.

Figure 3.2: Mapping of δe and δi to expected 1σ SSAT uncertainty for Keplerian motion.

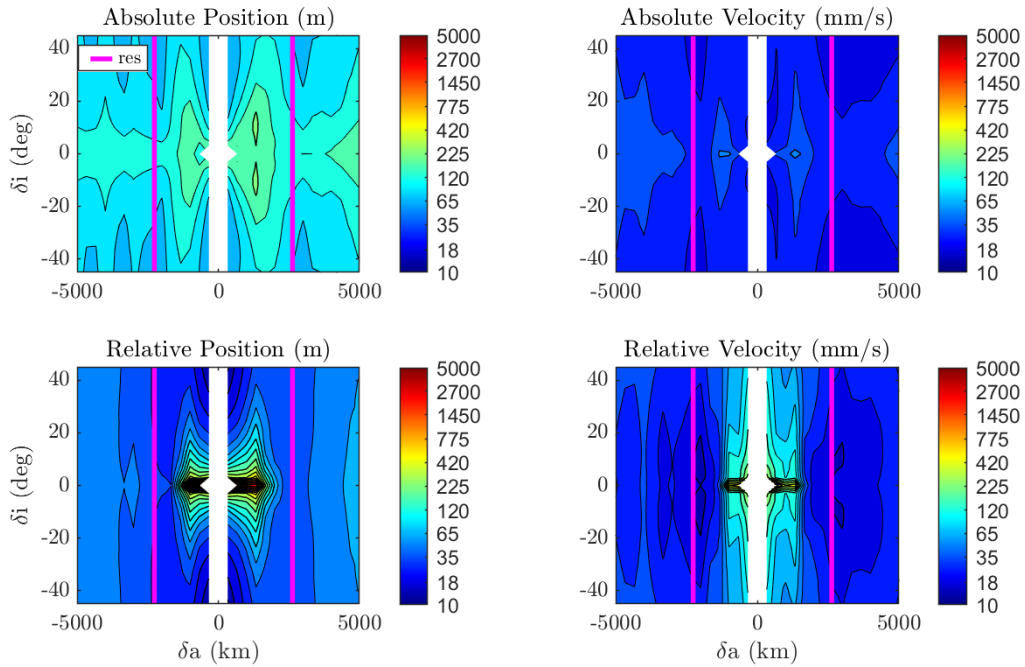


The simulations from Figure 3.2 have the matched semi-major axis for the agent and target so that they are period matched and remain in constant formation for all time. An immediate conclusion from the figure is that when the difference in eccentricity (δe) is too low, then the problem is unobservable because the absolute state cannot be estimated. This unobservable set is a validation of Markley’s conclusion that matched semi-major axis, eccentricity, and phased orbits are not estimable. Outside of zero δe there are cases with low δi that are also unobservable, this time due to relative range issues. The remaining cases always achieve observability. The fact that there is a large variety of period matched orbits that still obtain observability is promising for actual application of optical-only SSAT.

For the absolute states it is seen that increasing δe generally leads to better estimates. The δi

variation on the absolute estimates on the other hand has two symmetric minima, which is expected because the two body problem is symmetric about inclination. The appearance of these minima is likely the balancing point between nonlinear motion creating information and optical sensor accuracy degrading over larger distances. Larger ranges lead to greater dynamical differences and better information, but optical sensors are also weaker over longer ranges because the angular noise converts to larger arc lengths. Examining the relative state, changes in inclination are more impactful than changes in eccentricity. In fact, the relative state is poorly conditioned when δi is zero.

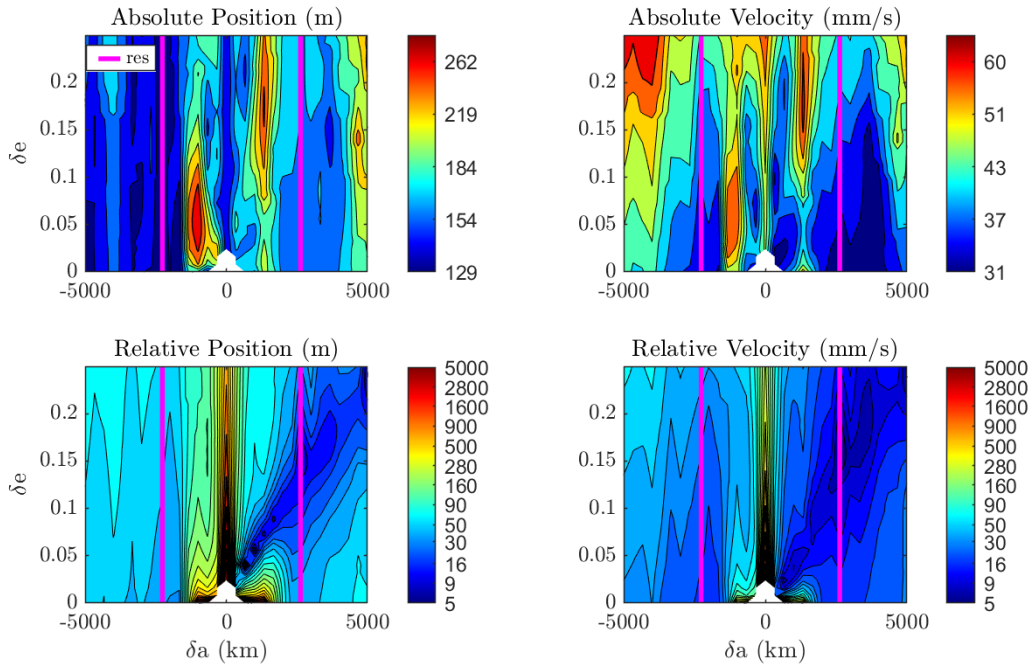
Figure 3.3: Mapping of δa and δi to expected 1σ SSAT uncertainty for Keplerian motion.



Now semi-major axis is allowed to vary with inclination to construct δa , δi maps which are displayed in Figure 3.3. Once again, there is an unobservable case when δa and δi are zero. Increasing the change in inclination effectively improves the state estimate for all cases. Changes to δa are more complex and create peaks and valleys which appear to correspond to resonances. The two valleys most clearly identified in the absolute positing navigation plots are closest to the

4/5 and 5/4 resonances, which are marked by magenta lines in the figure. These resonances are likely the most obvious because the simulation length is 5 orbit periods and therefore creates a closed viewing geometry. Finally, the general range of uncertainties are smaller than the previous inclination and eccentricity map which is expected because changing semi-major axis leads to spacecraft drift allowing for further nonlinearity buildup to enhances state estimates.

Figure 3.4: Mapping of δa and δe to expected 1σ SSAT uncertainty for Keplerian motion.



Finally, maps of changes in semi-major axis and eccentricity to navigation accuracy are displayed in Figure 3.4. Please note that the colormap scale for the absolute estimates has been changed because the range of values are much smaller and would not show the appropriate information if left in the previous scale. In these maps the origin is unobservable as expected. In the absolute state the interplay between δa and δe create pocketed maps. There are two peaks of higher uncertainties centered around δa at -1000 km and δe of 0.05, and δa at 1400 km and δe of 0.17, which should be avoided. There are also two valleys, again around the 4/5 and 5/4 resonances, which leads to lower uncertainties which should be taken advantage of.

The relative states behave more consistently than the absolute states with increases in δa leading to general navigation improvement. Additionally, the resonances make themselves very clear with local valleys on the semi-major axis. Finally, and unique to this mapping, there is a unique linear relationship between eccentricity and semi-major axis which produces the best results. The reason for this is unclear, but these geometries are best for breaking the range ambiguity and warrants further study.

The range of navigation accuracies across all maps has been tabulated in Table 3.1 for a more direct comparison. When semi-major axis is allowed to vary, then the relative states have lower bounds than the absolute state which is expected because the relative state is being directly measured. If the semi-major axis is matched though it appears that the fixed orbits do not establish a relative range estimate as easily and produce worse estimates.

The table also depicts the importance of differences in $\delta a, \delta i$ for obtaining absolute information, while differences in $\delta a, \delta e$ aid in obtaining relative information. In both cases δa is an important factor for getting better estimates, which is expected because allowing drift between the spacecraft means their respective dynamics will have greater differences that lead to more information accumulation.

Table 3.1: Range of navigation uncertainty for delta orbit element maps.

| Model | Maps | 1- σ ranges | | | |
|----------|----------------------|--------------------------------|----------------------------------|-----------------------|-------------------------|
| | | $\ \sigma_{\rho_\alpha}\ $ (m) | $\ \sigma_{\nu_\alpha}\ $ (mm/s) | $\ \sigma_\rho\ $ (m) | $\ \sigma_\nu\ $ (mm/s) |
| 2BP | $\delta e, \delta i$ | 108 - 2514 | 45 - 1443 | 148 - 8944 | 194 - 5887 |
| | $\delta a, \delta i$ | 68 - 241 | 19 - 54 | 13 - 6702 | 16 - 978 |
| | $\delta a, \delta e$ | 129 - 286 | 31 - 65 | 6 - 8944 | 5 - 2728 |
| 2BP + J2 | $\delta e, \delta i$ | 108 - 570 | 44 - 1046 | 42 - 8979 | 193 - 9809 |
| | $\delta a, \delta i$ | 67 - 498 | 19 - 1046 | 13 - 6008 | 16 - 9809 |
| | $\delta a, \delta e$ | 129 - 286 | 31 - 65 | 6 - 8979 | 5 - 2739 |

3.4.2 SSAT With J2 Perturbation

The J2 spherical harmonic is added to assess how its perturbation changes the navigation maps developed in the previous section. The inclusion of J2 was predicted to enhance observability

by Psiaki [8]. To verify this prediction empirically the two body acceleration has the perturbing acceleration $\mathbf{a}_{pert} = \mathbf{a}_{j2}$ from Eq. (2.4).

The navigation maps from the previous section are reconstructed over the same delta elements with the new force model. The addition of J2 means that orbit elements will oscillate and drift with time, but the simulations are initialized using instantaneous elements that produce identical initial conditions to the previously simulated results. Again, the agent is initialized in a 10000 km altitude circular orbit with 0 degrees inclination, and the initial target elements are varied with respect to the agent. The maps are contained in Figure 3.5-3.7. The range of the uncertainties are also reported in Table 3.1.

There are two primary verdicts from the figures and table. First, adding J2 allows for additional configurations to be observable, as predicted by Psiaki, but the navigation estimate for these newly observable configurations are rather poor. Second, for already observable orbits, the addition of J2 does not make a noticeable impact on overall navigation accuracy. Qualitatively, this is likely because J2 is only a perturbation to Keplerian motion and therefore only lightly perturbs the overall results.

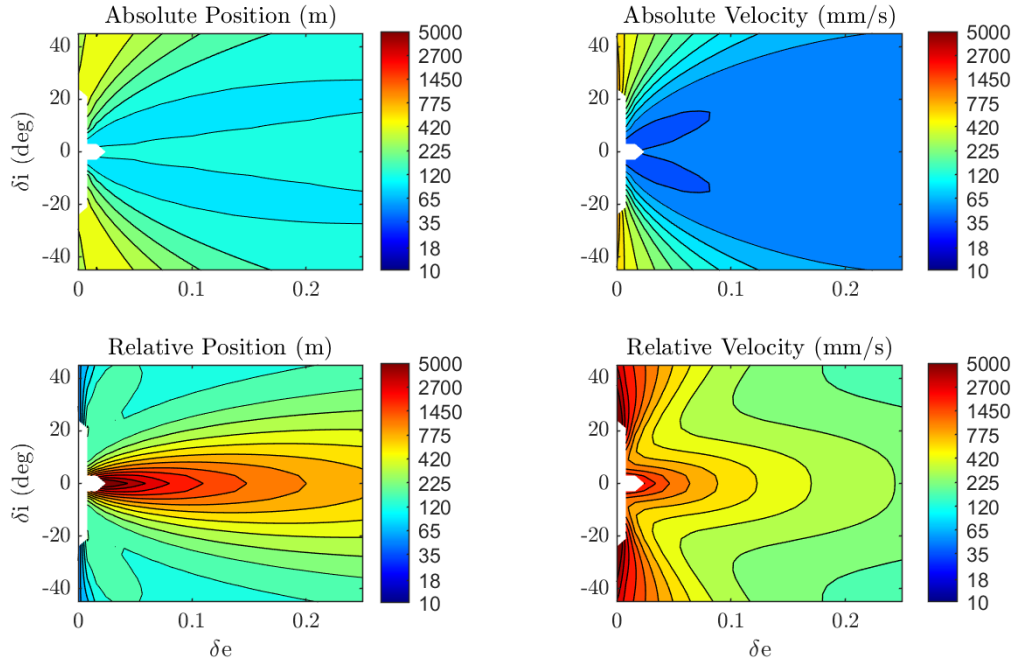


Figure 3.5: Mapping of δe and δi to expected 1σ SSAT uncertainty with J2 perturbed Keplerian motion.

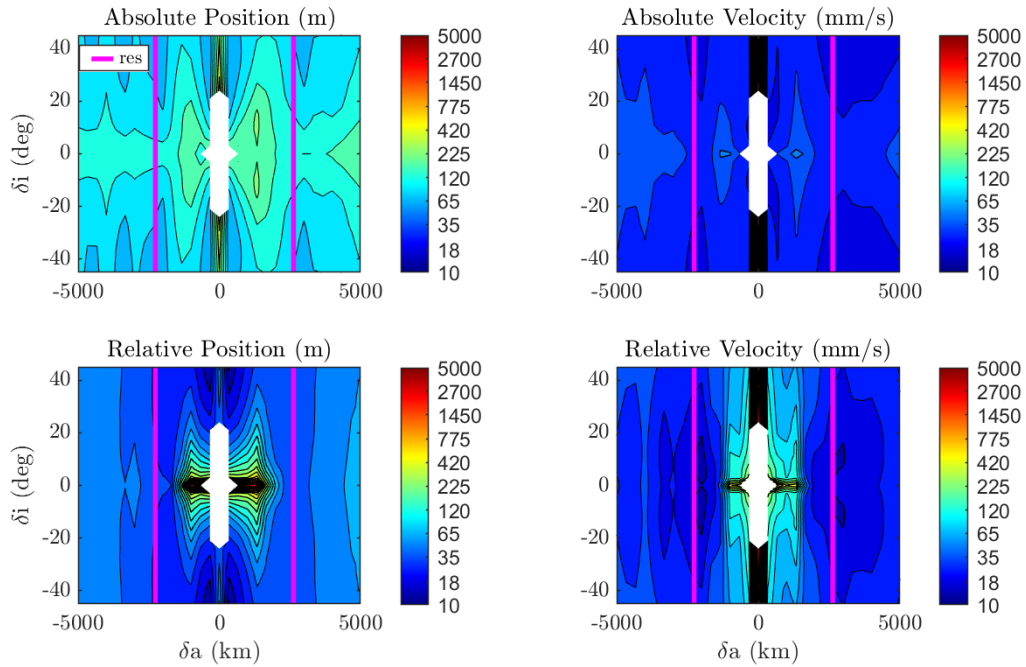


Figure 3.6: Mapping of δa and δi to expected 1σ SSAT uncertainty with J2 perturbed Keplerian motion.

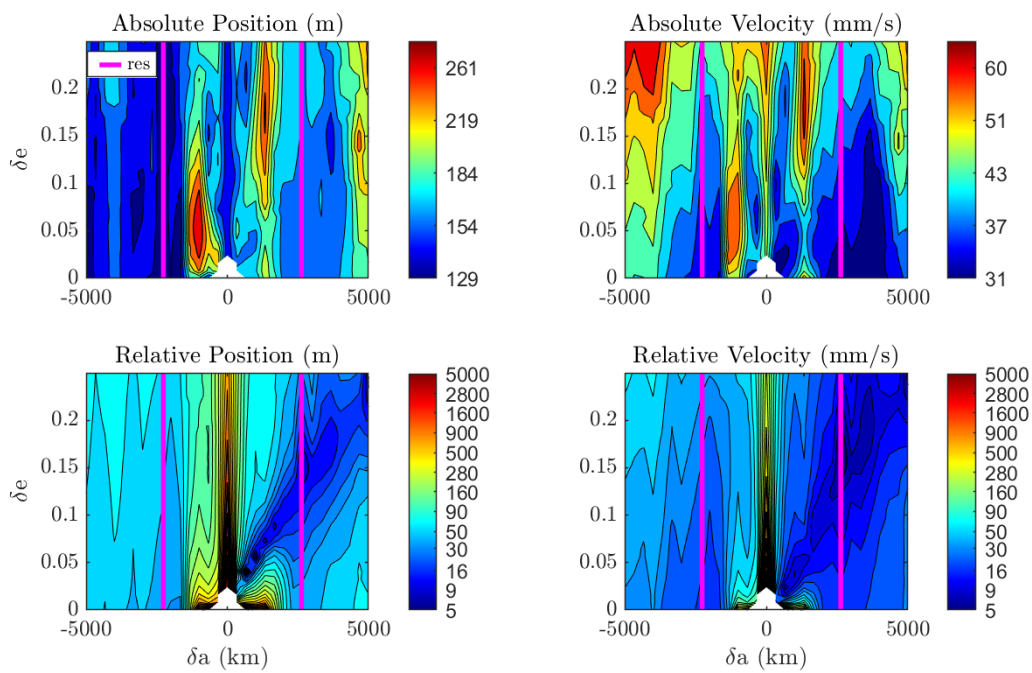


Figure 3.7: Mapping of δa and δe to expected 1σ SSAT uncertainty with J2 perturbed Keplerian motion.

3.4.3 SSAT in the CR3BP

Moving beyond a near-Earth space, we examine SSAT in the cislunar environment by using CR3BP dynamical model. The dimensional CR3BP force equations in the rotating frame can be found in Eq (2.9) and the non-dimensional equations in Eq (2.14). The following work uses the non-dimensional equations of motion to reduce numerical issues and convert to dimensional products for reporting results.

Because the CR3BP has no analog to orbit elements the navigation maps are constructed by varying the target over orbit families while the agent is fixed to an orbit of interest. The selected orbit for the agent is the 9:2 synodic southern L2 near rectilinear halo orbit (NRHO) which is currently the selected orbit NASA’s Lunar Gateway. This 9:2 NRHO is depicted as a magenta orbit in the following maps. The maps encode the final covariance after 2 agent orbit periods with 60 measurements per orbit. The use of fewer orbits with more measurements reflects the fact that these orbits have much longer periods than the explored Earth orbits, on the order of weeks instead of hours. The final uncertainty is used as a color map on the agent orbit along the family.

The simulations vary the target over the southern L2 orbit family, and the results are given in Figure 3.8. There are three conclusions we have drawn from these simulations. First, all cases are always observable and have accuracies less than 1 km in position and 10 mm/s in velocity. It is very important to note that these results are for apoapsis where orbital velocity is low. Second, as distance increases the relative navigation accuracy decreases monotonically. This is expected since optical measurements are stronger over shorter ranges. Finally, the absolute state accuracy contains 2 local peaks and valleys before beginning to monotonically increase. The existence of the minima likely are products of a balancing between larger distances enhancing nonlinear motion while decreasing the efficacy of optical sensor accuracy in position space. Eventually moving far enough away degrades sensor performance enough that it outweighs the increase in nonlinearity.

Next, the target is varied over the northern L1 halos and the planar Lyapunov L2 orbits. The results for each are in Figures 3.9 and 3.10 respectively. The northern L1 orbits show similar trends

to the southern L2 family but with generally higher uncertainties overall. Reduced navigation performance is expected because these orbits have greater ranges, and as range increases optical sensors noise has a greater impact as seen in the previous results. Critically, all orbit combinations remain observable with desirable accuracies.

The planar L2 families show less promising results than both the northern L1 and southern L2 orbit families. The poor accuracy is likely a reflection on the highly linear and completely planar motion of the planar orbits which makes observability difficult and numerical differencing poorly conditioned. The hypothesis that linear motion is a culprit is corroborated by the fact that the largest and most irregularly shaped orbits have the lowest uncertainties.

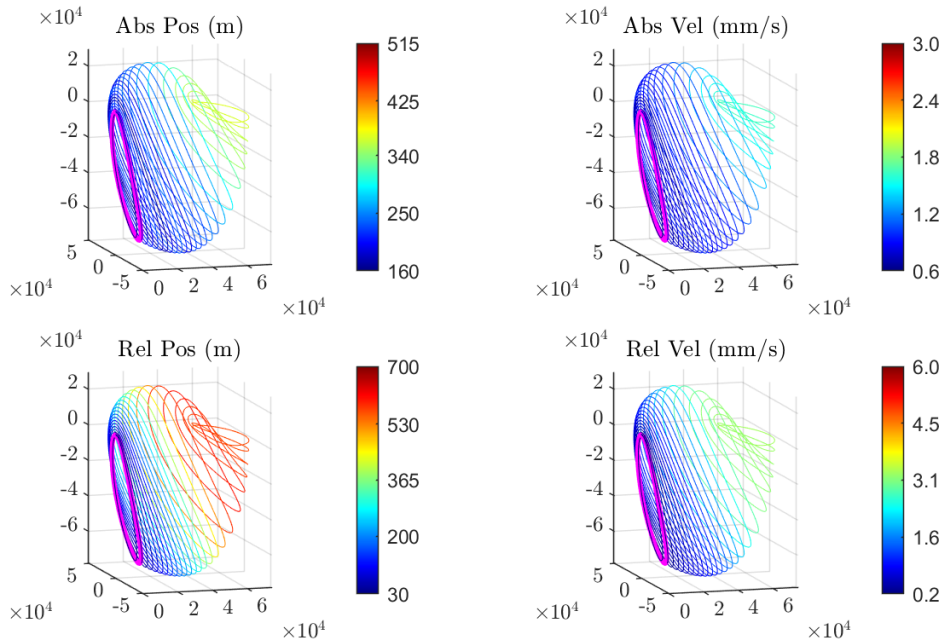


Figure 3.8: SSAT navigation map for an agent on the 9:2 NRHO and target on the southern L2 halo family in the CR3BP.




Figure 3.9: SSAT navigation map for an agent on the 9:2 NRHO and target on the northern L1 halo family in the CR3BP.

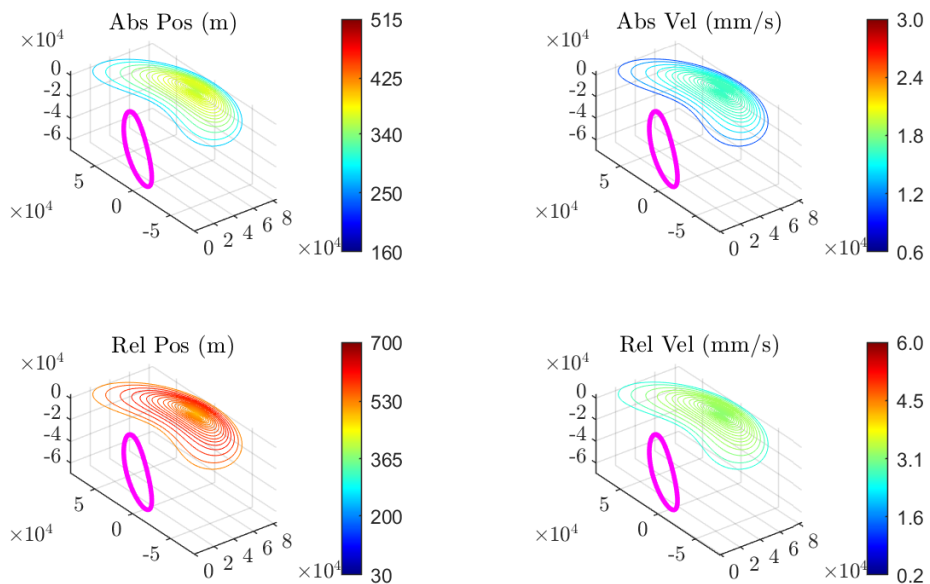


Figure 3.10: SSAT navigation map for an agent on the 9:2 NRHO and target on the planar L2 Lyapunov family in the CR3BP.

3.4.4 SSAT of a Multi-Target Constellation

The previous simulations only considered a single agent-target pair, but a DSS often has multiple target crafts to consider. This section now considers including multiple targets to be observed in a constellation of spacecraft. This multi-target simulation is assessed with an agent-based absolute state with the notation and equations introduced in Section 3.3.3. Introducing multiple targets creates a combinatoric problem that quickly leads to intractable maps, so only two simulations under pure Keplerian motion are examined. The focus of these simulations is to assess the impact of adding successive vehicles into a constant formation DSS. Constant formation means that only closed relative orbits, equivalently expressed as period matched orbits with the same semi-major axis, are considered.

The period matching constraint means these simulations are expanding on the $\delta e, \delta i$ maps that were previously built, and thus the map from Figure 3.2 is leveraged to instantiate a configuration which has favorable navigation properties. To implement this design choice the agent must inhabit a circular orbit with a semi-major axis of 16378 km and an inclination of 0 degrees to match the map simulations. The targets are given the same semi-major axis to implement the period matching condition. Then the targets are given a δi of -30 degrees and varying eccentricities from 0.01 to 0.2 degrees to balance the absolute and relative accuracy. Again, the observing agent takes 10 observations per orbit for a total of 5 orbit periods. Figure 3.11 depicts all the targets centered about the agent in a LVLV frame for 5 orbit periods.

Table 3.2 reveals the expected navigation accuracies for the multi-target simulations, and highlights two primary trends. First, when observation power is not a concern and all vehicles can be measured as in the 'all at once' simulations, then including additional targets always improves state estimation uncertainties for all vehicles. The systematic improvement is due to the fact that the agent is better able to locate itself with multiple targets, which translates into better estimates for the targets as well. The second trend is that when observation power is limited, and not all targets can be easily acquired as in the 'one at a time' simulations, then adding targets can degrade




Figure 3.11: Relative orbits of targets 1-4 plotted in a LVLH frame centered about the agent.

estimation. The worsening of the estimates is likely because insufficient information is gathered on each target and the agent is never capable of localizing itself relative to any of them. Clearly in the latter case optimal sensor tasking is of critical importance to determine a schedule which will provide the best navigation solution. Modern techniques for optimal sensor tasking are broad, but Monte-Carlo tree search methods are a promising approach which can be applied on top of these simulations [44].

3.5 Conclusion

Spacecraft to spacecraft absolute tracking poses absolute navigation for a distributed systems using relative only sensing and shows that observability for all vehicles can be obtained if the systems is properly designed to leverage nonlinear motion. This conclusion is reached by deriving the linear observability matrix for the SSAT system where the target is approximated as a deviation from the agent. Doing so allows the target dynamics to be re-written as a function of the target dynamics and directly leads to the necessity of second order moments in the dynamics for observability. Additionally, it is analytically shown that range measurements will have unobservable out of plane

Table 3.2: Navigation accuracies with multiple targets given orbit geometries in Figure 3.11.

| Targets Observed | # of Targets | 1- σ position, velocity norms (m,mm/s) | | | | |
|------------------|--------------|---|----------|----------|----------|----------|
| | | Agent | Target 1 | Target 2 | Target 3 | Target 4 |
| All at Once | 1 | 94, 121 | 15, 230 | - | - | - |
| | 2 | 65, 83 | 11, 156 | 51, 163 | - | - |
| | 3 | 56, 65 | 9, 120 | 40, 125 | 78, 131 | - |
| | 4 | 51, 50 | 7, 90 | 31, 94 | 60, 99 | 110, 107 |
| One at a Time | 1 | 94, 121 | 15,230 | - | - | - |
| | 2 | 113, 131 | 23, 243 | 80, 253 | - | - |
| | 3 | 100, 114 | 27, 209 | 70, 219 | 137, 229 | - |
| | 4 | 130, 116 | 33, 205 | 70, 216 | 140, 227 | 252, 243 |

state when given planar motion, while optical measurements will still retain full observability.

Results for SSAT are generated via a linear covariance analysis to demonstrate potential navigation estimates using optical-only tracking. First, pure Keplerian motion is examined and shown to have an unobservable absolute state if eccentricity and semi-major axis are matched. Additionally, low changes in inclination lead to poor relative estimation. Otherwise, given sufficiently different orbit elements, a wide variety of simulations show observability and favorable navigation estimates. Then the J2 spherical harmonic is added, and while some more cases become observable, it does not qualitatively change the accuracy of the results. Cislunar space, approximated with CR3BP dynamics, always shows observability and favorable navigation estimates for nonplanar orbits. Finally, including multiple targets has the opportunity to enhance navigation if sufficient observational capabilities are available. If observation power is limited, then a sensor tasking algorithm would greatly benefit a multi-target scenario.

Overall, we show that relative optical-only measurements can be used to simultaneously estimate the absolute and relative state of a distributed space system. This navigation technique is termed spacecraft to spacecraft absolute tracking and operates autonomously of external observation networks. The SSAT method is a promising approach for autonomous navigation of distributed systems as long as spacecraft motion is different enough. This may preclude proximity operations unless methods to gather information are considered, which is the topic of the next chapter.

Chapter 4

Autonomous Guidance for Information Gathering with Optics

4.1 Introduction and Background on Information Gathering

Information gathering to improve spacecraft state estimation is a critical capability for deep space spaceflight missions where orbit determination is not readily available, or for autonomous operations where real time navigation is necessary. State estimates can be improved through the intelligent use of control to reposition sensors in a manner to better observe desired targets. The idea of using control to obtain state knowledge has been explored since at least the 1960's starting with dual control theory, which has been first credited to A.A. Feldbaum [45]. The fundamental principle of dual control is to shape trajectories to initially minimize system uncertainty such that later control effort more accurately reduces tracking errors.

The principles of dual control are widely applicable and have been used in many systems to improve overall performance [46, 47, 48, 49]. These principles also appear across many diverse fields under various names. For example, within the artificial intelligence and machine learning domain is the closely related concept of exploration vs exploitation. The only way to discover an optimal policy is to explore, and once enough information is gained then the optimal policy can be exploited to gain the maximal reward [50, 51]. This topic also has a background in robotics [52], controls [53], and missile guidance [54]. Broadly speaking, gathering information to better understand the environment before acting tends to improve overall system performance.

The concept of information gathering is especially important for systems which are slow to accrue information, such as single sensor systems using optical-only measurements. Optical-only

measurements, also commonly referred to as angle-only or bearing-only measurements, are appealing because they use lightweight, low-cost, and widely available technology with an abundance of literature and application. The primary weakness is that they lack range information, which can lead to observability issues in linear systems [9, 35]. Optical-only observability has been studied at length for various systems, and while some systems may be locally observable, there is a consensus that obtaining range information is always an important consideration [55, 10, 56, 57, 11].

Optical-only observability and tracking is an influential problem within spaceflight navigation literature because optics often play a primary role during rendezvous and proximity operations. If a DSS can demonstrate that it only requires onboard optical sensors then it would reduce reliance on Earth-based resources, increase reliability under limited communication scenarios, and potentially enable novel missions. Fortunately, as demonstrated in the previous chapter, SSAT offers a navigation solution which can provide a full state estimate of every vehicle in a distributed space system without the need for external sensors. Unfortunately, there are conditions which are completely unobservable, or which may require significant time for information to accumulate. An important example of such conditions include proximity operations where linearized motion is valid. Thus, methods to gather information for single sensor systems is still a critical capability even in a SSAT setting.

Obtaining range information with optical sensors during spaceflight proximity operations has been well studied in low Earth orbit. Woffinden proved that optical-only measurements cannot produce range observability under linearized motion [35]. Then Woffinden and J. Grzymisch went on to develop optimal maneuvers to gather the necessary range information [12, 13]. More recent work has extended information gathering for optics into new environments, such as cislunar space, with a variety of approaches [14, 58, 59, 60]. Crucially, the previous works either provide a situation specific solution, or a nonlinear solution that is costly to compute. Additionally, all the results assumed a well-known target state which is not as easily adopted into a SSAT setting. Therefore, there is room for a more generalized analytical approach which can be used in SSAT scenario for little computation cost.

The focus of this work is to develop low computation cost guidance policies to minimize a desired projection of state uncertainty, with the main contribution being an analytic guidance policy which maximizes measurement deviations as a surrogate to maximizing information orthogonal to the measurement space. The policy is developed through the lens of optical sensors which provides additional insight into range information gathering with optics. The paper is organized as follows. Section 4.2 outlines the problem formulation including the models and relevant geometry. Section 4.3 presents two developed guidance policies: a heuristic policy for optical sensors, and the main analytic information gathering policy. Section 4.4 simulates a covariance analysis of the policies in cislunar space with optical sensors to demonstrate expected performances. Cislunar space is used as a simulated proving ground for the developed guidance policies because it is a region of interest where traditional near-Earth methods may not be applicable. Section 4.5 concludes the paper with key results and future work.

4.2 Problem Formulation

The problem is formulated by first defining the system models, then the cost function, and finally a geometric analysis of optical variations due to a maneuver. The system model is a two vehicle SSAT state model with cislunar dynamics approximated by the CR3BP. The agent has access to optical measurements of the target to estimate the state of both vehicles. With the state space formulated, a generalized cost function for information gathering is presented. The general form of the cost function is refined specifically for optical sensors to minimize relative range uncertainty, which is the portion of the position space that is orthogonal to the measurement. The defined cost function is nonlinear and non-trivial to optimize, so in pursuit of analytic policies an analysis of optical variations due to a maneuver is carried out. The geometric analysis produces a relationship between angular measurement uncertainty and estimated range uncertainty which serves as a basis for guidance policies.

4.2.1 Dynamical Models

The spacecraft to spacecraft absolute tracking problem estimates the state of multiple objects in a distributed system while only relying on relative measurements between them. Let $\mathbf{s}_a \in \mathbb{R}^6$ be the state of the agent in the dynamic reference frame, which is responsible for generating measurements and is actively managed. Then let $\mathbf{s}_t \in \mathbb{R}^6$ be the state of the target in the dynamic reference frame, which operates independently of the agent so that cooperative planning is not required for this architecture. The states of the spacecraft are Cartesian position $\mathbf{r} \in \mathbb{R}^3$ and velocity $\mathbf{v} \in \mathbb{R}^3$ such that a single spacecraft state is $\mathbf{s} = [\mathbf{r}^T, \mathbf{v}^T]^T$. Because the SSAT problem seeks to provide state estimates for both vehicles, the SSAT system state is comprised of both the target and the agent sub-states and is defined by:

$$\mathbf{x} = \begin{bmatrix} \mathbf{s}_a^T & \mathbf{s}_t^T \end{bmatrix}^T \quad (4.1)$$

Estimating the dynamic state of both vehicles is equivalent to estimating a relative and absolute state simultaneously. Defining the absolute state as of the vehicles average $\mathbf{s}_\alpha = \mathbf{s}_a + \mathbf{s}_t / 2$ and the relative state as the difference $\mathbf{s}_\rho = \mathbf{s}_t - \mathbf{s}_a$ leads to a linear transformation \mathbf{T} to swap from the SSAT system to an absolute-relative system. The absolute-relative system and the linear transformation are defined as:

$$\tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{s}_\alpha^T & \mathbf{s}_\rho^T \end{bmatrix}^T \quad (4.2)$$

$$\mathbf{T} = \begin{bmatrix} \frac{1}{2}\mathbf{I}_{n \times n} & \frac{1}{2}\mathbf{I}_{n \times n} \\ -\mathbf{I}_{n \times n} & \mathbf{I}_{n \times n} \end{bmatrix} \quad (4.3)$$

The natural dynamical model that dictates spacecraft motion is given by the function $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{s})$. The natural dynamics model used to approximate cislunar motion is the CR3BP as depicted in Figure 2.2 with equations of motion given by Eq. 2.9. The agent spacecraft is equipped with the capability to execute control in the form of impulsive velocity changes. Control is denoted by the vector $\mathbf{u} \in \mathbb{R}^p$ and influences the state through the agent's control input matrix \mathbf{B}_a . The discontinuous state before and after the impulse is \mathbf{x}^- and \mathbf{x}^+ respectively. The state equations

for control are:

$$\mathbf{f}(\mathbf{s}) = \begin{bmatrix} \mathbf{v}^T & \mathbf{a}_{cr3bp}^T \end{bmatrix}^T \quad (4.4)$$

$$\mathbf{x}^+ = \mathbf{x}^- + \mathbf{B}_a \mathbf{u} \quad (4.5)$$

$$\mathbf{B}_a = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_3 & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \end{bmatrix}^T \quad (4.6)$$

Maneuver execution errors are exempted from simulations because this work seeks to identify ideal performance given perfect maneuvers. However, maneuver errors are briefly considered in the geometric analysis, and we show that execution errors and measurement errors sum in a root-mean-square to inflate the final uncertainty. In application execution errors are non-negligible, but because the agent is cooperative it should have access to accelerometers which can measure these errors. Previous works have shown that given a sufficiently accurate accelerometer then maneuvers can effectively improve state estimation [59]. Practically speaking, the maneuver errors will need to be balanced against information gain to determine the overall efficacy. Future work will explore the necessary accuracy required for maneuver executions and/or accelerometers such that errors do not overtake the information gain.

It is assumed that there is a nominal trajectory and that the linear variation of the SSAT system is estimated with a Gaussian uncertainty model with covariance \mathbf{P} . Because of the linear Gaussian assumption, the covariance is linearly propagated via:

$$\mathbf{P}_{i+1|i} = \mathbf{\Phi}(\tau_i, \tau_{i+1}) \mathbf{P}_{i|i} \mathbf{\Phi}(\tau_i, \tau_{i+1})^T \quad (4.7)$$

and transformed from the inertial to relative SSAT system using the linear transformation found in Eq. 3.6 via:

$$\tilde{\mathbf{P}} = \mathbf{T} \mathbf{P} \mathbf{T}^T \quad (4.8)$$

4.2.2 Measurement Models

Optical measurements are simulated via azimuth and elevation angles which are taken from the agent to the target. Measurement vectors $\mathbf{y} \in \mathbb{R}^m$ are obtained from the measurement function

$\mathbf{y}_\theta = \mathbf{h}(\mathbf{x})$ where the equations for optical measurements were given in Eq. 2.36. By using the above representation it is assumed that there is a separate image processing system that has identified the target and converted the image data into the dynamic reference frame via the necessary transformations, which in our case is the rotating Earth-Moon frame.

Measurements are taken at a set measurement cadence with zero-mean Gaussian noise. The measurement covariance is determined by a root-mean-square uncertainty model to account for partially resolved objects. The root-mean-square model combines a minimum angular uncertainty parameter $\sigma_{\theta,min}$ with a reference resolution uncertainty parameter D similar to work by Reference [59]. The equation for the angular uncertainty model is:

$$\rho^2 \sigma_\theta^2 = \rho^2 \sigma_{\theta,min}^2 + D^2 \quad (4.9)$$

$$\sigma_\theta = \sqrt{\sigma_{\theta,min}^2 + \frac{D^2}{\rho^2}} \quad (4.10)$$

The minimum angular uncertainty parameter is meant to account for all pointing errors, including spacecraft attitude errors which are not directly modeled. The value of $\sigma_{\theta,min}$ is set to 10 micro-rad, which was chosen to emulate a high accuracy pointing system based on Reference [33]. The resolution uncertainty D is 4 meters, which is approximately half the size of Lunar Gateway modules [61]. Note that even though these uncertainties are small, they are obtainable, and reducing the accuracy will reduce state accuracy results without invalidating the methods. The measurement cadence is set to once per hour, allowing for sufficient processing time and allowing for relative motion between updates. Finally, it is assumed that the agent always has line-of-sight to the target because the simulated scenarios are close proximity operations, and we are seeking to characterize ideal performance.

Once again utilizing the assumption that the linear variation of the state is estimated, a linear covariance analysis is adopted to incorporate measurement information into the uncertainty

matrix. The measurement update is therefore:

$$\mathbf{P}_{i+1|i+1} = (\mathbf{I}_n - \mathbf{K}\mathbf{H})\mathbf{P}_{i+1|i}(\mathbf{I}_n - \mathbf{K}\mathbf{H})^T + \mathbf{K}\mathbf{R}_{i+1}\mathbf{K}^T \quad (4.11)$$

$$\mathbf{K} = (\mathbf{P}_{i+1|i}\mathbf{H}^T)/(\mathbf{H}\mathbf{P}_{i+1|i}\mathbf{H}^T + \mathbf{R}_{i+1}) \quad (4.12)$$

4.2.3 Uncertainty Cost Function

Estimation challenges, such as unobservable sub-spaces, often arise when using a single sensor. Specifically, relative range is unobservable given optical-only sensors in a linear dynamical system [35]. For nonlinear dynamical systems range may or may not be locally observable [6]. Range is unobservable given linear dynamics because range is orthogonal to the optical measurement space and various initial conditions can produce equivalent measurement profiles since range is never directly observed. In general, the portion of the state space which is the most difficult to estimate is orthogonal to the measurement space and that space tends to have the greatest associated uncertainty. Therefore, it is desirable to maximize information gain on the portion of the state that is orthogonal to the measurement space to improve the current state estimate. This idea to maximize information, or conversely minimize uncertainty, on a desired projection of the state space forms the basis of the cost function to minimize.

First, a general cost function that aims to minimize an arbitrary set of projections is defined. Given an initial condition \mathbf{x}_0 and uncertainty \mathbf{P}_0 , find the optimal impulse \mathbf{u} that satisfies the following:

$$\begin{aligned} \min_{\mathbf{u}} \quad & J := \text{trace}(\mathbf{M}^T \mathbf{P}_f \mathbf{M}) \\ \text{subject to} \quad & \mathbf{u}^T \mathbf{u} - v^2 \leq 0 \end{aligned} \quad (4.13)$$

where \mathbf{P}_f is the final uncertainty, \mathbf{M} is a matrix of vectors containing the projections to be minimized, and $v \in \mathbb{R} \geq 0$ is the maximum allowable control imposed as an inequality. If the vectors in \mathbf{M} are unit vectors then the cost J is the sum of the variances along each projection. Note that the matrix \mathbf{M} needs to balance dimensional scale and tuning weight considerations if multiple projections are to be minimized.

When appraising an optical-only system then the generic cost function from Eq. 4.13 takes on a specific form. The sub-space orthogonal to optical measurements is along the relative range vector, and thus the projection of uncertainty to minimize is along the relative range vector. Relative range uncertainty is easiest to obtain from the absolute-relative SSAT uncertainty $\tilde{\mathbf{P}}$. To isolate relative range variance as a cost function the relative position vector is projected on the relative uncertainty matrix as follows:

$$J_\rho(\mathbf{u}) = \mathbf{M}_\rho^T \tilde{\mathbf{P}}_f \mathbf{M}_\rho \quad (4.14)$$

$$\mathbf{M}_\rho = \begin{bmatrix} \mathbf{0}_{n \times 1}^T & \hat{\boldsymbol{\rho}}_{\tau_f}^T & \mathbf{0}_{(n-3) \times 1}^T \end{bmatrix}^T \quad (4.15)$$

where $\boldsymbol{\rho} = \mathbf{r}_t - \mathbf{r}_a$ is the relative range vector from the agent to the target. It is important to highlight that this cost function only focuses on the relative position and does not minimize the entire orthogonal sub-space. The position sub-space is primary chosen because it is directly measurable by the sensor and minimizing positional uncertainty should carry over into velocity space over several measurements through the correlated dynamics. Second, removing the relative velocity component from the cost also removes tuning and scaling issues that arise between position and velocity magnitudes. Finally, it is assumed that the absolute state estimate has already been established through prior SSAT or a given prior estimate.


Equation 4.14 is the applied cost function we wish to optimize for optical-only sensing. The functional J_ρ is a nonlinear problem which requires complex solvers that are not suitable for autonomous guidance to truly optimize. Therefore, the following guidance methods will define surrogate models and cost functions that are analytically solvable, suitable for online guidance, and nearly approximate the true cost J_ρ . The first step towards developing these guidance policies is to examine the impact of control on relative geometry and how it relates to information gathering.

4.2.4 Angular Geometry and Variation

Information gathering is predicated on the basis that maneuvers cause changes in viewing geometry, which in turn allows the observer to measure new portions of the state space. Thus,

having a geometric understanding of how maneuvers relate to optical measurement deviations to generate range information is critical. Figure 4.1 illustrates how a maneuver δv causes changes in optical viewing geometry by the angle θ . In this figure the quantity ρ^* is the nominal range without a maneuver, δr is the change in agent position, and γ is the agent displacement angle. The value θ is the measured quantity with uncertainty σ_θ , while γ and δr are a function of the control due to δv , and ρ and ρ^* are the unknowns in this scenario.

Figure 4.1: Depiction of a change in viewing geometry due to a maneuver.



Given the geometry from Figure 4.1, the law of sines relates the unknown ρ to the measured quantity θ :

$$\rho = \frac{\sin(\gamma)}{\sin(\theta)} \delta r \quad (4.16)$$

Taking the first variation, denoted Δ , of ρ produces a relationship with its dependent quantities $\theta, \delta r, \gamma$:

$$\Delta \rho = \frac{\sin(\gamma)}{\sin(\theta)} \Delta \delta r + \frac{\cos(\gamma)}{\sin(\theta)} \delta r \Delta \gamma - \frac{\sin(\gamma) \cos(\theta)}{\sin^2(\theta)} \delta r \Delta \theta \quad (4.17)$$

Now, if perfect control is assumed, then the quantities $\gamma, \delta r$ are known and fixed values such that $\Delta \gamma = \Delta \delta r = 0$. This is equivalent to only considering a variation in the angle measurement θ and is akin to assumptions in previous works [12]. Substituting these assumption into Eq. 4.17, along

with the relationship for δr from Eq. 4.16, yields:

$$\begin{aligned}\Delta\rho &= -\frac{\sin(\gamma)\cos(\theta)}{\sin^2(\theta)}\delta r \Delta\theta \\ &= -\frac{\rho}{\tan(\theta)}\Delta\theta\end{aligned}\tag{4.18}$$

Equation 4.18 provides us with a direct relationship between a variation in the measurement space θ and the unknown ρ which we use to calculate statistical variances. The range standard deviation σ_ρ is the square root of the variance of the range variation $\Delta\rho$. The equation for variance is $\text{Var}(x) = E[x]^2 + E[x^2]$. Thus, the range variance is:

$$\sigma_\rho^2 = E[\Delta\rho^2] - E[\Delta\rho]^2\tag{4.19}$$

Next we also assume zero-mean variations $E[\Delta\rho] = E[\Delta\theta] = 0$, which leads to:

$$\begin{aligned}\sigma_\rho &= \sqrt{E\left(\left(-\frac{\rho\Delta\theta}{\tan(\theta)}\right)^2\right)} \\ &= \left|\frac{\rho}{\tan(\theta)}\right|\sqrt{E[\Delta\theta^2]}\end{aligned}\tag{4.20}$$

Noting that $E[\Delta\theta^2]$ is simply the measurement variance σ_θ^2 , again due to the zero-mean Gaussian measurement assumption, then the range uncertainty as a function of measurement uncertainty σ_θ is:

$$\sigma_\rho = \left|\frac{\rho}{\tan(\theta)}\right|\sigma_\theta\tag{4.21}$$

This expression is analogous to results from Reference [12] and has the same minima at $\theta = \pm\frac{\pi}{2}$. The primary difference is the replacement of the sine function with the tangent function, which is negligible for small angles. A secondary difference is that the newly derived expression has an absolute value, ensuring the range uncertainty is positive definite. With the relevant geometry established, the equation states that moving orthogonal to the original nominal trajectory, and minimizing range, will minimize the range estimate uncertainty. Seeking orthogonal viewing and minimum range aligns with intuition as both of these effectively work to triangulate the target's position with minimal uncertainty given optical measurements.

4.3 Information Gathering Guidance

Two guidance policies which use maneuvers to improve the current navigation state estimate are presented here. First, a heuristic policy for optical sensors is defined based on a simplified surrogate model that is numerically optimized. The second policy uses insights gained from the heuristic to define a surrogate cost function which is solved analytically. Both policies dictate the pointing of a burn to obtain near-maximal information and allow the burn magnitude to be a free parameter chosen by the user. Additionally, while both policies are designed for perfect maneuvers an analytic analysis of maneuver noise is included for the heuristic.

In practice maneuvers have execution errors which will inflate the estimated velocity uncertainty and cause sub-optimal information gain. Fortunately, maneuver execution error models, such as the Gates model [62], primarily align the error parallel to maneuver direction and only introduce small pointing offsets. Thus, even erroneous impulses should mostly align with the guidance specified pointing and lead to information gain along the desired projection. If the thruster noise is so great that it outweighs the information gained from the new geometry, then accelerometers should be considered to measure the erroneous maneuvers. Previous works have shown that precise accelerometers can reduce maneuver uncertainty and increase supplemental information gain [59].

4.3.1 Heuristic Guidance for Optical Sensors

The geometric analysis based on Figure 4.1 resulted in Eq. 4.21, which is an expression for range uncertainty that is minimized by minimizing range and moving orthogonal to the nominal viewing angle. The previous geometry and uncertainty equation can now be used to define a simplified surrogate model which is easily optimized. The problem is now to minimize range uncertainty given an optical sensor and a perfect maneuver such that the cost function is Eq. 4.21 and is written as:


$$J_h = \sigma_\rho = \left| \frac{\rho}{\tan(\theta)} \right| \sigma_\theta \quad (4.22)$$

When first approaching the new cost function, the problem is non-dimensionalized by the

nominal range vector $\boldsymbol{\rho}^*$. Non-dimensionalizing the problem reduces the free parameters and generalizes the analysis, similar to non-dimensionalizing dynamics in the CR3BP. Then a planar coordinate system is defined by the vector set $\{\hat{\boldsymbol{\rho}}^*, \hat{\boldsymbol{\rho}}_\perp^*\}$, where $\boldsymbol{\rho}_\perp^*$ is perpendicular to $\boldsymbol{\rho}^*$. A planar system is used instead of a three dimensional Cartesian space because a third axis only introduces symmetry with respect to the triangular geometry.

In the non-dimensional system $\boldsymbol{\rho}^*$ becomes a fixed unit vector and the geometry is fully defined by the two parameters $\frac{\delta r}{\rho^*}$ and γ via the side-angle-side theorem. Therefore, we can sweep over $\frac{\delta r}{\rho^*}$ as the independent variable and use the displacement angle γ as the free variable to minimize the cost J_h . The optimization is carried out with a nonlinear solver and the optimal pointing is reported in Figure 4.2. Translating the optimization within the figure to spacecraft operations requires

Figure 4.2: Optimal displacement angle γ to minimize range uncertainty σ_ρ given the geometry from Figure 4.1.



implementing the change in position vector $\delta \mathbf{r}$ through a maneuver. Making the assumption that executed maneuvers are small with short coast times, the maneuver displacement can be linearly approximated as $\delta \mathbf{r} = \mathbf{u} \delta \tau$.

The linear maneuver approximation allows the displacement magnitude δr to become a proxy for maneuver magnitude and coast time, while γ becomes the pointing of the burn. Additionally, because we are only considering small burns we assert that the nominal range is much larger than

the perturbation caused by the maneuver such that $\delta r \ll \rho^*$, which places the optimal maneuver on the left hand side of Figure 4.1. Therefore, from the small linear maneuver assumption precipitates the heuristic that the optimal maneuver to minimize range uncertainty σ_ρ is pointed perpendicular to $\boldsymbol{\rho}^*$ such that $\gamma \approx \pi/2$. Also, in order to minimize σ_ρ then the angle θ must be maximized, indicating that the burn magnitude should be set to the constraint v . Therefore, the heuristic policy is:

$$\mathbf{u}_h = v\hat{\mathbf{u}}_h \quad (4.23)$$

where the control pointing is defined by:

$$\hat{\mathbf{u}}_h \in \ker(\hat{\boldsymbol{\rho}}^*) \quad (4.24)$$

As a point of validation, this heuristic encompasses previously solved optimal policies for proximity operations with optics in the Clohessy-Wiltshire dynamical model [12, 60]. This heuristic also provides additional valuable insight into information gathering under small maneuvers. The heuristic assumes the maneuver is small so that the resulting displacement is also small, and therefore minimizing range is less valuable than maximizing a change in viewing geometry. Maximizing a change in viewing geometry is equivalent to maximizing a change in the measurement space in this case, which is intuitively appealing because it equates to allowing the agent to obtain as much new information as possible.

Now, the impact of maneuver execution errors on the heuristic are analytically considered by removing the perfect maneuver assumption and allowing the maneuver magnitude to vary inside of Eq. 4.17. The first step is to use the small displacement $\delta r \ll \rho^*$ to assert that $\theta \ll 1$ and that the small angle approximation for θ is valid. The small angle approximation together with $\gamma = \frac{\pi}{2}$ eliminates the $\Delta\theta$ term, as well as sine and cosine functions. Thus, Eq. 4.17 reduces to:

$$\Delta\rho = \frac{1}{\theta}\Delta\delta r - \frac{\delta r}{\theta^2}\Delta\theta$$

Then, inserting the linear maneuver equation produces:

$$\Delta\rho = \frac{\delta\tau}{\theta}\Delta u - \frac{u\delta\tau}{\theta^2}\Delta\theta$$

Next, the small angle approximation is inserted into the linear maneuver equation to obtain $\rho \theta = u \delta \tau$ based on the right triangle geometry which is enforced by the heuristic. Now control can be substituted into the variation equation to get:

$$\Delta \rho = \frac{\rho}{u} \Delta u - \frac{\rho^2}{u \delta \tau} \Delta \theta$$

Finally, the range variance under the heuristic policy, denoted σ_{ρ_h} , is the given by the variance of its variation again:

$$\sigma_{\rho_h} = \sqrt{E \left[\left(\frac{\rho}{u} \Delta u - \frac{\rho^2}{u \delta \tau} \Delta \theta \right)^2 \right]}$$

but given the expectations of Δu and $\Delta \theta$ are uncorrelated such that $E[\Delta u \Delta \theta] = E[\Delta u]E[\Delta \theta]$, and the zero-mean $E[\Delta \theta]$, we arrive at:

$$\sigma_{\rho_h} = \frac{\rho^2}{u \delta \tau} \sqrt{\sigma_{\theta}^2 + \left(\frac{\delta \tau}{\rho} \right)^2 \sigma_u^2} \quad (4.25)$$

This final equation states the range uncertainty is proportional to the root-mean-square of the measurement noise plus the maneuver magnitude error scaled by burn time over range. Thus, if the maneuver errors are small then the measurement noise will dominate. On the other hand, thrusters with larger errors cannot neglect this term and should consider accelerometers to reduce maneuver error and uncertainty. Additionally, the whole term is inversely proportional to maneuver magnitude and coast time. This coincides with our understanding of the problem because larger maneuvers that are allowed to coast longer create greater changes in geometry which lead to more information.

Overall, the heuristic policy provides three primary benefits. First, it is extremely easy to calculate which makes it useful for seeding optimal solvers. Second, the true optimal solution should approach the heuristic under the stated assumptions which is a practical sanity check. Third, the assumptions allowed us to more easily incorporate maneuver execution errors to show that the final range uncertainty is a root-mean-sum of maneuver errors and measurement noise. The major drawback is that it defines an admissible control region and not an explicit control. Additionally, the linear maneuver assumption is only good for small maneuvers and short coast times.

4.3.2 Maximum Measurement Deviation Guidance

The heuristic policy revealed that maximizing a change in the measurement space is an effective method for obtaining state information given a single sensor system. Therefore, a new surrogate cost function is adopted to maximize a measurement deviation defined as $\delta \mathbf{y} = \mathbf{y} - \mathbf{y}^*$ along with a Lagrange multiplier λ to enforce the control magnitude v . The new surrogate cost is:

$$J_y = \delta \mathbf{y}^T \delta \mathbf{y} - \lambda(\mathbf{u}^T \mathbf{u} - v^2) \quad (4.26)$$

We have solved the optimal control policy that satisfies the cost J_y analytically given $\delta \mathbf{y}$ is approximated with the linear term of its Taylor series expansion.

Start by defining $\delta \mathbf{y}$ by its Taylor series:

$$\begin{aligned} \delta \mathbf{y} &= \mathbf{y} - \mathbf{y}^* \\ &= \frac{\partial \mathbf{y}}{\partial \mathbf{u}} \mathbf{u} + \text{H.O.T.}(\mathbf{u}^2) \end{aligned} \quad (4.27)$$

Given that \mathbf{u} is small, the higher order terms are dropped and the term $\frac{\partial \mathbf{y}}{\partial \mathbf{u}}$ is expanded via the product rule:

$$\begin{aligned} \delta \mathbf{y} &= \frac{\partial \mathbf{y}_{\tau_f}}{\partial \mathbf{x}_f} \frac{\partial \mathbf{x}_f}{\partial \mathbf{x}_0} \frac{\partial \mathbf{x}_0}{\partial \mathbf{u}} \mathbf{u} \\ &= \mathbf{H}_f \Phi(\tau_f, \tau_0) \mathbf{B}_a \mathbf{u} \end{aligned} \quad (4.28)$$

Note that \mathbf{H}_f is the measurement sensitivity matrix evaluated at τ_f , $\Phi(\tau_f, \tau_0)$ is the state transition matrix from τ_0 to τ_f , and that the expression $\mathbf{H}_f \Phi(\tau_f, \tau_0) \mathbf{B}_a$ is the linear mapping from control space to measurement space. Now the expression for $\delta \mathbf{y}$ from Eq. 4.28 can be replaced in Eq. 4.26 to get:

$$J_y = \mathbf{u}^T \mathbf{Q} \mathbf{u} - \lambda(\mathbf{u}^T \mathbf{u} - v^2) \quad (4.29)$$

where \mathbf{Q} is the symmetric matrix:

$$\mathbf{Q} = \mathbf{B}_a^T \Phi(\tau_f, \tau_0)^T \mathbf{H}_f^T \mathbf{H}_f \Phi(\tau_f, \tau_0) \mathbf{B}_a \quad (4.30)$$

The necessary conditions for optimality require that $\frac{\partial J_y}{\partial \mathbf{u}} = 0$. Writing out necessary condition

and solving for \mathbf{u} results in:

$$\begin{aligned} 0 &= \mathbf{u}^T \mathbf{Q} - \lambda \mathbf{u}^T \\ \mathbf{Q}\mathbf{u} &= \lambda \mathbf{u} \end{aligned} \tag{4.31}$$

Equation 4.31 is an eigenproblem on the symmetric matrix \mathbf{Q} , so $\mathbf{u} = \pm \text{eig}(\mathbf{Q})$ satisfies the necessary condition but does not define the optimal eigenpair. To identify the optimal eigenpair, Eq. 4.31 is substituted into Eq. 4.29 which leads to:

$$\begin{aligned} J_y &= \lambda \mathbf{u}^T \mathbf{u} \\ &= \lambda v^2 \end{aligned} \tag{4.32}$$

Because v is fixed, J_y is maximized by choosing the largest eigenvalue λ . Therefore, the solution is:

$$\mathbf{u}_y = \pm v \max(\text{eig}(\mathbf{Q})) \tag{4.33}$$

given that the eigenfunction returns a unit vector. The guidance control \mathbf{u}_y , which we call the maximum measurement deviation (MMD) guidance, is a purely analytic solution for a generic system which should maximize information to the first order. Unfortunately, the sign is not given so both solutions must be examined for the optimal sign. Still, checking between two analytic options is significantly easier to compute than implementing a numerical optimizer.

An interesting aside is to note that the eigenvalues of the matrix \mathbf{Q} are equal to the square of the singular values of the term $\mathbf{H}_f \Phi(\tau_f, \tau_0) \mathbf{B}_a$, which would be another way to solve the problem. The singular value approach is equivalent to selecting the most stretching direction of the linear transformation from control to measurement space.

4.3.3 Combined Information Gathering and Station Keeping

Fuel expenditure is one of the most important aspects of a spaceflight mission and must be conserved carefully. While information gathering may be critical to ensure safe and accurate navigation, spending fuel to purely improve filter certainty may not always be justifiable. Fortunately, information gathering can be combined with station keeping if the station keeping policy has a

non-singular set of candidate solutions. The operations are combined by selecting the maneuver in the station keeping set that most closely aligns with the information gathering guidance.

The first step to selecting the station keeping maneuver, denoted \mathbf{u}_{sk} , that most closely aligns with the MMD guidance policy $\pm\mathbf{u}_y$ is to calculate the projection:

$$\varphi_u = \frac{\mathbf{u}_y^T \mathbf{u}_{sk}}{\|\mathbf{u}_y^T\| \|\mathbf{u}_{sk}\|} \quad (4.34)$$

Because the MMD guidance policy has a \pm ambiguity, the projections closest to $\varphi_u = 1$ and $\varphi_u = -1$ must both be considered as candidate solutions to maximize information inside the station keeping set. Thus, the two potential options for the informed station keeping maneuver is:

$$\begin{aligned} \mathbf{u}_+^* &= \max_{\mathbf{u}_{sk}} \varphi_u \\ \mathbf{u}_-^* &= \min_{\mathbf{u}_{sk}} \varphi_u \end{aligned} \quad (4.35)$$

To choose between \mathbf{u}_+^* and \mathbf{u}_-^* a linear covariance simulation is executed for both solutions, and the maneuver which produces a lower uncertainty is chosen as the informed station keeping maneuver denoted \mathbf{u}^* .

It is important to note that this guidance technique is time dependent because \mathbf{u}_y has a time dependent STM in it. Therefore, if time is free in the station keeping policy then the informed station keeping maneuver \mathbf{u}^* must be computed over the acceptable times. Additionally, In order for a nontrivial solution to exist there must be a non-singular set of station keeping maneuvers \mathbf{u}_{sk} .

4.4 Simulation and Results

The developed guidance strategies are tested in two scenarios mimicking proximity operations about the Lunar Gateway. The first scenario is a pure information gathering operation aimed at improving the relative range state estimate. The second scenario analyzes the proposed combined information gathering and station keeping guidance to determine its effectiveness.

In both scenarios the CR3BP dynamics are used to approximate cislunar motion, with the equations of motion given in Section 2.1.2. In the CR3BP model the target spacecraft is placed

on a near rectilinear halo orbit (NRHO) that has a 9:2 resonance with the lunar synodic cycle, which is the the Lunar Gateway’s currently selected orbit. The agent spacecraft is located on a quasi-periodic orbit (QPO) about the 9:2 NRHO with a stroboscopic period equal to the NRHO’s orbit period. Giving the spacecraft the same period ensures naturally bounded relative motion, which is ideal for observation and staging before further action. The average range between the natural trajectories of the two spacecraft for the missions outlined below is 215 km.

Both the agent and target spacecraft are estimated in a SSAT state space model and given an initial $3\text{-}\sigma$ uncertainty of 60 km and 60 cm/s per axis. The uncertainty is propagated and updated via a linear covariance analysis where the nominal trajectory is the true trajectory. The covariance analysis produces an upper bound on filter performance given that the implemented filter does not diverge from the true trajectory. The covariance analysis is deterministic and repeatable, which allows for quick evaluation and direct comparisons between simulations.

In all simulations the agent has access to optical measurements that are taken once an hour, with a root-mean-square uncertainty model as described by the measurement modeling in Section 4.2.2. Because only optical measurements are considered, the objective is to improve the nonlinear relative range estimate as quantified by the cost J_ρ from Eq. 4.14. Note that the square root of the cost function $\sqrt{J_\rho}$ is the $1\text{-}\sigma$ relative range uncertainty.

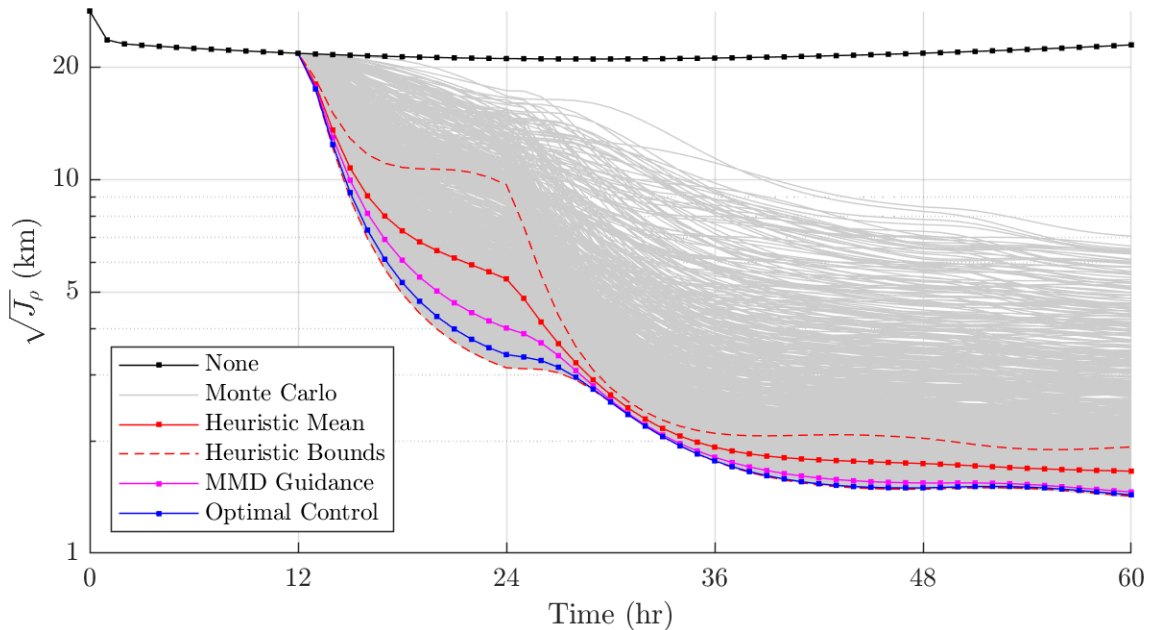
4.4.1 Information Gathering Operation

The first simulated scenario is a pure information gathering operation intended to compare the effectiveness of the developed strategies. The operation begins 24 hours before the target’s apoapsis. Then the agent’s initial guidance maneuver is performed 12 hours into the mission, with a maximum fuel expenditure of 10 mm/s. After, both vehicles coast for 12 hours until the target reaches apoapsis. This coasting period allows sufficient time for geometric changes to build up and be observed by the agent. Then, at apoapsis, the agent performs a second maneuver to return to the nominal trajectory 24 hours after the targets apoapsis. Finally, once the position has been restored, the agent executes a third burn to match the nominal velocity. A final coasting period

of 12 hours is included to allow for a steady state to be reached. Optical measurements are taken once an hour for the entirety of the scenario.

The information gathering operations are simulated with the heuristic guidance, MMD guidance, a true optimal solver, and a Monte Carlo analysis randomly drawing 10 mm/s maneuvers. The evolution of range uncertainty $\sqrt{J_\rho}$ for all simulations is displayed in Figure 4.3. When simulating the heuristic guidance it is important to identify there is an admissible control set. In order to assess the policy as a whole, uncertainty quantification of the full set was performed. The resulting data was aggregated into an upper/lower bound, as dashed lines, as well as a mean value which is a solid line. The Monte Carlo maneuver analysis was carried out to identify global performance of applying a maneuver and appear as the light gray lines in the figure. Finally, a baseline simulation with no maneuver was also included to showcase the importance of applying a maneuver to gaining information.

Figure 4.3: Range uncertainty evolution over information gathering operation. Note, the y-axis is in log-scale.



The range uncertainty profiles from Figure 4.3 contain several important qualities. The

optimal policy in blue, found by a nonlinear numeric solver and agreed upon by the Monte Carlo analysis, scores a final range uncertainty of 1.42 km which is over an order of magnitude more certain than the baseline which finishes with an uncertainty of 22.95 km. The heuristic policy in red has a lower bound which captures the optimal trajectory, but it also has an upper bound of 1.93 km and an average value of 1.66 km. Finally, the analytic maximum measurement deviation policy improves on the average heuristic and is only marginally off the optimal trajectory with a final cost of 1.46 km.


There are several notes of interest regarding these simulations. First, all the trajectories evolve nonlinearly and there is no one trajectory that fits the bottom bound for all time. This means the optimal policy which minimizes the uncertainty at the final time does not minimize uncertainty for the whole trajectory. The time dependence on optimality is expected because the trajectories and associated uncertainties evolve nonlinearly with time. The fact that the heuristic captures the bottom bound for all time indicates its admissible set contains all the optimal maneuvers for this time period, which truly endorses the heuristic policy and indicates its development was accurate.

The second note regarding these simulations is that the first measurement after instantiation obtains range information, but subsequent measurements provide little to no improvement without a maneuver. This lack of range information is expected because it takes considerable time for the geometry to naturally deviate and for the information to be collected. Similarly, the first measurement after a maneuver produces the most information, but subsequent measurements taper off to a steady state with time. The steady state seeking nature of uncertainty after a maneuver demonstrates that coasting is necessary to accrue information, but there is a fundamental limit to the amount of information that can be gained.

All of these simulations had the initial burn constrained to 10 mm/s, but the two return maneuvers were not constrained since they were deterministic maneuvers. To assess total fuel usage, the cost of all 3 maneuvers were summed and are presented in Figure 4.4. As the figure shows, the total cost of all 3 maneuvers never exceeds 30.5 mm/s, which is approximately three times the guidance maneuver cost. Therefore, this information gathering operation expends about

30 mm/s to improve the range uncertainty by an order of magnitude. It is very important to note that these results are dependent on the simulation design. Relative orbits with greater distances will require more fuel to achieve the same level of improvement, while closer orbits will require less.

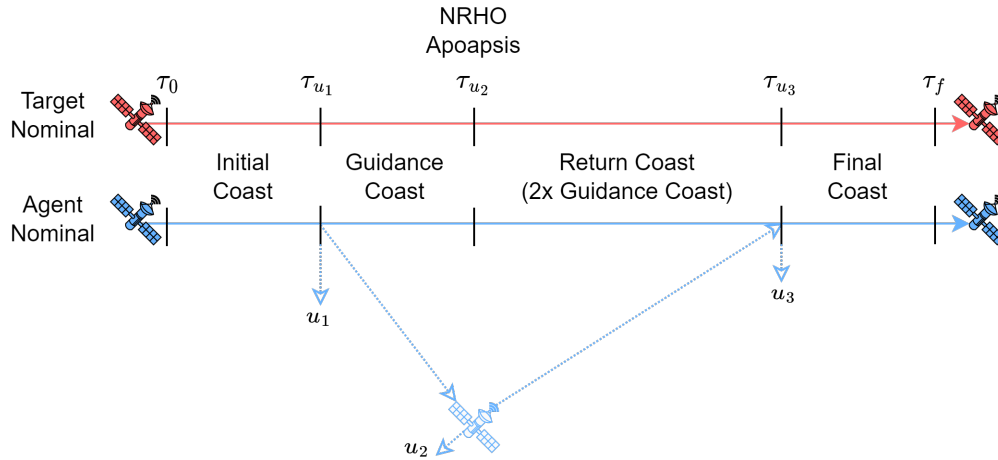
Figure 4.4: Total δv fuel usage for guided simulations.



Fuel constraints and coasting times are the primary factors in the performance of these simulations. Thus, these parameters are varied to determine their impact on final uncertainty reduction. The following simulations have the same structure as before: one information gathering burn followed by two more burns to return to the nominal. The guidance burn is calculated using the maximum measurement deviation policy \mathbf{u}_y and the fuel constraint v is varied as a free parameter. The subsequent return burns are deterministically calculated at apoapsis to return the spacecraft to its nominal trajectory. The initial coasting period is reduced to 3 hours based on the results from the previous simulation. The coast between the guidance burn and the first return burn is called the guidance coast period and is a free parameter that is varied. The period between

the first and second return burns is always twice as long as the guidance coast period because it was found that this limited the total fuel spend to approximately 3 times v . Then, there is a final coast after the final burn which is 12 hours. This outline is depicted in Figure 4.5.


Figure 4.5: Timeline for information gathering operations.



The simulations are run by varying guidance maneuver magnitude and coasting time between burns and presented in Figure 4.6. The y-axis is the resulting final range uncertainty reduction, which is the difference between the non-maneuvering uncertainty and the final uncertainty obtained with a guidance burn. The x-axis is the total fuel spent over of all 3 burns together. The figure depicts a clear Pareto front between total fuel usage and uncertainty reduction for a fixed coasting time as represented by a single line. The color bar encodes the coasting time between the guidance maneuver and the first return maneuver, and as time increases then uncertainty is also reduced.

The data in Figure 4.6 has several important qualities. First, the fuel-uncertainty reduction relationship is Pareto optimal and must be balanced in application. Second, increasing the coast time is an effective method to increase information gain for no fuel cost, but using infinite time is not desirable either. Finally, even small maneuvers on the order of 1 mm/s still produce tangible improvements given the simulated geometry.

Figure 4.6: Pareto front of total δv vs the reduction in range uncertainty $\sqrt{J_\rho}$. The guidance coast is encoded by the color bar.



4.4.2 Informed Station Keeping Operation


This section simulates the combined information gathering and station keeping guidance. To achieve this, a station keeping policy with a continuous set of maneuvers on a 2 dimensional surface is created. Then MMD guidance policy is projected onto the station keeping set to select the informed station keeping maneuver \mathbf{u}^* . Then all maneuvers in the set, including \mathbf{u}^* as a special case, are simulated with a linear covariance analysis to determine where the informed station keeping maneuver falls with regards to the set in reducing uncertainty.

As stated in the Section 4.3.3, there must be a non-singular set of station keeping maneuvers for a nontrivial combination of information gathering and station keeping to exist. Therefore, for the purposes of this work, the implemented station keeping policy is to place a spacecraft on the surface of a QPO in cislunar space rather than at a single point on an orbit. To target the surface of a QPO, we begin with an invariant ring as computed using the GMOS algorithm [21]. The ring sweeps out the surface of the QPO over the stroboscopic period which is matched with the NRHO orbit period. Because the ring and orbit have the same period, the ring maintains a closed geometry about a point on the NRHO orbit.

The station keeping maneuver places the agent spacecraft on the invariant ring as it evolves over time such that the agent retains its closed geometry. Thus, the two targeting parameters for the station keeping policy are the angle on the ring, which is normalized from 0 to 1, and the propagation time that the ring is moved forwards from the initial condition. The station keeping policy only targets a time about the NRHO apoapsis, so it only sweeps out a section of the QPO as depicted in Figure 4.7.

The maneuvers were computed using a single shooting algorithm given a fixed point initial condition, with the final states parameterized by the invariant ring propagated one orbit period plus or minus a small time variation. For each propagation time there is a unique informed station keeping maneuver \mathbf{u}^* because of the time dependence in \mathbf{u}_y . By allowing the propagation time as well as the point on the ring vary, a two dimensional maneuver surface is constructed. Note, since

Figure 4.7: Section of the QPO surface which is targeted by the station keeping maneuvers.



this is a shooting algorithm two burns must be executed, an initial position matching burn and secondary velocity matching burn.

The surface of all station keeping maneuvers, as well as the maneuver the informed station keeping maneuvers \mathbf{u}_{\pm}^* in magenta, are plotted in Figure 4.8. The left hand surface are the initial maneuvers, and the right hand plot are the secondary velocity matching maneuvers. Note that the informed station keeping maneuver is only predicted by the first maneuver.

All of these burns are simulated under the same initial conditions that were mentioned in the header of this section. The simulations start 3 hours before apoapsis where the initial burn is scheduled. There is also a final 12 hour coast after the second burn. The resulting final range uncertainties all station keeping maneuvers and the informed station keeping maneuvers are depicted in Figure 4.9. The surface represents the final range uncertainty $\sqrt{J_{\rho}}$, and the magenta line is the informed maneuver. The ring angle is the point on the invariant ring that is targeted by the station keeping maneuver, which is normalized from 0 to 1. The station keeping maneuver also propagates the invariant ring one orbit period plus the time from apsis to sweep out a section of the QPO near apoapsis.

The figure depicts that the informed station keeping maneuvers fall in the valley of minimum

Figure 4.8: Acceptable station keeping maneuvers to place the agent on a QPO surface.

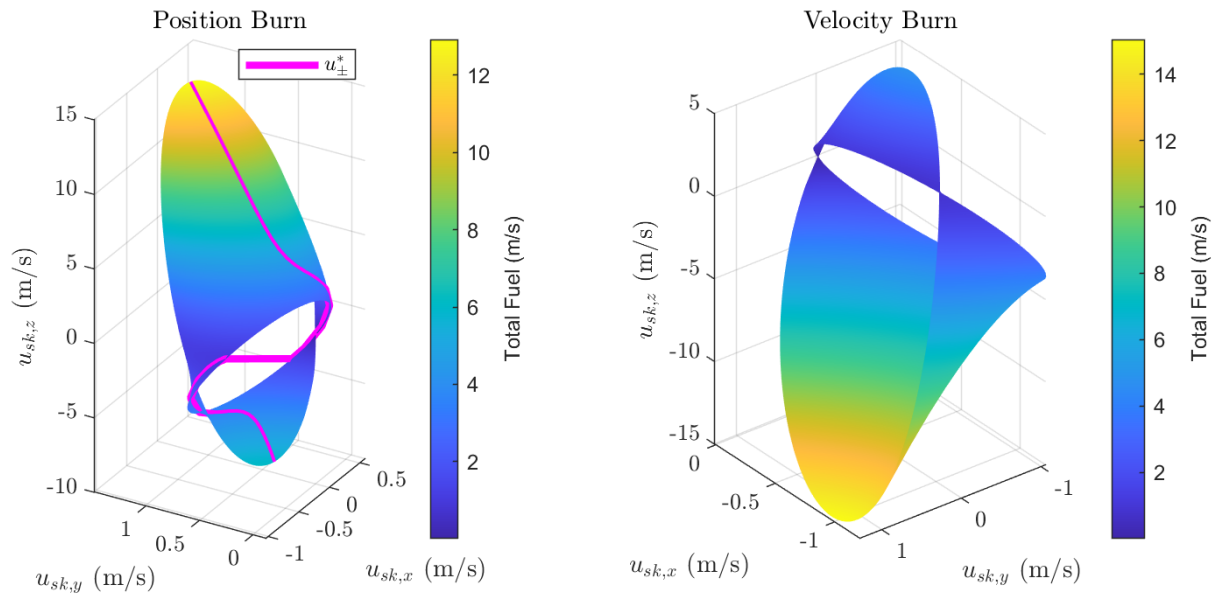




Figure 4.9: Resulting range uncertainties from station keeping maneuvers in Figure 4.8.



range uncertainty as desired, with a single point being raised by a small local peak. The small local peak near the center of the valley is induced by the secondary maneuver which is not accounted for in the prediction, but it only marginally impacts the overall performance. Therefore, if a surface of potential station keeping maneuvers is available, then the projection of the information gathering guidance policy \mathbf{u}_y can be used to quickly select the maneuver that improves the state estimate.

As a final analysis it is worth examining the fuel cost for these maneuvers to determine if the guidance policy selects the largest maneuvers. The total fuel costs for all simulations are displayed in Figure 4.10. Note the time from apoapsis axis has been flipped to get a better view of the predicted maneuvers. The results show that the chosen maneuver actually avoids the peak fuel cost in favor of alignment with the policy \mathbf{u}_y .

Figure 4.10: Resulting fuel usage from station keeping maneuvers in Figure 4.8.



4.5 Conclusion

While optical sensors are light weight, low power, and effective measurement sources, they inherently lack range information which may take substantial time to obtain. Instead of waiting for the natural dynamics to create the geometric changes necessary for range information to be estimated, one satellite may perform a maneuver to immediately change the relative geometry. By incorporating maneuvers to obtain new information an autonomous system will be able reduce uncertainty in its environment and actions. Thus, we have developed analytic guidance policies which enhance a portion of the state estimate that is orthogonal to the measurement space because that space is generally the least certain.

The first step towards developing the guidance policies was an analysis on optical information gathering to improve the relative range estimate. The optical analysis was used to define a surrogate model which resulted in a heuristic policy. The heuristic policy motivated a more general approach which led to the maximum measurement deviation surrogate cost function. The measurement deviation cost function was solved analytically which produced a light weight guidance policy that is suitable for onboard use and is a first order optimization of information orthogonal to the measurement space.

The guidance policies were simulated in the cislunar environment with optical-only measurements to test their performance. At first, a pure information gather campaign was simulated where all maneuvers were limited to a maximum maneuver magnitude of 10 mm/s. The heuristic guidance policy successfully captured the true optimal policy for the duration of the simulation time and maximum measurement deviation policy was nearly optimal for a fraction of the computation cost. Every simulation had a total fuel cost that was approximately 3 times the initial guidance burn. Then, a trade between fuel cost and uncertainty reduction was simulated and formed a Pareto front. Finally, it was found that increasing the coasting time between maneuvers can circumvent the need for large maneuvers if the time is available.

Finally, the maximum measurement deviation policy was combined with a station keeping

policy by projecting the MMD policy onto a surface of potential station keeping maneuvers. It was found that the station keeping policy that is most closely aligned with the MMD policy is a fair approximation of the optimal maneuver to reduce state uncertainty in the set. It must be noted that while the station keeping maneuvers themselves are not fuel optimal, they can be replaced with a more complex and lower cost set if found.

There are several avenues for future work. First, maneuver execution errors should be included in simulation and measured with accelerometers to determine the level of accuracy necessary within the maneuver/sensor that ensures noise does not overtake the information gain. Second, alternative station keeping policies that have lower fuel costs should be examined. Finally, these guidance maneuvers should be simulated in a full filter simulation to assess filter consistency.

Chapter 5

Spacecraft to Spacecraft Absolute Tracking and Maneuver Classification in Cislunar Space

5.1 Introduction

Cislunar space has been receiving significant attention to the region's suitability as a proving ground and waypoint for deep space exploration. The multinational Artemis program is one of many programs for cislunar development which will likely give rise to an influx of spacecraft cohabiting the region [63]. The myriad of vehicles traversing cislunar space will likely require enhanced observation and navigation capacities over traditional Earth-based approaches to ensure the collective safety of all vehicles in the region.

Traditional cislunar estimation relies on terrestrial resources which are subject to various observational challenges. Earth-based radar, such as the Deep Space Network, are already over-subscribed which means that equitably scheduling observation time for all future vehicles will be prohibitive. Earth-based optics must avoid exclusion zones around the Moon because excessive light can damage sensors. All Earth-based observers have a limited viewing geometry of cislunar objects, which results in sparse information and potentially long periods of occultation. The extensive list of considerations regarding Earth-based observing of cislunar objects warrants the investigation of space-based observation to alleviate Earth-based sensors and hopefully eliminate various tracking issues [64].

Fortunately, recent studies have shown that space-based observers show significant promise to provide a viable solution for cislunar navigation and observation [29, 65, 66, 44]. In particular, a

space-based observer using SSAT presents the opportunity to provide autonomous navigation and simultaneous tracking of nearby objects [67]. The benefits of SSAT result from the fact that it does not assume a given observer state which would require external orbit determination. Instead, both the observer and target are simultaneously estimated in SSAT. Cislunar space provides an ideal proving ground for optical-only SSAT due to the success of range-only measurements as proven by Liaison [43]. Still, even though it is theoretically possible, it must be demonstrated that optical-only SSAT can maintain a consistent filter solution when challenged by strongly nonlinear dynamics and unmodeled events such as maneuvers.

The premise of this chapter is to examine the feasibility of optical-only SSAT and maneuver classification within a filter as a robust tracking and navigation method for cislunar operations. To achieve this premise requires two major objectives: demonstration of successful optical-only SSAT in cislunar space, and accurate maneuver classification to ensure tracking under imperfect modeling. The chapter continues with Section 5.2 which provides additional background and supporting literature on these two topics. Section 5.3 describes the setting for the analysis and simulations. Section 5.4 demonstrates optical-only SSAT filter results in the given problem setting. Finally, Section 5.5 obtains control estimates from the optimal control based estimator to train and test maneuver classification algorithms. In total, this work presents a method of state estimation that successfully provides maneuver detection and classification with optical-only measurements that is completely independent of ground-based resources.

5.2 Background Literature

5.2.1 Space-Based Optical Tracking in the Cislunar Environment

Space-based observation of other space objects has a significant collection of background literature and applications. While extensive, we focus on SSAT, which is the use of relative measurements between space objects to generate an absolute state estimate for a distributed space system. The background for SSAT was covered in chapter 3. In addition to estimating the abso-

lute state from relative measurements, the full relative state must also be observable which is not guaranteed for single sensor systems. Specifically, obtaining range information from optical-only measurements under linear dynamics is unobservable. Fortunately, it has been shown in previous work, and in chapter 3, that relative range is locally observable given optical measurements for appropriately designed DSS configurations [37, 28, 68, 38, 69, 39]. As a result, it is possible for optical-only measurements to produce complete observability for both the absolute and relative states of a DSS.

However, even if a system is theoretically observable, it may not be easily filtered in the face of strong nonlinearities. Nonlinear dynamics pose a considerable concern for cislunar tracking because this environment is known to have reduced regions of linear operating validity that impact filter performance [17, 15, 18]. Regions with reduced linear validity are especially concerning for SSAT because estimating the states for two spacecraft can lead to large uncertainty volumes which further exacerbate performance issues for linearized filters. Even when only estimating the state of a single spacecraft, the nonlinear dynamics of cislunar space are known to be strong enough to cause filter divergence [15, 18, 69].

There are many higher order filtering methods, and additional exposition on this topic can be found in References [70, 71]. Potential options include higher order Kalman filters that use tensors to calculate nonlinearities, Gaussian mixture models to refine probability estimates, and particle filters which do both. While all these methods improve on the basic linear Kalman filter they come at the cost of additional computational load which is a limited resource for autonomous applications. The unscented Kalman filter (UKF) is effectively an instance of the particle filter which intelligently selects points to approximate probability distributions for minimal computation cost [72, 20]. The balance of computation cost and fidelity in the UKF makes it a logical step forward over linear filters if a linear filter is unable to meet demands. It must be mentioned that the selection of the UKF does not insinuate that it is optimal, only that it may be sufficient when other filters falter.

5.2.2 Maneuver Detection and Classification

Nonlinear cislunar dynamics necessitate station keeping in order to maintain an orbit, which also needs to be accounted for in tracking solutions [15, 16]. If station keeping maneuvers are not identified, or worse, confused with orbit departure maneuvers, then it is highly likely that tracking solutions will diverge. Therefore, maneuver detection is a critical capability because it enables robust tracking solutions that would otherwise be lost if a maneuver is not communicated. Similar to nonlinear filtering, numerous methods are available for maneuver detection but they often come at the cost of additional computational requirements.

Maneuver detection has been addressed in the past for various systems and applications. Much of the foundational work dealt with missiles and had two approaches: utilizing multiple models and/or estimating accelerations in the state [73]. For spacecraft, a variety of approaches have been developed, but most focus on pure detection and do not provide maneuver reconstruction [74, 75, 76, 77, 78]. There are, however, two promising methods for maneuver reconstruction: multiple model filters and control distance metrics.

Multiple model filters function by weighing banks of filters with different dynamical models to provide the best state estimate and most likely dynamical model at any given instance [79]. The downside of multiple model approaches is that a significant number of filters must be computed simultaneously in order to cover all potential actions. On the other hand, control distance metrics have shown that they are capable of reconstructing mismodeling while incorporated into a single generalized Kalman filter in the form of the optimal control based estimator (OCBE) [80, 19]. The ballistic OCBE is remarkably attractive because it runs as a single linear filter but is still able to reconstruct unmodeled events without having to specify maneuver models. As with most navigation methods, the ballistic OCBE is a linearized filter but it also contains favorable properties that allows its maneuver reconstruction capabilities to be applied in other filter settings [69]. While some work has been done with the OCBE towards the goal of maneuver detection, the opportunity to assess the application of the OCBE for optical-only SSAT and maneuver classification in cislunar

space remains open [77, 81, 82]. Appendix A provides an abbreviated explanation of the ballistic linear OCBE along with its the equations for implementation.

Previous work was able to provide maneuver detection with the OCBE by creating probability distributions from estimated control profiles, thereby enabling a binary hypothesis test, which is covered in Appendix C. Thus, even though the OCBE control estimate can detect maneuvers, it is unknown if it can accurately reconstruct maneuvers for classification purposes because the optimality assumptions may misrepresent true events [83, 69]. Fortunately, while the event reconstruction may not be accurate, it is consistent. The consistency of control reconstruction makes it well suited for classification techniques which can label unknown data. It is fair to consider most modern approaches to the classification problem to fall under machine learning, and a well rounded overview of this approach can be found in References [84, 85]. This work will use classification algorithms to label spacecraft behavior based on the OCBE control estimate, providing a higher level and more insightful picture of spacecraft operations over pure detection. Further still, we will attempt maneuver classification while accounting for two unknown vehicles via SSAT instead of only estimating a single vehicle.

5.3 Problem Formulation

This section covers the models used in this chapter to simulate optical-only SSAT and maneuver detection in cislunar space. As in the previous chapter, the state space is a two vehicle SSAT state space with cislunar motion approximated via the CR3BP. Unlike the previous chapter, only the target is the maneuvering vehicle in this work. Target maneuvers are modeled as instantaneous impulses with two levels of execution errors and belonging to one of four classes: no maneuver, station keeping, orbit departure, and generic maneuver. Then several orbits of interest are selected from various orbit families within the CR3BP for simulation purposes. Finally, measurements are introduced based on the equations from Section 2.2 along with noise and constraint considerations.

5.3.1 Dynamical Models

The state space is a two vehicle SSAT model with a cooperative agent $\mathbf{s}_a \in \mathbb{R}^6$ and a non-cooperative target $\mathbf{s}_t \in \mathbb{R}^6$. The state of each vehicle is comprised of Cartesian position $\mathbf{r} \in \mathbb{R}^3$ and velocity $\mathbf{v} \in \mathbb{R}^3$ such that a single spacecraft state is $\mathbf{s} = [\mathbf{r}^T, \mathbf{v}^T]^T$. The natural dynamical model that dictates spacecraft motion is given by the function $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{s})$. The natural dynamics model used to approximate cislunar motion is the CR3BP as depicted in Figure 2.2 with equations of motion given by Eq. 2.9. The full state and dynamics are:

$$\mathbf{x} = \begin{bmatrix} \mathbf{s}_a^T & \mathbf{s}_t^T \end{bmatrix}^T \quad (5.1)$$

$$\mathbf{f}(\mathbf{s}) = \begin{bmatrix} \mathbf{v}^T & \mathbf{a}_{cr3bp}^T \end{bmatrix}^T \quad (5.2)$$

Now, in this chapter, the target is the maneuvering vehicle. The considerations of modeling maneuvers on the target are quite different from the agent because only the agent is cooperative. Because the agent is cooperative, its maneuvers are known and may be employed to gather information with minimal mismodeling the spacecraft has access to either accurate engines or accelerometers to measure the maneuvers [59]. On the other hand, if the target maneuvers, then the maneuver can only be reconstructed based on observations and are highly uncertain. For these reasons, only target maneuvers are considered in this chapter. The maneuvers are implemented as impulsive $\delta\mathbf{v}$'s given by the equations:

$$\mathbf{x}^+ = \mathbf{x}^- + \mathbf{B}_t \delta\mathbf{v} \quad (5.3)$$

$$\mathbf{B}_t = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{I}_3 \end{bmatrix}^T \quad (5.4)$$

When a maneuver is executed it is drawn from one of three considered maneuver classes: station keeping (SK), orbit departure (OD), and general maneuvers (GM). Station keeping is necessary for spacecraft to maintain their desired orbits, and optimal $\delta\mathbf{v}$'s on halo orbits are known to approximately lie along the stable manifold which is denoted $\bar{\mathbf{e}}_s$ [86, 87]. The cost of station keeping varies depending on navigational accuracy, but typical impulses are on the order of mm/s to m/s [15, 16]. Orbital departures are approximated as maneuvers along the unstable manifold,

denoted $\bar{\mathbf{e}}_u$, which lead to low-cost orbit transfers. Placing a spacecraft on an unstable manifold takes advantage of natural dynamics to move away from an orbit for no fuel. Finally, generalized maneuvers cover the control space and are used to identify maneuvers that may not be aligned with a manifold. Note that while there are three classes of maneuvers, the target may also choose to not maneuver which effectively creates four potential courses of action.

All target maneuvers are modeled to occur at lunar apoapsis because various studies have shown lunar apoapsis to be an ideal control point [15, 16]. These maneuvers are subject to a statistical execution error with probabilities depending on either a noisy or quiet error distribution [15]. The noisy maneuvers have a mean of 1 m/s with a 300 mm/s standard deviation. The quiet maneuvers have a mean of 40 mm/s with a 10 mm/s standard deviation. The error is calculated using the Gates model which are characterized as zero-mean Gaussians that are oriented in a frame which aligns the first axis parallel to the maneuver direction [62]. The Gates covariance, which is composed of both fixed and proportional components, is:

$$\mathbf{P}_g = \begin{bmatrix} \sigma_{g1}^2 + \|\delta\mathbf{v}\|\sigma_{g2}^2 & 0 & 0 \\ 0 & \sigma_{g3}^2 + \|\delta\mathbf{v}\|\sigma_{g4}^2 & 0 \\ 0 & 0 & \sigma_{g3}^2 + \|\delta\mathbf{v}\|\sigma_{g4}^2 \end{bmatrix} \quad (5.5)$$

The error is rotated from the $\delta\mathbf{v}$ frame (V) to the CR3BP dynamic rotating frame (D) via the cosine rotation matrix $[\mathbf{DV}]$. Using the Gates error model and the definitions of the maneuver classes above, the distributions for each maneuver class are:

$$\delta\mathbf{v}_{SK} = \pm\mathcal{N}(\mu_{\delta v}, \sigma_{\delta v})\bar{\mathbf{e}}_s + [\mathbf{DV}]\mathcal{N}(\mathbf{0}, \mathbf{P}_g) \quad (5.6)$$

$$\delta\mathbf{v}_{OD} = \pm\mathcal{N}(\mu_{\delta v}, \sigma_{\delta v})\bar{\mathbf{e}}_u + [\mathbf{DV}]\mathcal{N}(\mathbf{0}, \mathbf{P}_g) \quad (5.7)$$

$$\delta\mathbf{v}_{GM} = \mathcal{N}(\mu_{\delta v}, \sigma_{\delta v})\mathbf{Sph2Cart}(\mathcal{U}(0, 2\pi), \mathcal{U}(-\pi/2, \pi/2)) + [\mathbf{DV}]\mathcal{N}(\mathbf{0}, \mathbf{P}_g) \quad (5.8)$$

Equations (5.6)-(5.8) define the three maneuver classes for the target given all of the aforementioned assumptions. In these equations $\mathcal{N}(\mu, \sigma)$ is a Gaussian distribution and $\mathcal{U}(\min, \max)$ is a uniform distribution. Table 5.1 summarizes the statistics associated with the maneuver model. For this work, the noisy policy parameters are set as Cassini's pre-flight Gates values, and the quiet

parameters are a mixture of the 2000 and 2006 models [62]. When these errors are applied to the given policies, proportional errors prevail in the noisy case, while fixed errors dominate in the quiet case.

Table 5.1: Maneuver policy statistics

| δv Policy | statistics | | | | | |
|-------------------|-------------------------|----------------------------|----------------------|---------------|----------------------|----------------------|
| | $\mu_{\delta v}$ (mm/s) | $\sigma_{\delta v}$ (mm/s) | σ_{g1} (mm/s) | σ_{g2} | σ_{g3} (mm/s) | σ_{g4} (mrad) |
| Noisy | 1000 | 300 | 10 | 0.35 | 17.5 | 10 |
| Quiet | 40 | 10 | 6.5 | 0.2 | 4.5 | 3.5 |

5.3.2 Selected Orbits


Now orbits of interest are selected out of several periodic orbit families to explore the effects of viewing geometry on tracking solutions. Included families are northern L1 halos (NL1), planar L2 halos (PL2), southern L2 halos (SL2), southern L2 halo Quasi-Periodic Orbits (QPO), and distant retrograde orbits (DRO). All orbits and associated families were calculated using the continuation or GMOS algorithms as explained by Reference [88]. Exploring all orbits in every family is impossible, since they are continuous families, thus a single representative orbit from each family is selected for simulation purposes.

Figure 5.1 contains all of the selected orbits plotted about the moon. Table 5.2 contains the orbital period, radius of lunar periapsis which is the orbit's closest lunar approach, and stability index for each selected orbit. Note in Table 5.2 lunar periapsis is simply periapsis. It is also important to note that the SL2 orbit and the QPO have an approximately 9:2 resonance with the

Table 5.2: Orbit parameters for selected orbits

| Parameter | Orbits | | | | |
|----------------|--------|-------|------|------|-------|
| Family | NL1 | PL2 | SL2 | QPO | DRO |
| Period (days) | 11.9 | 14.6 | 6.55 | 6.55 | 14.9 |
| Periapsis (km) | 51111 | 62902 | 3212 | 2346 | 74309 |
| Stability | 1172 | 724 | 1.32 | 1.33 | 1.00 |

Figure 5.1: Plots of all the selected orbits examined in this chapter, displayed in the CR3BP rotating frame.



lunar synodic cycle, which is currently the target orbit for NASA’s Lunar Gateway. The selected 9:2 synodic SL2 orbit is also commonly called a near rectilinear halo orbit (NRHO) because of its long and highly eccentric shape, so SL2 and NRHO are used interchangeably. The stability index is calculated as $\frac{1}{2}(\lambda + \frac{1}{\lambda})$, where λ is the largest magnitude eigenvalue of the monodromy matrix [15].

5.3.3 Measurement Models

Similar to the previous chapter, all measurements are relative measurements which are taken by the agent and are directed at the target. The agent is the sole measurement source because it is the only controlled vehicle, which allows the target to operate independently without any cooperative planning necessary. General measurements are denoted by the vector $\mathbf{y} \in \mathbb{R}^m$ and are processed by a filtering algorithm to generate a state estimate. Both range-only and optical-only simulations are considered so that their performances can be compared, but optical-only sensing is the primary focus. The measurement equations for both sensors are found in Section 2.2.

The range sensor is used to recreate Liasion results, and therefore equivalent noise statistics are applied here [43]. Liasion used range measurements which were received approximately once every 6 minutes and had Gaussian uncertainty model with a 1 meter standard deviation. The optical sensor is simulated as a measurement vector containing azimuth and elevation elements which are statistically uncorrelated. Optical measurements are taken every 2 hours and have a 0 mean noise with a standard deviation given by a root-mean-square uncertainty model to account for partially resolved object uncertainty. The equations for the root-mean-square optical uncertainty are given in Eq. 4.10. Again, the value of $\sigma_{\theta, min}$ is set to 10 micro-rad and the resolution uncertainty D is 4 meters, which is the same as the previous chapter. Note that even though these uncertainties are small, they are obtainable, and reducing the accuracy will only reduce state accuracy results but will not invalidate the methods applied.

There are several new measurement constraints included in this chapter. All measurements are limited by line-of-sight which may be broken by the primary bodies. The optical measurements

have additional occultation constraints which will invalidate any measurements if the target is in shadow or if the image angle is in a solar exclusion region of $\psi = 30$ deg [33]. Figure 2.2 qualitatively depicts these restrictions. Note that visible brightness is not considered because it is assumed that an appropriate imaging system has been selected for the desired agent-targets combination. Cislunar lighting analyses have also been studied by other authors and are only applicable on a case-by-case basis [65]. Thus, as long as there is line-of-sight, no shadowing, and the image angle is outside of the exclusion region then an optical measurement is captured.

5.4 Optical-Only Spacecraft to Spacecraft Absolute Tracking in Cislunar Space

This section of the chapter seeks to answer the following question: can optical-only cislunar SSAT be successfully filtered to generate accurate state estimates and tracking? To do this the problem must be theoretically observable. Observability is assessed for the cislunar orbits of interest by checking the rank of the observability matrix as discussed in Section 3.2.2. If relative orbits are full rank observable then a filter must show it can produce consistent solutions. Filter consistency is determined by the true error being bounded by the filter 3-sigma covariance for 95% of the simulation time for non-maneuvering cases and boundedness in the final 12 hours for maneuvering cases. While this is not a statistical check, it practically ensures that the error is bounded. Additionally, we have found that statistical checks such as normalized estimation error squared (NEES) tests tend to struggle with matrix inversions and give poor results.

5.4.1 Observability Results

The first step to ascertaining the feasibility of optical-only SSAT in cislunar space is to check the observability of the selected nominal trajectories. The results in this section report the first instance a simulation achieves observability because once a system becomes observable it remains so for the rest of time. The observability rank test is applied to all agent and target trajectory combinations, excluding the case where both crafts inhabit the same orbit.

The Liaison results are recreated by studying the observability of systems with range measurements. It is important to note that the results obtained here are expected to differ numerically from the Liaison because they used an information matrix metric instead of an observability matrix metric. The information matrix approach produces longer times to achieve observability because it requires a level of equality in state information, while the observability matrix only requires that information for all states exists. The left side of Table 5.3 displays the time required to obtain observability with range measurements given the specified orbit pairs, with N/O meaning "not observable". Note that swapping target and agent orbits produce symmetric dynamics and measurements, thus making the table diagonally symmetric.

The obtained observability results coincide with Liaison's conclusion that relative range measurements lead to observable systems for orbits that are not co-planar [43]. It is worth noting that the unobservable subspace of the the co-planar PL2 and DRO combination is the \bar{e}_z projection of position and velocity for both spacecraft. Being able to deduce the null space is for this case is beneficial because the system would be considered observable if a planar model is used.

The observability matrix test is applied to the same orbit combinations, but with optical-only measurements, and results are accumulated on the right side of Table 5.3. The angle measurements are promising for two major reasons. First, angle measurements lead to fully observable systems in every case, including the co-planar case. Second, observability was achieved faster in all cases, even with a measurement cadence that is twenty times longer.

Collectively, these results demonstrate that optical-only measurements produce fully observable cislunar SSAT systems and achieve observability faster than range-only measurements. Another notable result is the predominance of 10 hours as the reported time, the rationale being that 10 hours is the time needed to obtain 12 independent measurements (with the first measurement at initialization time), which is the minimum number of measurements required to achieve a full state space. It is likely observability could be achieved faster for those orbits given a faster measurement cadence. On the other hand, it appears that the NL1 and PL2 orbit combination is the slowest to achieve observability regardless of measurement type. We hypothesize this slow observability time

is likely due to the comparatively large inter-satellite range leading to small relative motion, which in turn requires additional time to produce significant change. Finally, we must note that these results are simulation specific, and that observability as well as filter results for other scenarios must be assessed on a case-by-case basis.

Table 5.3: Hours to reach observability for range measurements (left) and optical measurements (right).

| | | Range Measurements | | | | | Optical Measurements | | | | |
|--------------|-----|--------------------|------|------|------|-----|----------------------|-----|-----|-----|-----|
| Target Orbit | | NL1 | PL2 | SL2 | QPO | DRO | NL1 | PL2 | SL2 | QPO | DRO |
| | | NL1 | - | | | | | - | | | |
| | PL2 | 61.7 | - | | | | 16 | - | | | |
| | SL2 | 32.9 | 38.1 | - | | | 10 | 10 | - | | |
| | QPO | 34.5 | 37.6 | 36.7 | - | | 10 | 10 | 10 | - | |
| | DRO | 48.5 | N/O | 30.1 | 29.1 | - | 12 | 14 | 10 | 10 | - |
| | | NL1 | PL2 | SL2 | QPO | DRO | NL1 | PL2 | SL2 | QPO | DRO |
| | | Agent Orbit | | | | | Agent Orbit | | | | |

5.4.2 State Estimation Solutions

With proof of theoretic observability obtained for the desired orbits of interest, we now seek to generate consistent filter solutions. First, non-maneuvering cases demonstrate the optical-only SSAT leads to long term convergence and reduced uncertainty for both vehicles over multiple orbit periods. These long term cases are run in a forwards filter and backed up with a Monte Carlo consistency analysis that examines multiple filter types. A filter is considered consistent in the non-maneuvering case if the true filter error norms are bounded by the filter’s 3-sigma covariance norms for over 95% of the simulation. Then, maneuvering cases are introduced and filtered over a single orbit period such that we can demonstrate nominal periodic operations. Because there is purposeful mismodeling in the maneuvering cases, filter consistency is reduced to bounded errors in the final 12 hours of simulation time.

For the remainder of the work, we down select to three relative orbit combinations of interest as a means to reduce the analysis effort. For every case the target starts at the lunar periapsis of the SL2 orbit, which is the planned location of the Lunar Gateway. The agent is initialized at lunar

periapsis of the NL1, PL2, and QPO orbits. Note that each subsequently listed orbit is closer than the prior and produces more favorable viewing geometries.

Initial uncertainty values assume that the agent is instantiated with ground-based navigation which it discontinues to start autonomous operations, while the target is known to some degree of accuracy. The initial uncertainty assumptions are realized by setting the agent 3-sigma to 30 km and 30 cm/s, and the target 3-sigma to 300 km and 300 cm/s. Reference [15] serves as the basis for these assumptions. The UKF hyperparameters which produced the most consistent filtering results had $\alpha = 0.5$, implying that the system is strongly nonlinear as expected. The beta and kappa hyperparameters were either 3 and 0 or 4 and 1 respectively, both of which seemed to produce approximately equivalent performance.

5.4.2.1 State Estimation of Non-Maneuvering Cases

While examining state estimation solutions for optical-only SSAT, it was immediately discovered that linearized filters are not reliably consistent. Of the many failed cases, it was quickly identified that most filter failures occur around lunar periapsis of the SL2. This point has been found to be highly nonlinear and has caused issues for filter convergence in previous studies. To abate the nonlinearity issues, we adopted the UKF and found that this filter is capable of successfully handling the nonlinearity in most cases where there is an appreciable range between vehicles.

Still, over short inter-satellite ranges the UKF may struggle such as with the QPO and SL2 case. Short ranges create challenges because SSAT requires differential motion between vehicles to obtain observability. The smaller the proximity between vehicles, the greater the similarity in their motion, and the harder it is for filters to reliably update. This is most prevalent at lunar periapsis of our QPO and SL2 case where the inter-satellite range is lowest. Therefore, for this case alone, we impose a 32 hr window centered at lunar apoapsis where no measurements are taken. With these considerations in mind, a Monte Carlo convergence analysis was performed for a duration of 3 SL2 orbit periods with an extended Kalman filter (EKF), extended square-root information filter (E-SRIF), and a UKF. The convergence results are tabulated in Table 5.4.

Table 5.4: Number of consistent filter solutions out of 100 Monte Carlo runs per case. There are no maneuvers considered here indicating optimal performance.

| Filter | Agent Orbit | | |
|--------|-------------|-----|-----|
| | NL1 | PL2 | QPO |
| EKF | 79 | 69 | 27 |
| E-SRIF | 86 | 65 | 30 |
| UKF | 100 | 98 | 84 |

To illustrate typical filter performances from the Monte Carlo analysis, we have plotted representative solutions from the UKF for each orbit combination in Figure 5.2-5.4. The plots show the true filter error in blue bounded by filter 3-sigma covariance in red as desired, and the covariance heading towards a steady state over the 3 orbit periods. Note the two jumps near $1/3$ and $2/3$ total simulation time for the QPO case is the blackout window with no measurements at lunar periapsis. Additionally, none of these cases include any filter mismodeling and therefore zero process noise, indicating "optimal" filter performance.

Figure 5.2: Typical UKF performance from Monte Carlo analysis with an agent on NL1 orbit and the target on SL2 orbit. Data is position and velocity norms of the filter errors and 3-sigma covariance.

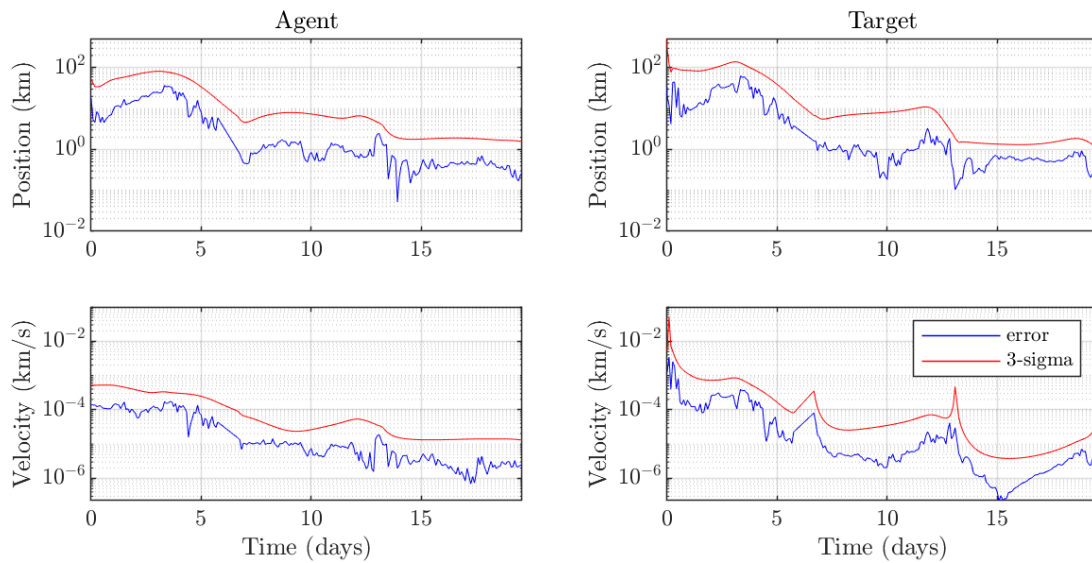


Figure 5.3: Typical UKF performance from Monte Carlo analysis with an agent on PL2 orbit and the target on SL2 orbit.

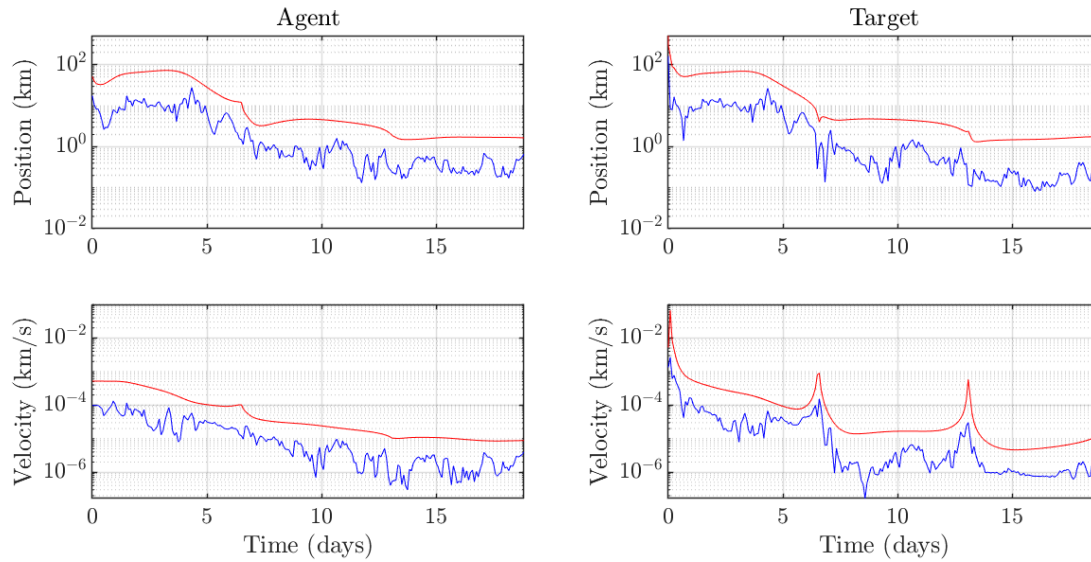
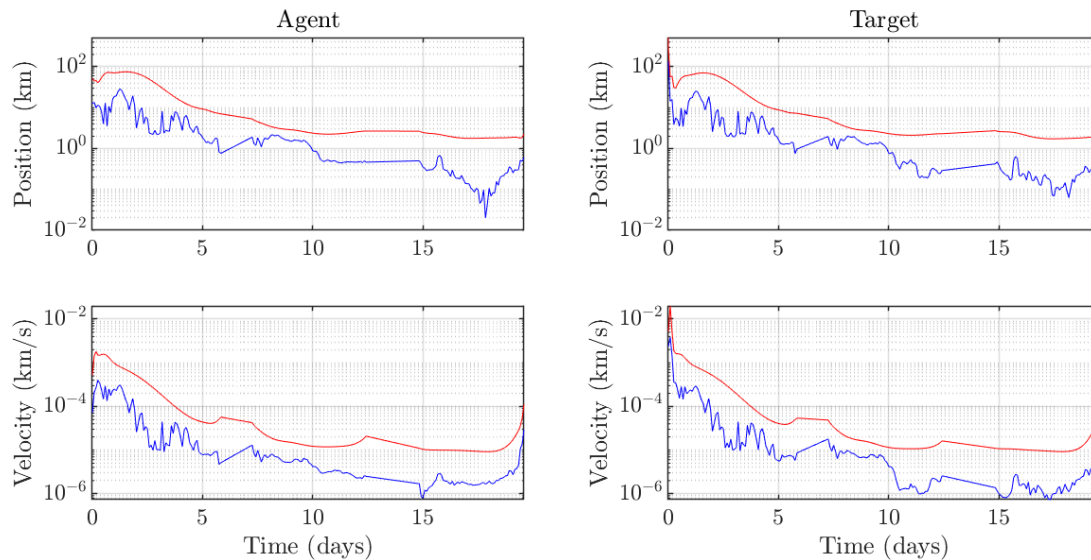


Figure 5.4: Typical UKF performance from Monte Carlo analysis with an agent on QPO and the target on SL2 orbit. Left plot is agent data, and the right plot is the target data.



These results show that nonlinear filtering is necessary for consistent convergence in this setting, but with a nonlinear filter optical-only SSAT is feasible over larger inter-satellite ranges.

The short range QPO to SL2 case struggles in comparison to other cases due to the similarity in first order motion between vehicles, and leaves room for improvement where a stronger nonlinear filter may be able better capture higher moments for further convergence. It is possible that even higher-order filters such as 3rd or 4th order EKF, Gaussian mixture model filters, or particle filters could further abate filter divergence and is an avenue of future work.

The need for higher order filters is highlighted when examining the failure points of non-convergent simulations. The majority of failures were caused by large initial errors (which are randomly drawn) or regions in which motion was highly similar. These failures emphasize two important and distinct difficulties associated with these simulations. First, the SSAT problem has a twelve dimensional state space with large associated uncertainties for both spacecraft, leading to uncertain filter behavior significantly outside of a linear operating regime. Second, differential motion, which is necessary for SSAT observability, dissipates as vehicles approach each other. Thus, there is some neighborhood of proximity operations where filters are unable to distinguish vehicle motion and where SSAT is not well posed. On the other hand, the neighborhood in which SSAT collapses is likely an ideal point to switch to a linearized relative motion model which would be easier to compute and more useful for proximity operations.

5.4.2.2 State and Control Estimation of Maneuvering Cases

With functional filters demonstrated for optical-only cislunar SSAT, the next objective is to apply the OCBE to estimate state and control solutions during simulations containing a maneuvering target in our desired setting. Unfortunately, the OCBE in its base state is a linearized filter, and as was illustrated in the previous section, linearized filters are insufficient in this setting. Fortunately, because the OCBE is a generalized Kalman filter, it can be modified to utilize unscented transformations. The unscented OCBE, henceforth referred to as the U-OCBE, is given in Appendix B.1.

Smoothed solutions from the forward U-OCBE are generated to estimate state and control vectors. Each simulation lasts for the duration of a single SL2 orbit period, which is approximately

6.5 days, to simulate periodic analysis of using smoothed data to estimate a potential maneuver. Additionally, since solutions are never processed through lunar periapsis we do not impose the 32 hr blackout window in the QPO for these cases.

The potential unmodeled maneuver is included at lunar apoapsis and has a magnitude of 40 mm/s. The chosen maneuver magnitude is selected because it is the average value of the quiet station keeping maneuvers to demonstrate that it can be detected without causing filter divergence. While the filter does not model these maneuvers it is instead given a 6 hour window of increased process noise around the event to account for them. The process noise is always included, whether a maneuver occurs or not. The lack of maneuver modeling is the only difference between the filter dynamics and the true simulation dynamics. Future work for higher fidelity models should consider other types errors, but increasing dynamical mismodeling is unlikely to cause substantial differences since the unmodeled maneuvers are orders-of-magnitude more significant than other errors (e.g. non-spherical gravity fields, solar radiation pressure, forth bodies, etc.).

To assess the U-OCBE performance of the filter, three quantities are examined. First, the true filter error is plotted alongside the filter covariance. For the filter to be functioning correctly the true error must be bounded by covariance. Second, the norm of the estimated optimal control profile will be displayed to ensure appropriate control estimate magnitudes. Finally, the OCBE metric is normalized by a 99% confidence bound and plotted. The OCBE metric is a residual based metric, and for the ballistic OCBE it is the filter Mahalanobis distance. Individual outliers are expected due to the stochastic measurement generation but repeated outliers in the metric indicate improper filter tuning. The normalization means that if the metric goes above 1 then there is statistical outlier.

Figure 5.5: Smoothed U-OCBE error and covariance norms with the agent on NL1 with target on SL2. Simulations start at periapsis, have an unmodeled maneuver at apoapsis at 3.25 days, and conclude after an orbit period of 6.5 days.

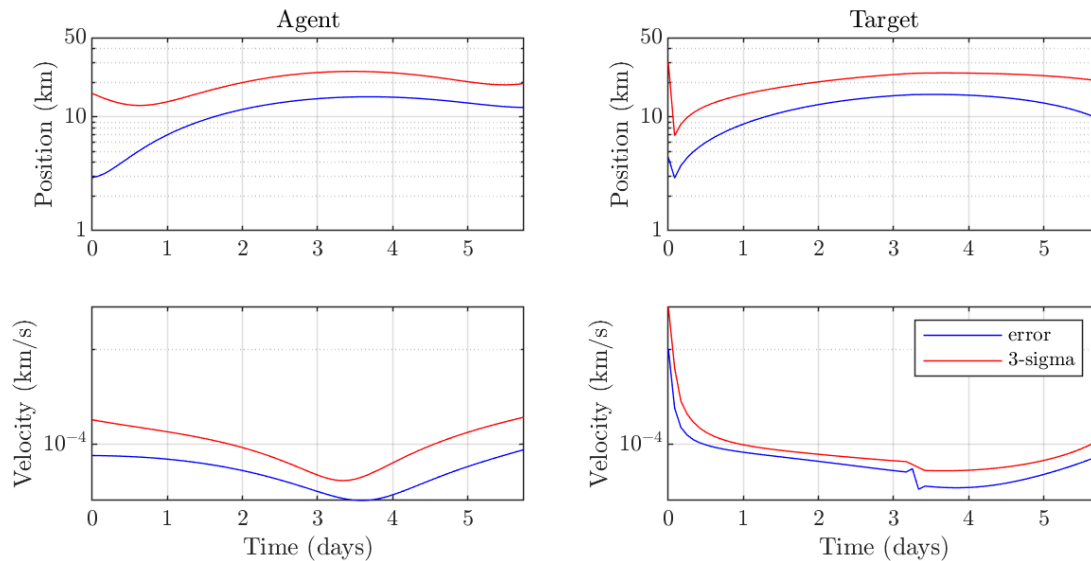



Figure 5.6: Smoothed control estimates (left) and OCBE metrics (right) with the agent on NL1 and target on SL2.



The first case studied is the scenario with the agent located on the NL1 orbit. Figure 5.5

illustrates the smoothed state error for both the agent and target. The errors in blue are bounded by the filter covariance in red, confirming a consistent filtering solution. Figure 5.6 plots the OCBE control estimate and metric. The control reflects that uncertainty in the state estimate is most significant on the target during lunar apoapsis as designed, and is roughly $6 * 3600s * 8 \times 10^{-4} mm/s^2 \approx 17 mm/s$ which is about half the true $40 mm/s$ maneuver. The reconstruction is underestimated but we expect these small maneuvers to be hard to capture over larger distances. We will perform a Monte Carlo analysis of the maneuver reconstruction properties in the maneuver classification section later. The filter metric exhibits variations about the mean as desired and is not indicative of any mismodeling.

Figure 5.7: Smoothed U-OCBE error and covariance norms with the agent on PL2 with target on SL2.





Figure 5.8: Smoothed control estimates (left) and OCBE metrics (right) with the agent on PL2 and target on SL2.



Next, we study the case where the agent is on the PL2 orbit. Figure 5.7 reveals that the errors are bounded by the covariance again, affirming proper filter behavior. Figure 5.8 confirms appropriate the control estimate behavior with an approximate value of $6 * 3600s * 7 \times 10^{-4} mm/s^2 \cong 15 mm/s$. The filter metric is again well behaved. These results are similar in nature to the NL1 orbit with only marginal differences.

Figure 5.9: Smoothed U-OCBE error and covariance norms with the agent on QPO with target on SL2.





Figure 5.10: Smoothed control estimates (left) and OCBE metrics (right) with the agent on QPO and target on SL2.



Finally, we simulate the case where the agent is placed on the QPO about the target on the SL2 which is a periodic orbit with the same orbital period. Figure 5.9 shows that the covariance

properly bounds the errors, demonstrating a consistent solution. Placing the agent on the QPO produces an order of magnitude better position estimates and marginally better velocity estimates. The estimate improvement is likely a result of the fact that angular measurements map to smaller position errors over shorter distances, and the QPO has a much smaller average range than both the NL1 and PL2 orbits. The improved state estimates impact the OCBE control estimate, plotted in Fig. 5.10, and shows better reconstruction with an approximate value of $6*3600s*1.75 \times 10^{-3}mm/s^2 \approx 38mm/s$ which is close to the $40mm/s$ maneuver. The filter metric does have two independent outliers, which is not statistically indicative of mismodeling since independent outliers are to be expected.

A Monte Carlo convergence analysis is executed to assess filter performance. With the unmodeled maneuver, state errors at lunar apoapsis are likely be outside the filter covariance bounds, but accurate on either side of the maneuver if the final forward tracking solution is accurate. To appropriately capture consistency in the final stage of filtering, the convergence criteria is relaxed such that the errors must only be bounded by the 3-sigma filter uncertainty during the final 12 hours of simulation. This time only an EKF and the U-OCBE are considered, since the E-SRIF was so similar to the EKF, and the results are reported in Table 5.5.

Table 5.5: Number of consistent filter solutions out of 100 Monte Carlo runs per case. There are maneuvers considered here indicating sub-optimal performance.

| δv Policy | Filter | Orbits | | |
|-------------------|--------|--------|-----|-----|
| | | NL1 | PL2 | QPO |
| Noisy | EKF | 0 | 0 | 1 |
| | U-OCBE | 100 | 100 | 94 |
| Quiet | EKF | 1 | 1 | 0 |
| | U-OCBE | 100 | 100 | 93 |

Table 5.5 shows strongly preferable convergence for the U-OCBE due to two possibilities. First, as shown in the previous section, the UKF modification significantly improves performance. Second, the U-OCBE internal process noise calculation more accurately represents mismodeling by tying the process noise to system dynamics. Additionally, these results report a higher consistency

rate than the non-maneuvering cases because the addition of process noise inflates the covariance and leads to more converged, but also under confident, filter solutions.

Overall, the results shown above demonstrate that it is feasible to generate state estimate solutions from the U-OCBE for SSAT in cislunar space with optical-only measurements on the orbits of interest. Furthermore, 3-sigma estimates generated under this framework are commonly on the order of 10 km and 10 cm/s for both vehicles. According to Reference [15] this translates to a marginally better than medium accuracy orbit determination solution from the Deep Space Network. The accuracy of these estimates is enhanced with longer observation windows as the nonlinearities build up to increase system information, which is a known phenomenon from Liaison results and demonstrated in the non-maneuvering simulations [43]. Since we have demonstrated that the U-OCBE can successfully generate state estimates and control estimates under the SSAT frameworks with optical-only measurements, we now seek classify the control estimates to label spacecraft behavior.

5.5 Maneuver Classification of Unmodeled Events

Maneuver classification provides insight into spacecraft operational behavior and demonstrates utility beyond simple state estimation or maneuver reconstruction. For example, being aware of a maneuver may assist in orbit determination, but understanding the maneuver's purposes provides predictive capabilities on future behaviors which includes the likelihood of another maneuver occurring again. Similarly, purely reconstructing a maneuver from measurements may fit the observed data, but to gain insight into how well the data creates a coherent representation of reality requires proper classification and assessment. A further benefit of classification is that it can take poorly reconstructed maneuvers that may not fit reality and align them with known events through training. Doing so makes perfect reconstruction unnecessary. Therefore, we seek to apply classification algorithms to take imperfect state and maneuver reconstructions from the U-OCBE, and use these to classify target spacecraft behaviors when its actions are unknown. Additionally, demonstrating the ability to properly classify the data is another indication that the filter solutions

are accurate.

This section outlines a maneuver classification method that utilizes the optimal control estimate from the U-OCBE. The classifier input is defined as the integral of the estimated control policy which is a surrogate for the total system mismodeling. Subsequently, two supervised learning classifiers are trained and tested to determine their effectiveness for maneuver classification given tracking solutions from cislunar SSAT with optical measurements. The first method is a custom algorithm called the one-hot maneuver classifier which is computationally efficient, quick to train, and intuitive. The second classifier is a support vector machine (SVM) which comes directly from machine learning literature.

In practice the maneuver classification procedure is a two-step process. To begin, compute and smooth the U-OCBE solutions to generate an estimated state and control. Then, given the state estimate is trusted, pass the control estimate to the classifier to label the solution for potential maneuvers. The training and Monte Carlo analysis must be run ahead of time in order to have a classifier that can run online. This process is primarily limited by the sigma-point integration of the U-OCBE and is only marginally more computationally intensive than a UKF.

5.5.1 Classifier Input and Setup

The simulated spacecraft have four actions available to them which have been described above in the section on problem modeling: no maneuver, station keeping, orbit departures, and generalized maneuvers. These actions form the four classification labels we wish to identify because they are generally indicative of long-term behavior. It is pertinent to note that even though we chose the dynamical manifolds as labels for their long-term behavioral properties, any vector in the control space could be associated with maneuvers if so desired.

Attempting to classify the whole, continuous, estimated control policy is not only intractable but not practical because the control is impulsive. Instead, the integral of the smoothed control

profile as defined as:

$$\mathbf{z} = \int_{t_0}^{t_f} \bar{\mathbf{u}}(t|t_f) dt \quad (5.9)$$

The control integral \mathbf{z} is used as the classifier input. This estimated control integral is more appropriate for classification because it represents the total estimated effort applied to the system. Additionally, if the U-OCBE provides an accurate reconstruction of the event then the estimated control integral should closely resemble the true maneuver.

Both the SVM and custom method are supervised learning algorithms which require data for training and testing. The data is generated from a Monte Carlo analysis of the U-OCBE event reconstruction capabilities. This serves the dual purposes of analyzing the accuracy of the reconstructed events and obtaining training data. The data will be split 80%-20% for training-testing respectively, with testing set being sequestered for validation.

5.5.2 Classification Algorithms

5.5.2.1 One-Hot Maneuver Manifold Classifier

As an alternative to traditional machine learning methods, we have developed a custom maneuver detection classifier based on the assumption that the estimated control integral \mathbf{z} is closely related to the true maneuvers. The custom classifier is referred to as the one-hot maneuver classifier, because the classification probabilities are encoded in one-hot vectors. It is important to note that this method will only be accurate for cases where the control reconstruction is close to the true control. As a tradeoff, this method does not require any hyperparameter tuning, is easy to implement, fast to train, and simple to understand.

The one-hot classifier functions on the same fundamental principle as an SVM; it projects the input data into a domain where it is easily separable. The one-hot method also takes advantage of the accurate reconstruction capabilities of the U-OCBE control estimates to define class boundaries in a hierarchical order. An incoming estimated control integral must be above α to be classified as a maneuver. Then, if the angle between the estimated control and a manifold is less than $\cos^{-1}(\beta)$,

it is assigned to that manifold. In order for the classifier to correctly label events, the parameters α, β must be appropriately selected which we do in the form of an optimization problem.

The one-hot classifier optimization problem minimizes false classification through the loss function:

$$J = - \sum_{i=1}^{N_L} \mathbf{L}_i^T \bar{\mathbf{L}}_i \quad (5.10)$$

where the term \mathbf{L}_i is the true probability that the sample i came from one of the four possible action classes which is defined as:

$$\mathbf{L}_i = \begin{bmatrix} pdf(\text{None}) \\ pdf(\text{SK}) \\ pdf(\text{OD}) \\ pdf(\text{GM}) \end{bmatrix} \quad (5.11)$$

The truth vector is set to unity for the true class, and zero for all other classes, making it a one-hot vector. The variable $\bar{\mathbf{L}}_i$ is the predicted probability, which we also assign a one-hot vector, though other approaches such as a soft max could be used. If the class probability vector is ordered in the same manner as Eq. (5.11), then the predicted probability is calculated:

$$\bar{\mathbf{L}}_i = \begin{cases} [1, 0, 0, 0]^T & \text{if } \|\mathbf{z}\| \leq \alpha \\ [0, 1, 0, 0]^T & \text{else if } \left| \frac{\hat{\mathbf{e}}_s^T \mathbf{z}}{\|\mathbf{z}\|} \right| \geq \beta \\ [0, 0, 1, 0]^T & \text{else if } \left| \frac{\hat{\mathbf{e}}_u^T \mathbf{z}}{\|\mathbf{z}\|} \right| \geq \beta \\ [0, 0, 0, 1]^T & \text{otherwise} \end{cases} \quad (5.12)$$

Given a number of training samples N_l , the parameters α and β can be optimized by an algorithm of choice to minimize the loss function. We have simply employed MATLAB's `fminsearch` optimizer.

5.5.2.2 Support Vector Machines

We will compare the one-hot manifold classifier to a classic machine learning algorithm, the support vector machine. The underlying principle of a SVM is to apply a kernel function that projects input data into a space where it is easily separable by a hyperplane [84, 85]. For the

implementation of the SVM algorithm and supporting utility functions, the authors used the Sci-Kit Learn library. Sci-Kit Learn is a powerful, well-documented, and easy to use open source library with the desired SVM functionality [89].

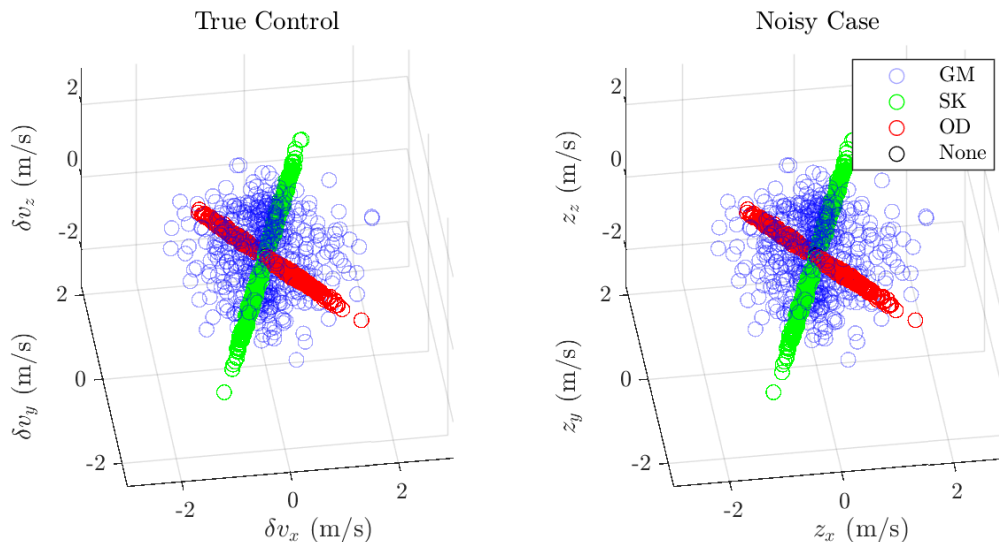
Our SVM is implemented as a pipeline which includes a normalizing scalar that is calculated during training of the SVM. The radial basis function is chosen as the kernel because it behaves as a similarity measure which will take advantage of the any clustering that occurs within the estimated control integral. The hyperparameters C and gamma were chosen by employing Sci-Kit's grid search function which returned the values 500,000 and 0.01 respectively.

5.5.3 Maneuver Classification Results

5.5.3.1 Maneuver Reconstruction Analysis


The simplicity of the one-hot maneuver classifier algorithm leverages accurate event reconstruction from the U-OCBE control estimate and is expected to deteriorate if the reconstruction does not align with true events. Hence, before jumping into the classification training, a Monte Carlo analysis of the U-OCBE estimated control integral is performed to determine the accuracy of its maneuver reconstruction abilities. As previously mentioned, the data generated from the Monte Carlo analysis also serves as the training-testing data for the classifiers later. All of the U-OCBE control estimates are created from filters that are computed in the same manner as the maneuvering case portion of the state estimation study. Thus, the U-OCBE is run for a single 9:2 NRHO orbit period starting at lunar periapsis where the target may include a maneuver at lunar apoapsis. The same three agent-target orbit pairs are used again with identical uncertainties and filter settings.

Figure 5.11: True maneuvers (left) and estimated integrals z (right) for an agent on the QPO with the target on the SL2 and noisy maneuvers.



For the first scenario, the agent is located on the QPO and the target is located on the SL2 and given noisy maneuvers. Figure 5.11 plots the resulting estimated control integral and true control side-by-side. This figure shows that the estimated control integral very closely approximates the actual maneuver. This precise reconstruction of true events, along with the minimal overlap between maneuver classes, suggests that classification on this system should be highly accurate.

Figure 5.12: Histogram of $\|z\|$ for an agent on the QPO with the target on the SL2 and noisy maneuvers.



Next, the estimated control integral norm is plotted in Fig. 5.12 to determine if the U-OCBE can correctly differentiate maneuvering and non-maneuvering cases. The histogram not only confirms that non-maneuvering cases have an order of magnitude smaller estimated control integrals, but that the estimated maneuvering distribution closely resembles the designed station keeping mean and covariance.

Figure 5.13: True maneuvers (left) and estimated integrals z (right) for an agent on the QPO with the target on the SL2 and quiet maneuvers.

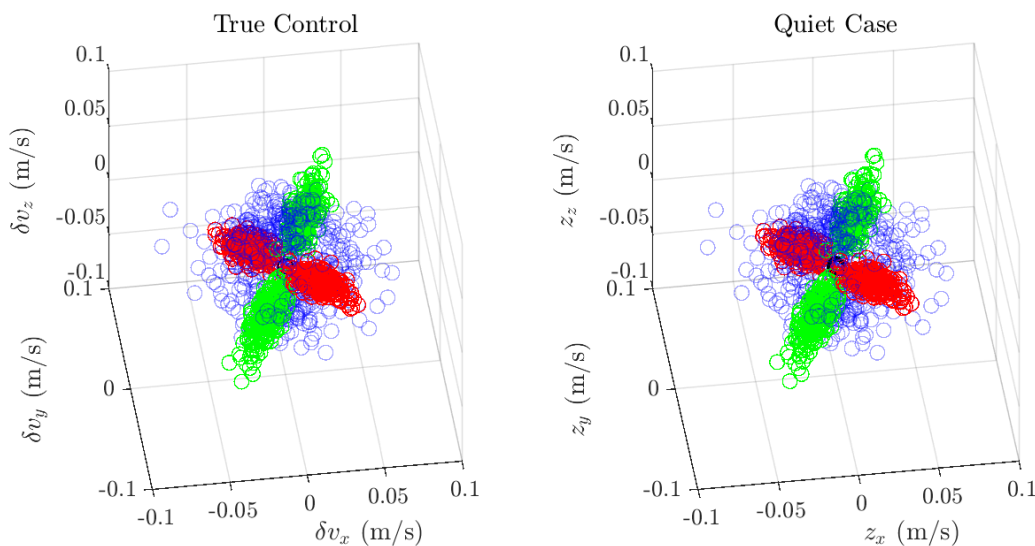

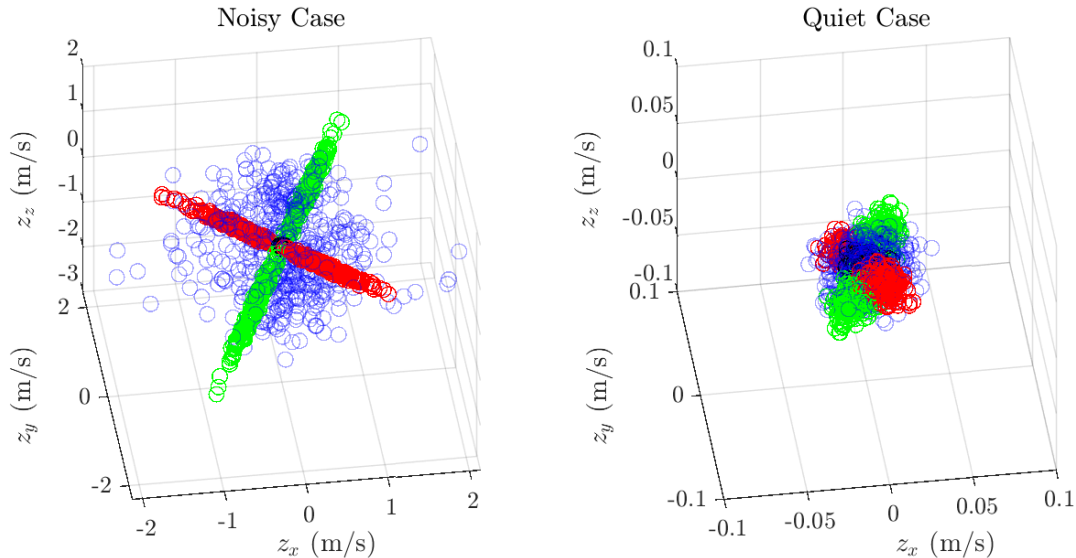


Figure 5.14: Histogram of $\|z\|$ for an agent on the QPO with the target on the SL2 and quiet maneuvers.



Subsequently, the same orbital geometry was analyzed given quiet target maneuvers. Figure 5.13 presents the estimated control integral and actual control side-by-side. From the true scatter plot, the maneuver execution error is much more impactful for the quiet maneuvers. This additional noise implies that classification should be more difficult on the quiet case but should still be feasible due to only minor errors in the U-OCBE maneuver reconstruction. Next, figure 5.14 shows the norm of the estimated control. Once again, the non-maneuvering cases are consistently an order of magnitude smaller, the maneuvering cases closely approximate the true distribution, and the plot has been cropped to focus on the maneuvering cases.

Figure 5.15: Estimated integrals z of the noisy (left) and quiet (right) maneuvers for an agent on the PL2 with the target on the SL2.



Now we examine the case with the agent located on the PL2. Figure 5.15 gives the noisy and quiet U-OCBE control estimates side-by-side. The true maneuvers are not displayed because they are equivalent to the QPO cases shown in Fig. 5.11-5.13. Once again, the U-OCBE estimated control integral closely approximates truth in the noisy case. However, unlike the QPO results, the estimated control integral for quiet maneuvers has an appreciable overlap between maneuver classes which are not seen in the true maneuvers. Fortunately, the structure of the estimated

control integrals still loosely aligns the stable and unstable manifolds. The general structure of the estimated control integral does appear separable such that classification should still be possible even if it is likely to be worse.

Figure 5.16: Histogram of $\|z\|$ of the noisy (left) and quiet (right) maneuvers for an agent on PL2 with target on SL2.


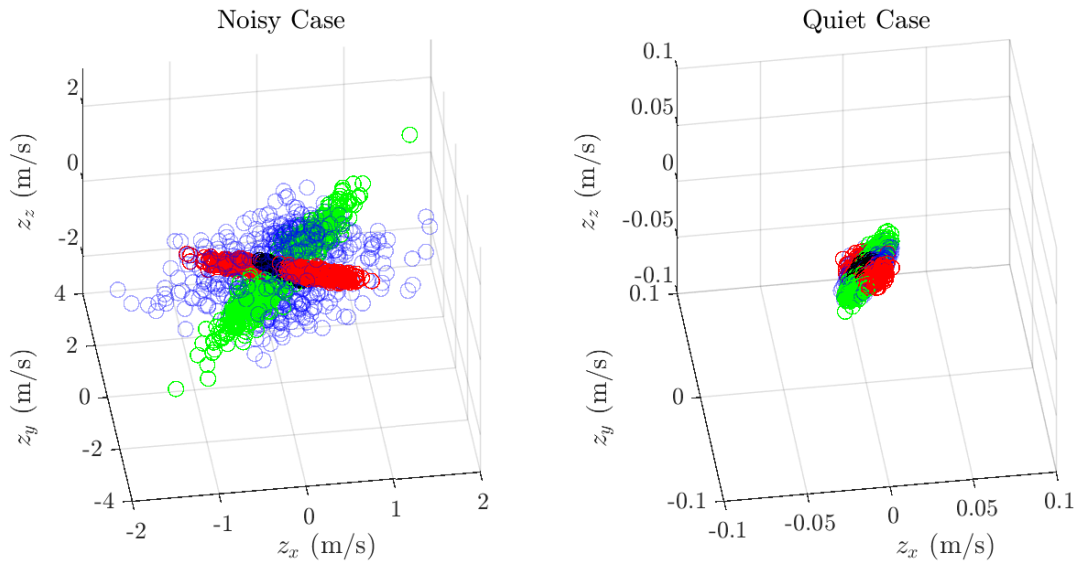


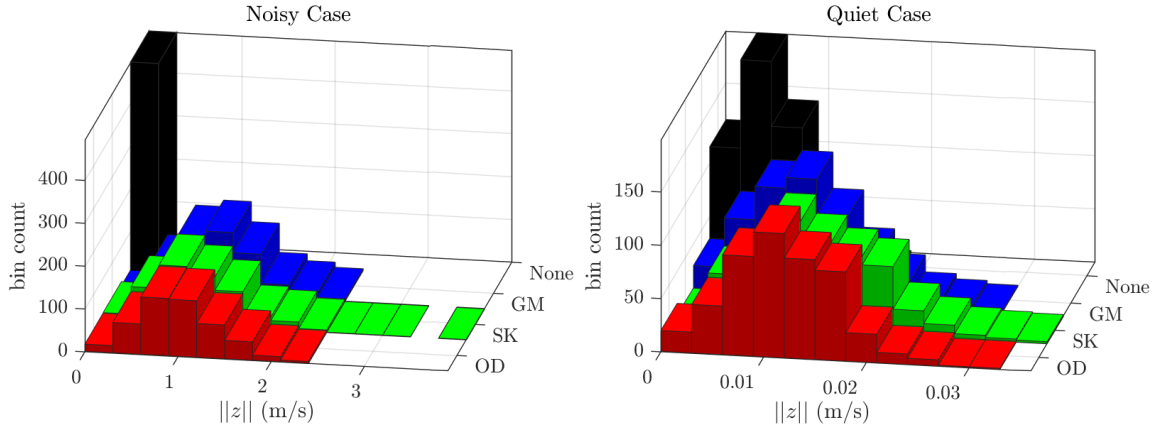
Figure 5.16 plots the norms of the two cases. The noisy case shows a clear distinction between estimated control integrals with and without maneuvers, but the quiet case shows significantly more overlap between maneuvering and non-maneuvering cases indicating an increased challenge for identifying if a maneuver occurs.

Figure 5.17: Estimated integrals \mathbf{z} of the noisy (left) and quiet (right) maneuvers for an agent on the NL1 with the target on the SL2.



Finally, the agent is placed on NL1 orbit and simulated. Figure 5.17 reveals that the U-OCBE estimated control integral is skewed when compared to the true control, particularly with minimal variance along the average range vector between the spacecraft. Note that the skew is particularly prevalent with the quiet data. The average range vector is relatively static over the observation window because of the large inter-satellite distance, which means very little control information is obtained along that direction. The lack of information is most likely the cause of the poor reconstruction seen in this case. Figure 5.17 also exposes that the quiet maneuvers are appreciably smaller than the true maneuvers. Again this is likely caused by the large distances, since optical measurements lead to larger positional uncertainty over longer distance and cannot detect such a small maneuver.

Figure 5.18: Histogram of $\|z\|$ of the noisy (left) and quiet (right) maneuvers for an agent on NL1 with target on SL2.



The histograms of the norms of the estimated control integral for both noisy and quiet maneuvers are plotted in Fig. 5.18. Similar to the PL1 case, the noisy case is well separated, and the quiet case is not, but the NL1 results have a greater overlapping set. Thus, we expect the NL1 set to produce the least accurate classification.

Overall, the uncertainty quantification of the estimated control integral clearly indicates that closer orbits generate more accurate event reconstruction with angular measurements. While these orbits with smaller relative distances lead to better event reconstruction, it also leads to more difficulties for filter convergence as indicated in the previous section. Therefore, a balance between event reconstruction and filter convergence must be considered when selecting relative orbits using SSAT. It is also evident that both the real maneuvers and the estimated control integrals contain structures that are readily identifiable, and that classification should be feasible even with low-complexity classifiers. In every case examined these structures were more clearly separated over shorter distances and larger maneuvers which we might expect to lead to better classifier performance in those cases.

5.5.3.2 Maneuver Classification Results

The data generated during the U-OCBE control quantification is now used to train and test the classification methods. The dataset contains 500 samples per action for a total of 2000 data points. Again, we use a random 80%-20% split for training-testing of the classification algorithms. The sequestered test data is then displayed in a confusion matrix that is row normalized. Note, the classes have been abbreviated for plotting purposes as follows: no maneuver (none), station keeping (SK), orbit departure (OD), and general maneuver (GM).

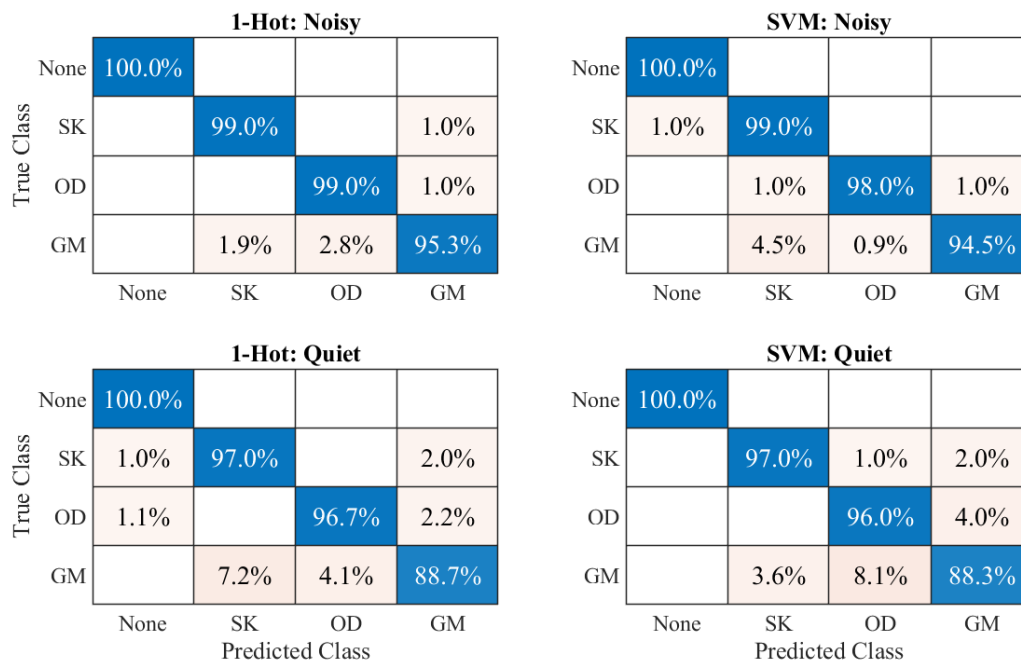


Figure 5.19: Confusion matrices for an agent on the QPO.

Figure 5.20: Confusion matrices for an agent on the PL2.

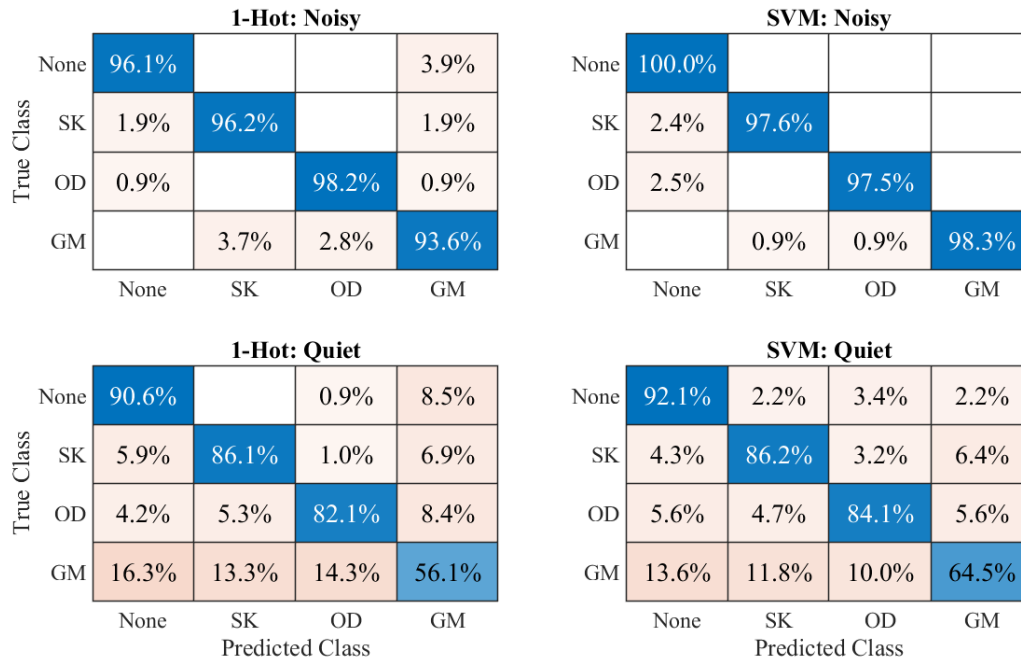
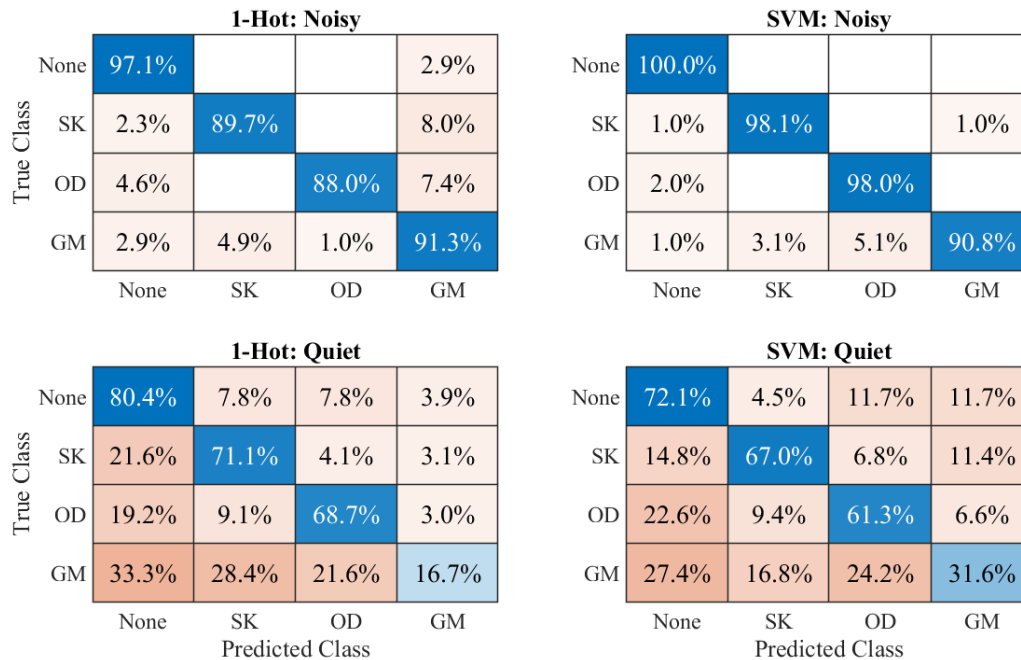


Figure 5.21: Confusion matrices for an agent on the NL1.



To begin with, we examine the case with the agent located on the QPO. Figure 5.19 gives the confusion matrices for both classification algorithms and both station keeping policies. The algorithms are separated by columns, and the station keeping is separated along the rows. This figure confirms that both algorithms perform well with nearly identical performance. This indicates that the dynamically motivated one-hot method is a good representation of the problem. It is also worth noting that the minimum accuracy is 88% across all cases. In summary, given an agent on a QPO, it is possible to achieve a high level of accuracy for maneuver classification.

Next, we processed the case where the agent was placed on the PL2 orbit. The confusion matrices in Fig. 5.20 again show that the two classifiers are mostly comparable but now the SVM can be seen to outperform the one-hot method, especially with respect to generalized maneuvers. This is expected because we saw greater maneuver reconstruction inaccuracies for this case during the uncertainty quantification analysis. The reconstruction error leads to reduced performance for the one-hot method because the metrics no longer align as accurately with the manifolds. This misclassification is more pronounced for the quiet case than the noisy case as expected, but overall, the PL2 orbit is less accurate than the QPO. It does appear that both algorithms are reasonable classifiers for this problem still, but performance is highly dependent on maneuver magnitude.

Finally, the confusion matrices with the agent located on the NL1 orbit are shown in Fig. 5.21. There are two major trends continuing with these results. First, maneuver magnitude and inter-spacecraft distance are major factors when determining the accuracy of classification results. Because this orbit is the farthest orbit it leads to the worst results, with the quiet case being particularly poor. Second, the SVM continues to outperform the one-hot method as the estimated control integral inaccuracies increase.

Cumulatively, the SVM classifier generates better results at the cost of additional algorithm complexity. However, the one-hot method is still appealing for maneuver detection if the orbits are close because of its accuracy and simplicity. Unfortunately, the one-hot method degrades as range increases because the U-OCBE reconstruction deviates from reality, and thus the classification problem doesn't adhere to dynamics as rigorously. It is also clear that closer relative orbits produce

better classification in general, which is expected since angle measurements produce lower positional uncertainty over shorter ranges due to shorter arc lengths.

5.6 Conclusion

This work sets out to determine if optical-only measurements are sufficient for SSAT in the cislunar environment. The first step in achieving this goal was proving that optical-only measurements produce a fully observable SSAT system via a linearized observability analysis. Then filtering solutions are generated for non-maneuvering and maneuvering cases to determine if accurate state solutions can be generated. The non-maneuvering cases show that nonlinear filtering is necessary for consistent filter convergence. Subsequently, the U-OCBE demonstrates desirable performance in the face of unmodeled maneuvers. Even though it was found that the UKF updates are an enabling factor for cislunar SSAT there are other higher-order filters such as second-order EKFs, Gaussian mixture models, or particle filters which will likely outperform the UKF updates at the cost of additional computation time.

With functional filters demonstrated, maneuver classification methods are tested to label U-OCBE estimated control integrals as non-maneuvering, station keeping, orbit departures, or general maneuvers. Uncertainty quantification of the estimated control integral shows that U-OCBE event reconstruction depends strongly on relative orbit distance, with shorter distance leading to better reconstructions. Then two classifiers are tested: a new one-hot maneuver manifold classifier and a SVM. The one-hot method is a dynamically grounded model that leverages proper event reconstruction and proves to be efficient when relative spacecraft distances are small. The SVM method was generally more accurate and provided effective classification even with poor event reconstruction. The classification results were best with the quasi-periodic orbit generated about a periodic orbit with the same base frequency. This observation scheme achieved the best accuracy at a minimum of 88% correct classification, even given an average maneuver size of only 40 mm/s.

There are various avenues of future work related to each objective described above. First and foremost, there is precedence to implement higher-order filters such as Gaussian mixture models

to further improve filter performance. Additionally, there are several enhancements to the classification work which include: adding continuous time maneuvers, applying neural network classifiers like perceptrons, and automating the filtering plus classifier algorithms. This dissertation has sought and developed methods towards the enhancement of autonomous navigation methods for distributed space systems. The fundamental enabling approach for this goal was spacecraft to spacecraft absolute tracking (SSAT), where relative measurements inside of a distributed space system (DSS) are used to estimate the full state of every vehicle in the system. SSAT is possible because the nonlinearities of spaceflight dynamics cause vehicles in the system to have differential motion which causes the absolute state to become locally observable. We used this knowledge to build navigation maps in two and three body dynamical settings and showed that larger state differences, generally implemented as large inter-satellite range, lead to well conditioned and more certain state estimates. The inclusion of perturbations such as J2 do theoretically improve system observability, but do not create tangible impacts on the level of information obtained. On the other hand, the inclusion of multiple targets can increase state information if each target can be measured to a sufficient degree. An unexpected result was that orbital resonances tend to lead to more information through repeated viewing patterns. A final note is that without nonlinearities this approach is not possible, and therefore nearby proximity operations where linear approximations are valid do not fare well in this scenario. Fortunately, where SSAT falters is an ideal location to swap to a linearized relative motion model which would be better suited for proximity operations anyways.

There is ample future work related to SSAT. First, a full nonlinear observability analysis should be applied to SSAT for a more rigorous derivation of observability. Nonlinear observability analysis should not only examine dynamics, but also the impact of nonlinear measurements as well. Second, filter simulations in full fidelity scenarios should be executed to simulate more realistic performance, including under mismodeled dynamics. We know from Chapter 5 that linearized filters struggle with SSAT, and there is a strong precedence to examine Gaussian mixture models and higher order filters in tandem with SSAT. Third, optimal target observation selection, or sensor

tasking, must be explored for multi-target SSAT simulations. One promising approach to sensor tasking is Monte Carlo tree search methods which have been applied to ground-based telescopes already. Finally, quantifying when to swap from a SSAT model to a linearized relative motion model should be examined in more detail.

While SSAT can produce fully observable systems, it may take significant time for information to accumulate. This is particularly true for single sensor systems such as optical-only observation. Optical sensors are lightweight, widely available, and fully autonomous sensors but they lack range information. Therefore, we developed analytic guidance strategies which can easily be calculated onboard a spacecraft to improve the state estimate along a desired projection, such as the range projection when using optics. We began with a geometric model to show how variations in viewing geometries due to maneuvers lead to a range estimate. Combining the geometric analysis with a linear maneuver model led to a heuristic policy which captured optimal maneuvers to reduce range uncertainty over short time periods. The heuristic inspired a surrogate cost function, one which maximizes measurement deviations, that was solved analytically and closely approximated the optimal maneuver. Because fuel is a critical resource, we then combined station keeping policies with our information gathering guidance which successfully predicted the near optimal maneuvers to maximize information.

Again, there is an abundance of future work related to information gathering. First and foremost, maneuver execution errors must be included and measured with an accelerometer to determine how much noise is allowable, or how accurate of a sensor is necessary, to still retain the benefits of changing viewing geometries. Next, alternative station keeping policies should be included and tested in combination with the information gathering guidance policies to reduce overall fuel costs. Another branch of research would be to examine different sensors, such as range sensors, to determine if the maximum measurement deviation policy also works well under those circumstances. Finally, these policies should be simulated in a filtering scenario to determine more realistic performance characteristics. Again, since this is a proximity operation, it may be that this is not well suited for SSAT unless one vehicle is very well known. Fortunately, the developed

policies work for single or multiple vehicle state estimates and therefore are still applicable in a relative formulation as well.

The last topic covered was implementing cislunar SSAT inside a filter which could also perform maneuver classification of unknown unmodeled events. We began by demonstrating successful filtering solutions using an unscented Kalman filter (UKF) when it was found that regular linearized filters struggled. Because of the success of the UKF, we modified the ballistic linear optimal control based estimator (OCBE) with the unscented transformation in a new algorithm called the ballistic unscented OCBE called the U-OCBE. Then the U-OCBE was applied to scenarios with unknown maneuvers and still demonstrated consistent filter solutions. The estimated control from the U-OCBE was used to train and test machine learning algorithms to classify unknown maneuvers along a dynamical manifold. Traditional support vector machines performed admirably but require significant effort to tune and train. As an alternative a one-hot maneuver manifold classifier was developed, and performed well when the U-OCBE was able to accurately reconstruct the unknown maneuver.

Next steps for applied SSAT and maneuver classification should start with higher order filters to further improve filter success rates. Once that has been completed, then alternative classifiers such as perceptrons or random forests should be investigated. To make matters more complex, continuous time maneuvers should also be included to determine if they can be classified as well. When classifying non-impulsive events, it may be beneficial to use the continuous time control estimate instead of the integral as was done here. Finally, integrating classification into the filtering process would greatly alleviate the workload caused by having them separate.

Bibliography

- [1] M. K. Ben-Larbi, K. F. Pozo, T. Haylok, M. Choi, B. Grzesik, A. Haas, D. Krupke, H. Konstanski, V. Schaus, S. P. Fekete, et al., “Towards the automated operations of large distributed satellite systems. part 1: Review and paradigm shifts,” Advances in Space Research, vol. 67, no. 11, pp. 3598–3619, 2021.
- [2] B. D. Tapley, S. Bettadpur, M. Watkins, and C. Reigber, “The gravity recovery and climate experiment: Mission overview and early results,” Geophysical research letters, vol. 31, no. 9, 2004.
- [3] J. Stuart, D. Velez, M. Vincent, J. Senent, J. Lee, R. R. Karimi, and K. Boudad, “Sunrise is coming: Mission design and navigation for the first space-based radio interferometer,” in 33rd AAS/AIAA Space Flight Mechanics Meeting, 2023.
- [4] F. Markley, “Autonomous navigation using landmark and intersatellite data,” in Astrodynamics conference, p. 1987, 1984.
- [5] K. Hill and G. H. Born, “Autonomous interplanetary orbit determination using satellite-to-satellite tracking,” Journal of guidance, control, and dynamics, vol. 30, no. 3, pp. 679–686, 2007.
- [6] Y. Li and A. Zhang, “Observability analysis and autonomous navigation for two satellites with relative position measurements,” Acta Astronautica, vol. 163, pp. 77–86, 2019.
- [7] J. A. Sullivan, Nonlinear Angles-Only Orbit Estimation for Autonomous Distributed Space Systems. PhD thesis, Stanford University, 2020.
- [8] M. L. Psiaki, “Autonomous orbit determination for two spacecraft from relative position measurements,” Journal of Guidance, Control, and Dynamics, vol. 22, no. 2, pp. 305–312, 1999.
- [9] S. C. Nardone and V. J. Aidala, “Observability criteria for bearings-only target motion analysis,” IEEE Transactions on Aerospace and Electronic systems, no. 2, pp. 162–166, 1981.
- [10] J. P. Helferty and D. R. Mudgett, “Optimal observer trajectories for bearings only tracking by minimizing the trace of the cramer-rao lower bound,” in Proceedings of 32nd IEEE Conference on Decision and Control, pp. 936–939, IEEE, 1993.
- [11] R. J. V. Chari, Autonomous orbital rendezvous using angles-only navigation. PhD thesis, Massachusetts Institute of Technology, 2001.

- [12] D. C. Woffinden and D. K. Geller, "Optimal orbital rendezvous maneuvering for angles-only navigation," Journal of guidance, control, and dynamics, vol. 32, no. 4, pp. 1382–1387, 2009.
- [13] J. Grzymisch and W. Fichter, "Analytic optimal observability maneuvers for in-orbit bearings-only rendezvous," Journal of Guidance, Control, and Dynamics, vol. 37, no. 5, pp. 1658–1664, 2014.
- [14] M. Ceresoli, G. Zanotti, and M. Lavagna, "Bearing-only navigation for proximity operations on cis-lunar non-keplerian orbits," in 72nd International Astronautical Congress (IAC 2021), pp. 1–10, 2021.
- [15] D. Davis, S. Bhatt, K. Howell, J.-W. Jang, R. Whitley, F. Clark, D. Guzzetti, E. Zimovan, and G. Barton, "Orbit maintenance and navigation of human spacecraft at cislunar near rectilinear halo orbits," 27th AAS/AIAA Space Flight Mechanics Meeting, 2017.
- [16] C. P. Newman, D. C. Davis, R. J. Whitley, J. R. Guinn, and M. S. Ryne, "Stationkeeping, orbit determination, and attitude control for spacecraft in near rectilinear halo orbits," 2018.
- [17] E. L. Jenson and D. J. Scheeres, "Semianalytical measures of nonlinearity based on tensor eigenpairs."
- [18] N. L. Parrish, M. J. Bolliger, E. Kayser, M. R. Thompson, J. S. Parker, B. W. Cheetham, D. C. Davis, and D. J. Sweeney, "Near rectilinear halo orbit determination with simulated dsn observations," in AIAA Scitech 2020 Forum, p. 1700, 2020.
- [19] D. P. Lubey, Maneuver Detection and Reconstruction in Data Sparse Systems with an Optimal Control Based Estimator. PhD thesis, University of Colorado at Boulder, 2015.
- [20] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," Proceedings of the IEEE, vol. 92, no. 3, pp. 401–422, 2004.
- [21] N. Baresi, Z. P. Olikara, and D. J. Scheeres, "Fully numerical methods for continuing families of quasi-periodic invariant tori in astrodynamics," The Journal of the astronautical sciences, vol. 65, pp. 157–182, 2018.
- [22] H. Schaub and J. L. Junkins, Analytical mechanics of space systems. Aiaa, 2003.
- [23] S. I. Roumeliotis and G. A. Bekey, "Distributed multirobot localization," IEEE transactions on robotics and automation, vol. 18, no. 5, pp. 781–795, 2002.
- [24] Y. Shen, S. Mazuelas, and M. Z. Win, "Network navigation: Theory and interpretation," IEEE Journal on Selected Areas in Communications, vol. 30, no. 9, pp. 1823–1834, 2012.
- [25] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," IEEE Transactions on robotics, vol. 32, no. 6, pp. 1309–1332, 2016.
- [26] F. Meyer, O. Hlinka, H. Wymeersch, E. Riegler, and F. Hlawatsch, "Distributed localization and tracking of mobile networks including noncooperative objects," IEEE Transactions on Signal and Information Processing over Networks, vol. 2, no. 1, pp. 57–71, 2015.

- [27] G. Chai, C. Lin, Z. Lin, and M. Fu, “Single landmark based collaborative multi-agent localization with time-varying range measurements and information sharing,” Systems & Control Letters, vol. 87, pp. 56–63, 2016.
- [28] F. Schiano and R. Tron, “The dynamic bearing observability matrix nonlinear observability and estimation for multi-agent systems,” in 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 3669–3676, IEEE, 2018.
- [29] K. A. Hill, Autonomous navigation in libration point orbits. PhD thesis, University of Colorado at Boulder, 2007.
- [30] Y. Ou, H. Zhang, and B. Li, “Absolute orbit determination using line-of-sight vector measurements between formation flying spacecraft,” Astrophysics and Space Science, vol. 363, no. 4, pp. 1–13, 2018.
- [31] P. R. Cachim, J. Gomes, and R. Ventura, “Autonomous orbit determination for satellite formations using relative sensing: Observability analysis and optimization,” Acta Astronautica, vol. 200, pp. 301–315, 2022.
- [32] J. A. Greaves and D. J. Scheeres, “Optical only co-estimation and maneuver classification in the cislunar regime,” in Rocky Mountain AAS GN&C Conference, 2022.
- [33] N. Bradley, Z. Olikara, S. Bhaskaran, and B. Young, “Cislunar navigation accuracy using optical observations of natural and artificial targets,” Journal of Spacecraft and Rockets, vol. 57, no. 4, pp. 777–792, 2020.
- [34] J. A. Christian and E. G. Lightsey, “Review of options for autonomous cislunar navigation,” Journal of Spacecraft and Rockets, vol. 46, no. 5, pp. 1023–1036, 2009.
- [35] D. C. Woffinden and D. K. Geller, “Observability criteria for angles-only navigation,” IEEE Transactions on Aerospace and Electronic Systems, vol. 45, no. 3, pp. 1194–1208, 2009.
- [36] J. R. Yim, J. L. Crassidis, and J. L. Junkins, “Autonomous orbit navigation of two spacecraft system using relative line of sight vector measurements,” in Proceedings of the AAS Space Flight Mechanics Meeting, 2004.
- [37] E. Kaufman, T. A. Lovell, and T. Lee, “Nonlinear observability for relative orbit determination with angles-only measurements,” The Journal of the Astronautical Sciences, vol. 63, pp. 60–80, 2016.
- [38] E. M. Ward and R. J. Patel, “On the observability of spacecraft navigation using landmarks,” in 2022 IEEE Aerospace Conference (AERO), pp. 1–14, IEEE, 2022.
- [39] J. Kulik and D. Savransky, “State transition tensors for passive angles-only relative orbit determination,” in 33rd AAS/AIAA Space Flight Mechanics Meeting, 2023.
- [40] B. Schutz, B. Tapley, and G. H. Born, Statistical orbit determination. Elsevier, 2004.
- [41] D. S. Bernstein, “Matrix mathematics,” in Matrix Mathematics, Princeton university press, 2009.

- [42] R. S. Park and D. J. Scheeres, “Nonlinear mapping of gaussian statistics: theory and applications to spacecraft trajectory design,” Journal of guidance, Control, and Dynamics, vol. 29, no. 6, pp. 1367–1375, 2006.
- [43] K. Hill, G. H. Born, and M. W. Lo, “Linked, autonomous, interplanetary satellite orbit navigation (liaison) in lunar halo orbits,” 2006.
- [44] S. Fedeler, M. Holzinger, and W. Whitacre, “Sensor tasking in the cislunar regime using monte carlo tree search,” Advances in Space Research, vol. 70, no. 3, pp. 792–811, 2022.
- [45] A. Feldbaum, “Theory of dual control,” Avtomatika i telemekhanika, vol. 21, no. 9, 1960.
- [46] E. Tse, Y. Bar-Shalom, and L. Meier, “Wide-sense adaptive dual control for nonlinear stochastic systems,” IEEE Transactions on Automatic Control, vol. 18, no. 2, pp. 98–108, 1973.
- [47] R. J. Casler Jr, “Dual-control guidance strategy for homing interceptors taking angle-only measurements,” Journal of Guidance and Control, vol. 1, no. 1, pp. 63–70, 1978.
- [48] N. M. Filatov and H. Unbehauen, “Survey of adaptive dual control methods,” IEE Proceedings-Control Theory and Applications, vol. 147, no. 1, pp. 118–128, 2000.
- [49] B. Wittenmark, “Adaptive dual control methods: An overview,” Adaptive Systems in Control and Signal Processing 1995, pp. 67–72, 1995.
- [50] L. P. Kaelbling, M. L. Littman, and A. W. Moore, “Reinforcement learning: A survey,” Journal of artificial intelligence research, vol. 4, pp. 237–285, 1996.
- [51] M. Coggan, “Exploration and exploitation in reinforcement learning,” Research supervised by Prof. Doina Precup, CRA-W DMP Project at McGill University, 2004.
- [52] G. L. Mariottini, F. Morbidi, D. Prattichizzo, N. Vander Valk, N. Michael, G. Pappas, and K. Daniilidis, “Vision-based localization for leader–follower formation control,” IEEE Transactions on Robotics, vol. 25, no. 6, pp. 1431–1438, 2009.
- [53] G. L. Mariottini, S. Martini, and M. B. Egerstedt, “A switching active sensing strategy to maintain observability for vision-based formation control,” in 2009 IEEE International Conference on Robotics and Automation, pp. 2637–2642, IEEE, 2009.
- [54] J. Le Cadre and C. Jauffret, “Discrete-time observability and estimability analysis for bearings-only target motion analysis,” IEEE Transactions on Aerospace and Electronic Systems, vol. 33, no. 1, pp. 178–201, 1997.
- [55] S. E. Hammel, Optimal observer motion for bearings-only localization and tracking. University of Rhode Island, 1988.
- [56] J.-M. Passerieux and D. Van Cappel, “Optimal observer maneuver for bearings-only tracking,” IEEE Transactions on Aerospace and Electronic Systems, vol. 34, no. 3, pp. 777–788, 1998.
- [57] J. Kawaguchi, T. Hashimoto, T. Kubota, S. Sawai, and G. Fujii, “Autonomous optical guidance and navigation strategy around a small body,” Journal of guidance, control, and dynamics, vol. 20, no. 5, pp. 1010–1017, 1997.

- [58] F. D’Onofrio, G. Bucchioni, and M. Innocenti, “Bearings-only guidance in cis-lunar rendezvous,” Journal of Guidance, Control, and Dynamics, vol. 44, no. 10, pp. 1862–1874, 2021.
- [59] S. Takahashi and D. J. Scheeres, “Autonomous exploration of a small near-earth asteroid,” Journal of Guidance, Control, and Dynamics, vol. 44, no. 4, pp. 701–718, 2021.
- [60] S.-H. Mok, J. Pi, and H. Bang, “One-step rendezvous guidance for improving observability in bearings-only navigation,” Advances in Space Research, vol. 66, no. 11, pp. 2689–2702, 2020.
- [61] D. Smitherman and A. Schnell, “Gateway lunar habitat modules as the basis for a modular mars transit habitat,” in 2020 IEEE Aerospace Conference, pp. 1–12, IEEE, 2020.
- [62] S. V. Wagner and T. D. Goodson, “Execution-error modeling and analysis of the cassini-huygens spacecraft through 2007,” 2008.
- [63] M. Smith, D. Craig, N. Herrmann, E. Mahoney, J. Krezel, N. McIntyre, and K. Goodliff, “The artemis program: An overview of nasa’s activities to return humans to the moon,” in 2020 IEEE Aerospace Conference, pp. 1–10, IEEE, 2020.
- [64] M. Holzinger, C. Chow, and P. Garretson, “A primer on cislunar space,” 2021.
- [65] J. K. Vendl and M. J. Holzinger, “Cislunar periodic orbit analysis for persistent space object detection capability,” Journal of Spacecraft and Rockets, pp. 1–12, 2021.
- [66] E. Fowler, D. Paley, and S. Hurtt, “Observability metrics for space-based cislunar domain awareness,” in 31st AIAA/AAS Space Flight Mechanics Meeting, 2021.
- [67] J. A. Greaves and D. J. Scheeres, “Absolute and autonomous navigation for distributed space systems from relative measurements,” in 33rd AAS/AIAA Space Flight Mechanics Meeting, 2023.
- [68] J. Sullivan, A. W. Koenig, J. Kruger, and S. D’Amico, “Generalized angles-only navigation architecture for autonomous distributed space systems,” Journal of Guidance, Control, and Dynamics, vol. 44, no. 6, pp. 1087–1105, 2021.
- [69] J. A. Greaves and D. J. Scheeres, “Observation and maneuver detection for cislunar vehicles: Using optical measurements and the optimal control based estimator,” The Journal of the Astronautical Sciences, vol. 68, no. 4, pp. 826–854, 2021.
- [70] B. Ristic, S. Arulampalam, and N. Gordon, Beyond the Kalman filter: Particle filters for tracking applications. Artech house, 2003.
- [71] D. Simon, Optimal State Estimation: Kalman, H Infinity, and Nonlinear Approaches. USA: Wiley-Interscience, 2006.
- [72] E. A. Wan and R. Van Der Merwe, “The unscented kalman filter for nonlinear estimation,” in Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373), pp. 153–158, Ieee, 2000.
- [73] G. A. McIntyre and K. J. Hintz, “Comparison of several maneuvering target tracking models,” in Signal processing, sensor fusion, and target recognition VII, vol. 3374, pp. 48–63, SPIE, 1998.

- [74] R. P. Patera, "Space event detection method," Journal of Spacecraft and Rockets, vol. 45, no. 3, pp. 554–559, 2008.
- [75] S. Lemmens and H. Krag, "Two-line-elements-based maneuver detection methods for satellites in low earth orbit," Journal of Guidance, Control, and Dynamics, vol. 37, no. 3, pp. 860–868, 2014.
- [76] S. Lee and I. Hwang, "Interacting multiple model estimation for spacecraft maneuver detection and characterization," in AIAA Guidance, Navigation, and Control Conference, p. 1333, 2015.
- [77] A. D. Jaunzemis, M. V. Mathew, and M. J. Holzinger, "Control cost and mahalanobis distance binary hypothesis testing for spacecraft maneuver detection," Journal of Guidance, Control, and Dynamics, vol. 39, no. 9, pp. 2058–2072, 2016.
- [78] D. Aguilar Marsillach and M. J. Holzinger, "Spacecraft custody maintenance and maneuver detection using robotic telescopes and reachable sets," Journal of Guidance, Control, and Dynamics, vol. 44, no. 4, pp. 667–683, 2021.
- [79] E. Mazor, A. Averbuch, Y. Bar-Shalom, and J. Dayan, "Interacting multiple model methods in target tracking: a survey," IEEE Transactions on aerospace and electronic systems, vol. 34, no. 1, pp. 103–123, 1998.
- [80] M. J. Holzinger, D. J. Scheeres, and K. T. Alfriend, "Object correlation, maneuver detection, and characterization using control distance metrics," Journal of Guidance, Control, and Dynamics, vol. 35, no. 4, pp. 1312–1325, 2012.
- [81] K. Rivera, A. Zara, D. Aguilar-Marsillach, M. Holzinger, I. Elliott, and N. Bosanac, "Patterns of life and maneuver detection for cislunar trajectory maintenance," in 22nd Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Maui, HI, 2021.
- [82] S. J. Fedeler, Optical Sensor Tasking using Monte Carlo Tree Search. PhD thesis, University of Colorado at Boulder, 2023.
- [83] D. P. Lubey and D. J. Scheeres, "Identifying and estimating mismodeled dynamics via optimal control policies and distance metrics," Journal of Guidance, Control, and Dynamics, vol. 37, no. 5, pp. 1512–1523, 2014.
- [84] S. B. Kotsiantis, I. Zaharakis, P. Pintelas, et al., "Supervised machine learning: A review of classification techniques," Emerging artificial intelligence applications in computer engineering, vol. 160, no. 1, pp. 3–24, 2007.
- [85] S. B. Kotsiantis, I. D. Zaharakis, and P. E. Pintelas, "Machine learning: a review of classification and combining techniques," Artificial Intelligence Review, vol. 26, no. 3, pp. 159–190, 2006.
- [86] D. Folta, M. Woodard, and D. Cosgrove, "Stationkeeping of the first earth-moon libration orbiters: The artemis mission," in AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska, 2011.
- [87] M. Bester, D. Cosgrove, S. Frey, J. Marchese, A. Burgart, M. Lewis, B. Roberts, J. Thorsness, J. McDonald, D. Pease, G. Picard, M. Eckert, and R. Dumlao, "Artemis operations - experiences and lessons learned," in 2014 IEEE Aerospace Conference, pp. 1–20, 2014.

- [88] N. Baresi, Spacecraft formation flight on quasi-periodic invariant tori. PhD thesis, University of Colorado at Boulder, 2017.
- [89] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
- [90] B. Schutz, B. Tapley, and G. H. Born, Statistical orbit determination. Elsevier, 2004.
- [91] C. D’Souza, T. Crain, F. Clark, and J. Getchius, “Orion cislunar guidance and navigation,” in AIAA Guidance, Navigation and Control Conference and Exhibit, p. 6681, 2007.
- [92] S. Yun, K. Tuggle, R. Zanetti, and C. D’Souza, “Sensor configuration trade study for navigation in near rectilinear halo orbits,” The Journal of the Astronautical Sciences, vol. 67, no. 4, pp. 1755–1774, 2020.
- [93] J. Williams, D. E. Lee, R. J. Whitley, K. A. Bokelmann, D. C. Davis, and C. F. Berry, “Targeting cislunar near rectilinear halo orbits for human space exploration,” 27th AAS/AIAA Space Flight Mechanics Meeting, 2017.
- [94] B. C. Levy, Principles of signal detection and parameter estimation. Springer Science & Business Media, 2008.
- [95] R. Whitley and R. Martinez, “Options for staging orbits in cislunar space,” in 2016 IEEE Aerospace Conference, pp. 1–9, IEEE, 2016.
- [96] V. Muralidharan, A. Weiss, and U. V. Kalabic, “Control strategy for long-term station-keeping on near-rectilinear halo orbits,” in AIAA Scitech 2020 Forum, p. 1459, 2020.

Appendix A

Ballistic Linear OCBE for Spacecraft to Spacecraft Absolute Tracking

This appendix covers the ballistic linear form of the OCBE. The ballistic assumption means that there is no known mismodeling or control estimate prior to the filter pass, which greatly simplifies the filter equations. The linear assumption means that a state deviation is estimated similar to other linearization methods and linearized filters. All of the following is an abbreviation from Reference [19], and it is highly recommended that that additional details are sought there. We also adopt a similar notation for estimated states to this reference which denotes the estimate of \mathbf{x} at time k given measurement information up to time $k - 1$ as $\bar{\mathbf{x}}_{k|k-1}$.

The underlying principle of the OCBE is that it dynamically links measurement updates with an estimated optimal control policy. The estimated optimal control policy, $\bar{\mathbf{u}}$, is assumed to be additive and affine in control such that the filter dynamics are $\dot{\bar{\mathbf{x}}} = \mathbf{f}(\bar{\mathbf{x}}) + \mathbf{B}_{\text{ce}}\bar{\mathbf{u}}$. The matrix \mathbf{B}_{ce} is the estimated control input matrix, which for space systems is assumed to be an unknown acceleration, and for the SSAT system is given by the matrix:

$$\mathbf{B}_{\text{ce}} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{I}_3 & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_3 \end{bmatrix} \quad (\text{A.1})$$

The OCBE estimated control is incorporated into the state estimate as a least squares Lagrangian term with the weight matrix $\mathbf{Q}_{\text{ce}}(\tau)$. The matrix $\mathbf{Q}_{\text{ce}}(\tau)$ is akin to a process noise matrix and should be tuned in similar manner. For the SSAT system the matrix $\mathbf{Q}_{\text{ce}}(\tau)$ is a matrix which

weights each axis equally and has a separate weighting for the agent and target vehicles which are denoted $q_a(t)$ and $q_t(t)$ and is given by:

$$\mathbf{Q}_{\text{ce}}(\tau) = \begin{bmatrix} q_a(\tau)\mathbf{I}_3 & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & q_t(\tau)\mathbf{I}_3 \end{bmatrix} \quad (\text{A.2})$$

The estimated control is calculated as a function of costates, denoted \mathbf{p} , which have their dynamics propagated inside of the STM to enforce optimality constraints. As a consequence, the OCBE STM Φ_{ce} is actually of size $2N$ where N is the number of system states (note for the SSAT system with n_{sc} spacecraft, each of state size n_{state} , gives a system state size $N = n_{sc}n_{state}$). The OCBE STM is partitioned into state and costate blocks:

$$\Phi_{\text{ce}}(\tau_k, \tau_{k-1}) = \begin{bmatrix} \Phi_{xx}(\tau_k, \tau_{k-1}) & \Phi_{xp}(\tau_k, \tau_{k-1}) \\ \Phi_{px}(\tau_k, \tau_{k-1}) & \Phi_{pp}(\tau_k, \tau_{k-1}) \end{bmatrix} \quad (\text{A.3})$$

with the following STM dynamics:

$$\dot{\Phi}_{\text{ce}}(\tau_i) = \mathbf{A}_{\text{ce}}(\mathbf{x})\Phi_{\text{ce}}(\tau_i, \tau_0) \quad (\text{A.4})$$

$$\mathbf{A}_{\text{ce}}(\mathbf{x}) = \begin{bmatrix} \mathbf{A}(\mathbf{x}) & -\mathbf{B}_{\text{ce}}\mathbf{Q}_{\text{ce}}(\tau)\mathbf{B}_{\text{ce}}^T \\ \mathbf{0}_{n \times n} & -\mathbf{A}(\mathbf{x})^T \end{bmatrix} \quad (\text{A.5})$$

From the aforementioned definitions it is self evident that the block Φ_{xx} is the standard state STM Φ . For notational convenience the time dependencies on the STM is assumed to be propagating from τ_{k-1} to τ_k and is dropped unless otherwise noted.

With the base principles of the ballistic linear OCBE defined the forwards time, measurement, and costate update equations can be presented. The time update is nearly identical to a Kalman time update, but the process noise matrix is prescribed by the state and costates dynamics. The time update equations for the state and covariance are:

$$\delta \bar{\mathbf{x}}_{k|k-1} = \Phi_{xx} \delta \bar{\mathbf{x}}_{k-1|k-1} \quad (\text{A.6})$$

$$\mathbf{P}_{k|k-1} = \Phi_{xx} \mathbf{P}_{k-1|k-1} \Phi_{xx}^T - \Phi_{xp} \Phi_{xx}^T \quad (\text{A.7})$$

The measurement update equations are exactly equivalent to the Kalman filter and are given by:

$$\mathbf{K}_k = \mathbf{P}_{k|k-1} \mathbf{H}_k^T (\mathbf{R}_k + \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^T)^{-1} \quad (\text{A.8})$$

$$\delta \bar{\mathbf{x}}_{k|k} = \delta \bar{\mathbf{x}}_{k|k-1} + \mathbf{K}_k (\delta \mathbf{y}_k - \mathbf{H}_k \delta \bar{\mathbf{x}}_{k|k-1}) \quad (\text{A.9})$$

$$\mathbf{P}_{k|k} = (\mathbf{I}_N - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_{k|k-1} (\mathbf{I}_N - \mathbf{K}_k \mathbf{H}_k)^T + \mathbf{K}_k \mathbf{R}_k \mathbf{K}_k^T \quad (\text{A.10})$$

Finally, the OCBE unique costate update and the control estimate are calculated:

$$\mathbf{K}_{k-1} = \mathbf{P}_{k-1|k-1} \Phi_{xx}^T \mathbf{H}_k^T (\mathbf{R}_k + \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^T)^{-1} \quad (\text{A.11})$$

$$\delta \bar{\mathbf{p}}_{k-1|k} = -\mathbf{P}_{k-1|k-1}^{-1} \mathbf{K}_{k-1} (\delta \mathbf{y}_k - \mathbf{H}_k \delta \bar{\mathbf{x}}_{k|k-1}) \quad (\text{A.12})$$

$$\delta \bar{\mathbf{u}}(\tau|\tau_k) = -\mathbf{Q}_{ce}(\tau) \mathbf{B}_{ce}^T \Phi_{pp}(\tau, \tau_{k-1}) \delta \bar{\mathbf{p}}_{k-1|k} \quad (\text{A.13})$$

Together the time, measurement, and costate update equations fully cover the ballistic linear OCBE as a forwards filter. The costate and control updates contains three important properties. First, the costate update is only a function of the state estimate and not a function of costate estimates. Second, the costate update produces a smoothed estimate of costate at the time prior to the measurement which is indicative of the one step smoothing property of the ballistic OCBE. Third, and finally, the control estimate is continuous in time and not discrete, therefore it must be integrated to approximate discrete events such as impulsive maneuvers.

The ballistic linear OCBE comes with a smoothing algorithm which is, once again, very similar the Kalman smoother. The major difference is costate update step which is unique to the OCBE. The smoother for measurements up to time τ_l is a set of equations which are calculated recursively:

$$\mathbf{S}_{k-1} = \mathbf{P}_{k-1|k-1} \Phi_{xx}^T \mathbf{P}_{k|k-1}^{-1} \quad (\text{A.14})$$

$$\delta \bar{\mathbf{x}}_{k-1|l} = \delta \bar{\mathbf{x}}_{k-1|k-1} + \mathbf{S}_{k-1} (\delta \bar{\mathbf{x}}_{k|l} - \Phi_{xx} \delta \bar{\mathbf{x}}_{k-1|k-1}) \quad (\text{A.15})$$

$$\mathbf{P}_{k-1|l} = \mathbf{P}_{k-1|k-1} + \mathbf{S}_{k-1} (\mathbf{P}_{k|l} - \mathbf{P}_{k|k-1}) \mathbf{S}_{k-1}^T \quad (\text{A.16})$$

$$\delta \bar{\mathbf{p}}_{k-1|l} = -\mathbf{P}_{k-1|k-1}^{-1} \mathbf{S}_{k-1} (\delta \bar{\mathbf{x}}_{k|l} - \Phi_{xx} \delta \bar{\mathbf{x}}_{k-1|k-1}) \quad (\text{A.17})$$

$$\delta \bar{\mathbf{u}}(\tau|\tau_l) = -\mathbf{Q}_{ce}(\tau) \mathbf{B}_{ce}^T \Phi_{pp}(\tau, \tau_{k-1}) \mathbf{P}_{k-1|k-1}^{-1} (\delta \bar{\mathbf{x}}_{k-1|k-1} - \delta \bar{\mathbf{x}}_{k-1|l}) \quad (\text{A.18})$$

There is one final piece of the ballistic linear OCBE that is used in this paper, which is the filter distance metric. The OCBE distance metrics are statistical tests which identify mismodeling in the filter, and the ballistic linear OCBE distance metric is effectively a Mahalanobis distance on the predicted measurement residual which is calculated:

$$D_k = \frac{1}{2}(\delta \mathbf{y}_k - \mathbf{H}_k \delta \mathbf{x}_{k|k-1})^T (\mathbf{R}_k + \mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k)^{-1} (\delta \mathbf{y}_k - \mathbf{H}_k \delta \mathbf{x}_{k|k-1}) \quad (\text{A.19})$$

Because of Gaussian assumptions the ballistic linear OCBE distance metric is actually a chi-squared test with m degrees of freedom, where m is the size of the measurement vector at time t_k , such that $D_k = \frac{1}{2} \chi_m^2$. Therefore the mean and variance of the distance metric are given by:

$$\mu_D = m/2 \quad (\text{A.20})$$

$$\sigma_D^2 = m/2 \quad (\text{A.21})$$

A probabilistic limit can be selected from the chi-squared distribution to easily identify if filter mismodeling is occurring.

Appendix B

Modifications for the ballistic OCBE

The OCBE is a generalization of the KF, and as such it can be modified to adhere to common KF techniques and practices. The necessary adaptations to the ballistic OCBE are made based on the fact that the OCBE time and measurement updates are only functions of the state and not dependent on the co-state. Hence, only the state needs to be accurately tracked and the co-state can be calculated based on the state afterwards. This realization means that an implementation of the OCBE can utilize any filter update as desired, as long as it carries with it the OCBE STM necessary for the co-state and control updates.

B.1 Unscented Modification

The value provided by the UKF is instrumental for optical-only cislunar SSAT consistency, so the implementation of a ballistic unscented OCBE to provide accurate state updates in conjunction with an estimated control is warranted. The modification process takes advantage of the single step smoothing property of the OCBE to effectively compute the control estimates as a linear profile along the center sigma-point of the unscented transformation. To appropriately present the unscented OCBE (U-OCBE) the unscented transformations must be introduced.

The forwards unscented transform the estimated mean state $\bar{\mathbf{x}}$ and covariance \mathbf{P} and returns the sigma-points $\bar{\chi}^{0:2n}$ using the function $\bar{\chi}^{0:2n} = S(\bar{\mathbf{x}}, \mathbf{P})$. The inverse transformation takes the sigma-points and includes an additional noise matrix \mathbf{R} to return a new mean and covariance, written as $\bar{\mathbf{x}}, \mathbf{P} = S^{-1}(\bar{\chi}^{0:2n}, \mathbf{R})$ (this is not a true inverse, just shorthand notation). Additionally,

the cross-covariance term between the state and measurement is denoted \mathbf{P}_{xy} . The implementation of the UKF and its associated functions were sourced from the References [71]. Finally, the a combined dynamics function is used to propagate both sigma-points and an associated STM along the sigma-point nominal using the notation $\bar{\chi}_k, \Phi_{\infty}(\tau_k, \tau_{k-1}) = \mathbf{F}(\chi_{k-1})$. This single function represents a single ODE solver call because the STM is solved simultaneously with along the point propagation nominal.

With the unscented transformations covered, we can now incorporate it into the ballistic OCBE. The intrinsic single-step smoothing property of the ballistic OCBE facilitates the use of higher-order filter update equations. The smoothing property establishes that the costate estimate is a function of the current state and smoothed prior state [19]. Therefore, the OCBE can effectively use any update equation and still calculate the costate by carrying the OCBE STM plus a single-step smoother. We apply this concept by utilizing the UKF time and measurement updates to produce an accurate state estimate, and then use a single-step smoother to calculate the costate. The combination of the UKF state updates and OCBE costate estimation is what defines the U-OCBE and effectively makes for a three step filter.

The first step of the U-OCBE is the time update, which is nearly identical to the UKF time update. The defining different is the mean sigma-point propagation used to calculate the OCBE STM which provides the estimated control noise term \mathbf{Q} which is defined by Eq. (B.4). The full U-OCBE time update is:

$$\bar{\chi}_{k-1|k-1}^{0:2n} = S(\bar{\mathbf{x}}_{k-1|k-1}, \mathbf{P}_{k-1|k-1}) \quad (\text{B.1})$$

$$\bar{\chi}_{k|k-1}^0, \Phi_{\infty} = \mathbf{F}(\bar{\chi}_{k-1|k-1}^0) \quad (\text{B.2})$$

$$\bar{\chi}_{k|k-1}^{1:2n} = \mathbf{F}(\bar{\chi}_{k-1|k-1}^{1:2n}) \quad (\text{B.3})$$

$$\mathbf{Q} = -\Phi_{xp} \Phi_{xx}^T \quad (\text{B.4})$$

$$\bar{\mathbf{x}}_{k|k-1}, \mathbf{P}_{k|k-1} = S^{-1}(\bar{\chi}_{k|k-1}^{0:2n}, \mathbf{Q}) \quad (\text{B.5})$$

The second step of the U-OCBE is the measurement update, which is the same as the UKF measurement step. Note that \mathbf{y}_k is the measurement at time t_k , \mathbf{h} is the nonlinear measurement

function, and \mathbf{R}_k is the measurements covariance matrix. The full U-OCBE measurement update is:

$$\bar{\boldsymbol{\chi}}_{k|k-1}^{0:2n} = S(\bar{\boldsymbol{x}}_{k|k-1}, \mathbf{P}_{k|k-1}) \quad (\text{B.6})$$

$$\bar{\boldsymbol{\gamma}}_k^{0:2n} = \mathbf{h}(\bar{\boldsymbol{\chi}}_{k|k-1}^{0:2n}) \quad (\text{B.7})$$

$$\bar{\mathbf{y}}_k, \mathbf{P}_{yy} = S^{-1}(\bar{\boldsymbol{\gamma}}_k^{0:2n}, \mathbf{R}_k) \quad (\text{B.8})$$

$$\mathbf{K} = \mathbf{P}_{xy} \mathbf{P}_{yy}^{-1} \quad (\text{B.9})$$

$$\bar{\boldsymbol{x}}_{k|k} = \bar{\boldsymbol{x}}_{k|k-1} + \mathbf{K}(\mathbf{y}_k - \bar{\mathbf{y}}_k) \quad (\text{B.10})$$

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \mathbf{K} \mathbf{P}_{yy} \mathbf{K}^T \quad (\text{B.11})$$

Finally, the third step is the costate update which utilizes the single-step smoothing properties of the OCBE. We only use a linear smoother because the costate estimate is not integral to filter performance and it was found to be sufficient for control estimation. The costate update is:

$$\bar{\boldsymbol{x}}_{k-1|k} = \bar{\boldsymbol{x}}_{k-1|k-1} + \mathbf{P}_{k-1|k-1} \boldsymbol{\Phi}_{xx}^T \mathbf{P}_{k|k-1}^{-1} (\bar{\boldsymbol{x}}_{k|k} - \bar{\boldsymbol{x}}_{k|k-1}) \quad (\text{B.12})$$

$$\bar{\boldsymbol{p}}_{k-1|k} = -\mathbf{P}_{k-1|k-1}^{-1} (\bar{\boldsymbol{x}}_{k-1|k} - \bar{\boldsymbol{x}}_{k-1|k-1}) \quad (\text{B.13})$$

$$\bar{\mathbf{u}}(\tau|\tau_k) = -\mathbf{Q}_{\text{ce}}(\tau) \mathbf{B}_{\text{ce}}^T \boldsymbol{\Phi}_{xx}(\tau, \tau_{k-1}) \bar{\boldsymbol{p}}_{k-1|k} \quad (\text{B.14})$$

In the costate update equations $\bar{\boldsymbol{p}}$ is the costate vector. Note, that the control estimate $\bar{\mathbf{u}}(t|t_k)$ is continuous between the measurement epochs and is only applicable over that time. To get improved control estimates at previous times a full smoother must be applied, which can be implemented from equations in Reference [19].

B.2 Square Root Information Modification

While the U-OCBE is powerful, it is more computationally expensive than a linearized filter. Thus, when a linear filter is sufficient then the U-OCBE is not necessary warranted. Unfortunately, linearized filter can still struggle with numerics and therefore square-root information updates from the square-root information filter (SRIF) may alleviate these issues. The square-root information

implementation of the OCBE, now referred to as the SRI-OCBE, takes advantage of the same properties that were used to implement the U-OCBE and leads to a three step filter: time update, measurement update, and control update. The base SRIF updates are from Statistical Orbit Determination by Tapley, Schutz, and Born [90]. Their implementation of the SRIF utilizes orthogonal transformations which are represented by the matrix \mathbf{T} .

The SRI-OCBE requires transforming the measurement uncertainty, information matrix, and state into SRI space. These transformations are respectively:

$$\mathbf{R} = \mathbf{S}_y^{-1} \mathbf{S}_y^{-T} \quad (\text{B.15})$$

$$\mathbf{P} = \mathbf{S}_x^{-1} \mathbf{S}_x^{-T} \quad (\text{B.16})$$

$$\mathbf{S}_x \delta \bar{\mathbf{x}} = \mathbf{b}_x \quad (\text{B.17})$$

The first step of the SRI-OCBE is the time update which effectively identical to the SRIF time update, except that the control noise term is now determined by the OCBE STM which dynamically propagates the estimated control noise. The control noise, commonly denoted as \mathbf{Q} and its SRI form as \mathbf{S}_q , is given by the OCBE STM terms:

$$\mathbf{S}_{q,k-1|k-1}^{-1} \mathbf{S}_{q,k-1|k-1}^{-T} = -\mathbf{\Phi}_{xp} \mathbf{\Phi}_{xx}^T \quad (\text{B.18})$$

With the OCBE noise in hand the standard SRIF time update is applied:

$$\mathbf{T} \begin{bmatrix} \mathbf{S}_{q,k-1|k-1} & \mathbf{0} & \mathbf{0} \\ -\mathbf{S}_{x,k-1|k-1} \mathbf{\Phi}_{xx}^{-1} & \mathbf{S}_{x,k-1|k-1} \mathbf{\Phi}_{xx}^{-1} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{S}_{q,k|k-1} & \mathbf{S}_{qx,k|k-1} & \mathbf{b}_{q,k|k-1} \\ \mathbf{0} & \mathbf{S}_{x,k|k-1} & \mathbf{b}_{x,k|k-1} \end{bmatrix} \quad (\text{B.19})$$

The second step of the SRI-OCBE is the measurement update, which is identical to its SRIF counterpart. The full U-OCBE measurement update is:

$$\mathbf{T} \begin{bmatrix} \mathbf{S}_{x,k|k-1} & \mathbf{b}_{x,k|k-1} \\ \mathbf{S}_y^{-1} \mathbf{H}_k & \mathbf{S}_y^{-1} \delta \mathbf{y}_k \end{bmatrix} = \begin{bmatrix} \mathbf{S}_{x,k|k} & \mathbf{b}_{x,k|k} \\ \mathbf{0} & \mathbf{e}_{x,k} \end{bmatrix} \quad (\text{B.20})$$

The last step is the costate update which utilizes the single-step smoothing property of the OCBE. Therefore, we apply a single step SRIF smoother and then use that to update the costate.

This is implemented as:

$$\mathbf{T} \begin{bmatrix} \mathbf{S}_{q,k|k-1} + \mathbf{S}_{qx,k|k-1} & \mathbf{S}_{qx,k|k-1} \Phi_{xx} & \mathbf{b}_{q,k|k-1} \\ \mathbf{S}_{x,k|k} & \mathbf{S}_{x,k|k} \Phi_{xx} & \mathbf{b}_{x,k|k} \end{bmatrix} = \begin{bmatrix} \mathbf{S}_{q,k-1|k} & \mathbf{S}_{qx,k-1|k} & \mathbf{b}_{q,k-1|k} \\ \mathbf{0} & \mathbf{S}_{x,k-1|k} & \mathbf{b}_{x,k-1|k} \end{bmatrix} \quad (\text{B.21})$$

$$\delta \bar{\mathbf{p}}_{k-1|k} = -\mathbf{P}_{k-1|k-1}^{-1} (\delta \bar{\mathbf{x}}_{k-1|k} - \delta \bar{\mathbf{x}}_{k-1|k-1}) \quad (\text{B.22})$$

$$\bar{\mathbf{u}}(t) = -\mathbf{Q}_{\text{ce}} \mathbf{B}_{\text{ce}}^T \Phi_{pp}(\tau, \tau_{k-1}) \delta \bar{\mathbf{p}}_{k-1|k} \quad (\text{B.23})$$

Once the forwards filter is complete, it is often desirable to smooth the solution to return a better estimate for the entire trajectory. Given information up to time t_l the smoothing process is performed by iteratively applying Equations B.24-B.26.

$$\mathbf{T} \begin{bmatrix} \mathbf{S}_{q,k|k-1} + \mathbf{S}_{qx,k|k-1} & \mathbf{S}_{qx,k|k-1} \Phi_{xx} & \mathbf{b}_{q,k|k-1} \\ \mathbf{S}_{x,k|l} & \mathbf{S}_{x,k|l} \Phi_{xx} & \mathbf{b}_{x,k|l} \end{bmatrix} = \begin{bmatrix} \mathbf{S}_{q,k-1|l} & \mathbf{S}_{qx,k-1|l} & \mathbf{b}_{q,k-1|l} \\ \mathbf{0} & \mathbf{S}_{x,k-1|l} & \mathbf{b}_{x,k-1|l} \end{bmatrix} \quad (\text{B.24})$$

$$\delta \bar{\mathbf{p}}_{k-1|l} = -\mathbf{P}_{n-1|n-1}^{-1} (\delta \bar{\mathbf{x}}_{k-1|l} - \delta \bar{\mathbf{x}}_{n-1|n-1}) \quad (\text{B.25})$$

$$\bar{\mathbf{u}}(t) = -\mathbf{Q}_{\text{ce}} \mathbf{B}_{\text{ce}}^T \Phi_{pp}(\tau, \tau_{n-1}) \delta \bar{\mathbf{p}}_{k-1|l} \quad (\text{B.26})$$

Appendix C

Space-Based Observation of Cislunar Vehicles with the Optimal Control Based Estimator

C.1 Introduction

There is significant international interest to begin long term operations in cislunar space for future exploration and scientific development purposes. The Artemis plan is exemplary model for the growing interest in this region, as it is a multinational program to develop an outpost and highway for deep space exploration through cislunar space [63]. This vision suggests that there will be a myriad of spacecraft flying through this zone; everything from crewed vehicles to robotic explorers. The burgeoning focus on this region of space will lead to a surge in traffic, and potential, debris in the region. In order to ensure long-term use and safety for all cislunar vehicles the occupation of this regions will need to be monitored carefully.

Traditional estimation schemes for cislunar spacecraft are predicated on Earth-based radar such as the Deep Space Network (DSN). The DSN has persisted since the Apollo era and it is still a primary source of high accuracy information for current studies and missions [15, 16, 18, 91]. While Earth-based radar provides accurate solutions for cooperative vehicles, it also leads to several issues for persistent observation of all vehicles. First, scheduling DSN time is difficult because it is already heavily subscribed. Alleviating DSN resources has become important enough that it has resulted in various onboard autonomous navigation studies for vehicles planning to use the network [33, 92]. Second, the DSN is not optimized for non-cooperative vehicles and therefore cannot provide the same level of accuracy in those cases. Third, taking measurement from Earth

severely limits viewing geometry of cislunar orbits and leads to inflated states estimates along the Earth-Moon line.

To alleviate Earth-based resources, cislunar optical observation platforms are examined as an alternative measurement source. Cislunar optical observers are an appealing solution because they utilize existing technology and generally yield an improved viewing geometry over Earth-based sensors [65, 44]. Optical measurements also provide a key benefit by supporting the same degree of accuracy for cooperative and non-cooperative systems. Ideally only a single observer placed near the L2 equilibrium point would be necessary to enable robust estimation, however a variety of observer locations can be explored if a single observer is not sufficient. Thus, this work explores cislunar optical observers to establish a feasible strategy that achieves both persistent and accurate state estimation.

It is important to understand that having deterministic observers does not guarantee an accurate state solution since target vehicles will experience strongly nonlinear motion and may pass through regions of chaotic motion. Nonlinearities and chaotic dynamics means that small perturbations and estimation inaccuracies lead to large predictive errors. These in turn may cause filter divergence if insufficient information is ingested into the estimation algorithm. Beyond filtering issues, these chaotic dynamics also underpin the necessity of a station keeping scheme for spacecraft to maintain its orbit [15, 16, 18]. If station keeping and other events are not accounted for, then the spacecraft will deviate from its nominal trajectory and the track may be lost. This necessitates some form of maneuver compensation, or ideally, detection for long-term estimation purposes.

The objective of this appendix is to examine cislunar tracking and maneuver detection with a known observer, making this a predecessor to Chapter 5's results. This appendix begins by detailing the modeling choices and assumptions for the subsequent simulations which covers: reference orbits, target maneuvers, and measurements. Next, the state estimation strategies are assessed to determine necessary design components for functional filter solutions. Finally, the maneuver detection is carried out by performing hypothesis tests on estimated control profiles from the optimal control based estimator (OCBE). The chapter concludes by summarizing the presented results and

discussing possible future work.

C.2 Problem Formulation

This work formulates space based cislunar observation by assuming known observers and unknown targets as a precursor to SSAT which is examined in the Chapter 5. By starting with a deterministic observer, it is easier to develop observational schemes without the worry of observability issues. Therefore, only the target spacecraft state is estimated. The dynamical model used propagate state motion in cislunar space is approximated with the CR3BP which is covered in Section 2.1.2. Several observer and target locations are then selected in the CR3BP model for results and simulations purposes. Subsequently, maneuver are included as instantaneous impulses with two levels of execution errors. Finally, the optical measurements are introduced based on the equations form Section 2.2 along with modifications necessary for this chapter alone.

C.2.1 Selected Orbits

There are 2 southern L2 NRHO reference orbits of interest; an approximate 9:2 resonance orbit which NASA is targeting for its Lunar Gateway and a stable NRHO with a stability index of 1. Both NRHOs are found using a standard single shooting algorithm with initial conditions at apoapsis. Because the initial condition is at apoapsis the initial y position, x velocity, and z velocity are 0. The details of these orbits including the stability index, radius of periapsis, and non-dimensional non-zero initial conditions are given in Table C.1.


Table C.1: Reference orbit details.

| Parameter | 9:2 Synodic NRHO | Stable NRHO |
|--------------------------|------------------|-------------|
| stability index | -1.32 | 1.0 |
| radius of periapsis (km) | 3,236.3 | 16,428.0 |
| initial r_i | 1.0219 | 1.0796 |
| initial \dot{r}_j | -0.10309 | -0.19739 |
| initial r_k | -0.18206 | -0.20237 |

The 2 reference orbits, as well as other family members, are displayed in Figure C.1 with

respect to the moon. The 9:2 resonance orbit is highlighted in blue and the stable orbit is highlighted in red. Note the axes in Figure C.1 are centered at the Moon's origin and not the CR3BP origin.

Figure C.1: Southern L2 NRHO family, with the 2 reference orbits highlighted in blue and red.



C.2.2 Target Maneuvers

The station keeping model used here is based on previous NRHO station keeping work [15]. The proposed station keeping model functions by burning at apoapsis nearly every orbit. Due to the small station keeping costs it may be impractical, or impossible, to execute the minor burns every revolution. The station keeping costs vary with navigation errors but are capped by 5 m/s impulses and are often on the order of 10's of mm/s. To approximate these station keeping maneuvers for uncertainty quantification purposes a randomly pointing impulse is added at apoapsis with a magnitude given by a 3-sigma truncated Gaussian. This policy is formally described by Equations C.1-C.3.

$$u_\theta = \mathcal{U}[-\pi, \pi] \quad (\text{C.1})$$

$$u_\psi = \mathcal{U}[-\pi/2, \pi/2] \quad (\text{C.2})$$

$$u_{mag} = \mathcal{N}(\mu_u, \sigma_u, \mu_u - 3\sigma_u, \mu_u + 3\sigma_u) \quad (\text{C.3})$$

Within the station keeping framework there are 2 policies that are examined. First, a noisy policy with a mean of 1 m/s and standard deviation of 0.3 m/s. Second, a quiet policy with a mean of 50 mm/s and a 15 mm/s standard deviation. This information is summarized in Table C.2. Since station keeping maneuvers are generally pointed along the stable manifold this policy encapsulates additional maneuver possibilities since the stable manifold direction is a subset of the fully random pointing. Ensuring the observational system can handle the boarder set includes additional robustness for state estimation purposes.

Table C.2: Station keeping policy statistics.

| Parameter | Noisy Policy | Quiet Policy |
|-------------------|--------------|--------------|
| μ_u (mm/s) | 1000 | 50 |
| σ_u (mm/s) | 300 | 15 |

C.2.3 Measurements

The optical measurements are commonly represented by azimuth and elevation angles provided by the observer. It is assumed that these angles are given in the CR3BP frame when being sent to the filter. The equations for azimuth and elevation are given in Equations C.5 and C.6 respectively. Note that these equations are only a function of the relative position vector from the observer to the object being tracked, which is defined by Equation C.4.

$$\boldsymbol{\rho} = \mathbf{r} - \mathbf{r}_{obs} \quad (\text{C.4})$$

$$\theta = \tan^{-1}(\rho_j, \rho_i) \quad (\text{C.5})$$

$$\psi = \sin^{-1}\left(\frac{\rho_k}{\|\boldsymbol{\rho}\|}\right) \quad (\text{C.6})$$

The rate of change of these angles are also introduced as measurements for filtering purposes and these rates are analytically derived in Equations C.8 and C.9. These angular rates contain velocity information which is useful since orbits in the CR3BP, particularly NRHOs, are highly

sensitive to velocity deviations. Specifically, constraining velocity information around perihelion of NRHOs is critical to maintaining custody of targets [15, 18].

$$\dot{\boldsymbol{\rho}} = \dot{\mathbf{r}} - \dot{\mathbf{r}}_{obs} \quad (\text{C.7})$$

$$\begin{aligned} \dot{\theta} &= \frac{1}{(\rho_j/\rho_i)^2 + 1} \left(\frac{\dot{\rho}_j}{\rho_i} - \frac{\rho_j \dot{\rho}_i}{\rho_i^2} \right) \\ &= \frac{\rho_i \dot{\rho}_j - \rho_j \dot{\rho}_i}{\rho_j^2 + \rho_i^2} \end{aligned} \quad (\text{C.8})$$

$$\begin{aligned} \dot{\psi} &= \frac{1}{(1 - (\rho_k/|\boldsymbol{\rho}|)^2)^{1/2}} \left(\frac{\dot{\rho}_k}{|\boldsymbol{\rho}|} - \frac{\rho_k |\dot{\boldsymbol{\rho}}|}{|\boldsymbol{\rho}|^2} \right) \\ &= \frac{1}{(1 - (\rho_k/|\boldsymbol{\rho}|)^2)^{1/2}} \left(\frac{\dot{\rho}_k}{|\boldsymbol{\rho}|} - \frac{\rho_k (\boldsymbol{\rho} \cdot \dot{\boldsymbol{\rho}})/|\boldsymbol{\rho}|}{|\boldsymbol{\rho}|^2} \right) \\ &= \frac{\dot{\rho}_k |\boldsymbol{\rho}|^2 - \rho_k (\boldsymbol{\rho} \cdot \dot{\boldsymbol{\rho}})}{|\boldsymbol{\rho}|^2 (|\boldsymbol{\rho}|^2 - \rho_k^2)^{1/2}} \end{aligned} \quad (\text{C.9})$$

Angular rate information can be provided by differencing subsequent optical observations, or if the exposure times are long enough, a light streak. The simplest way to incorporate angular rate information would be to apply a numerical differentiation scheme, such as the midpoint rule, over several successive observations. Given that the measurements are independent identically distributed zero mean Gaussian measurements, then the noise of numerical derivative calculated using Equation C.10 is derived following Equation C.11. This derivation holds for σ_ψ^2 as well.

$$\dot{\theta} = \frac{\theta_2 - \theta_1}{t_2 - t_1} \quad (\text{C.10})$$

$$\begin{aligned} \sigma_{\dot{\theta}}^2 &= \frac{1}{\Delta t^2} (\sigma_{\theta}^2 + \sigma_{\theta}^2) \\ &= \frac{2}{\Delta t^2} \sigma_{\theta}^2 \end{aligned} \quad (\text{C.11})$$

For the purpose of this paper it is assumed that numerical errors associated with calculating the angular rate are negligible in comparison to the angular measurement noise, and therefore the angular rate noise is only a product of the angular measurements used for the angular rate calculation as shown in Equation C.11. The angular rate measurements are assumed to be taken 1 second

apart to conservatively preserve a higher noise estimate. These measurements are also assumed to be processed outside the state estimator and made available to the filter with no correlation to the incoming angular measurements. In a higher fidelity case with real data correlation could be added in the measurement uncertainty matrix.

Note, the presence of angular rate measurements is not meant to add significant information to the system. This is reflected by the fact that chosen rate measurement uncertainty is roughly equivalent to angular measurements. The purpose of explicitly adding the rates as measurements is to stabilize filter updates. It is known that linearized filters struggle with the strongly non-linear dynamics near periapsis and can cause filter failure [15, 18]. The strongly non-linear dynamics cause time updates to produce poor covariance propagation estimates, which in turn causes filter failure. Adding the rate as an explicit measurement forces the filter back to the true solution after the poor time update.

The measurement noise for all following simulations is shown in Table C.3. Note that the angular uncertainty is chosen to match previous studies that were based on current optical equipment aimed at autonomous optical navigation, which is likely conservative compared to a dedicated optical observation system [33].


Table C.3: Measurement noise statistics.

| Parameter | Value |
|---|--------------|
| σ_θ (micro-radians) | 10 |
| σ_ψ (micro-radians) | 10 |
| $\sigma_{\dot{\theta}}$ (micro-radians/s) | $10\sqrt{2}$ |
| $\sigma_{\dot{\psi}}$ (micro-radians/s) | $10\sqrt{2}$ |

The measurements are taken from a host of potential observer locations in the lunar vicinity. A natural placement for a cislunar observer would be in the vicinity of the L2 equilibrium point for its uninterrupted view of L2 NRHOs. Other attractive observer locations include the lunar surface

or on a Distant Retrograde Orbit (DRO). The spacecraft observers, such as the DRO vehicle, include lunar occultations. It was found that while the L2 observers do not experience line-of-sight interruptions, the DRO experiences limited outages that are infrequent and short. The lunar surface observatories are limited by their horizon and therefore require 3 observatories to view the full orbit 9:2 orbit due to the low periapsis coming so close to the lunar surface. All observer locations and the NRHOs are displayed with respect to the Moon's center in Figure C.2. It is assumed that the state of these observers is known well enough to be considered deterministic for simulation purposes.

Figure C.2: Observer locations in red, and NRHOs in blue.



C.3 Optical-Only Tracking with a Known Observer in Cislunar Space

All simulations initialize the OCBE with a random initial error and fixed uncertainty dictated by Earth-radar observing in order to approximate hand-off from Earth-based resources to an optical observer. To be approximately consistent with previous navigation studies the 3-sigma position uncertainty is 10 km and the 3-sigma velocity uncertainty is 10 cm/s [93, 92]. With the filter selected and initialization specified, modifications to the ballistic OCBE are now presented to aid state estimation in the CR3BP. While this section focuses on the modification the full algorithm

is given in the appendix for reference. Subsequently, an effective observation scheme is sought to provide consistent and robust estimation given no mismodeling. Finally, a filter tuning strategy to handle maneuvers is presented and its expected results are shown.

C.3.1 Observational Analysis

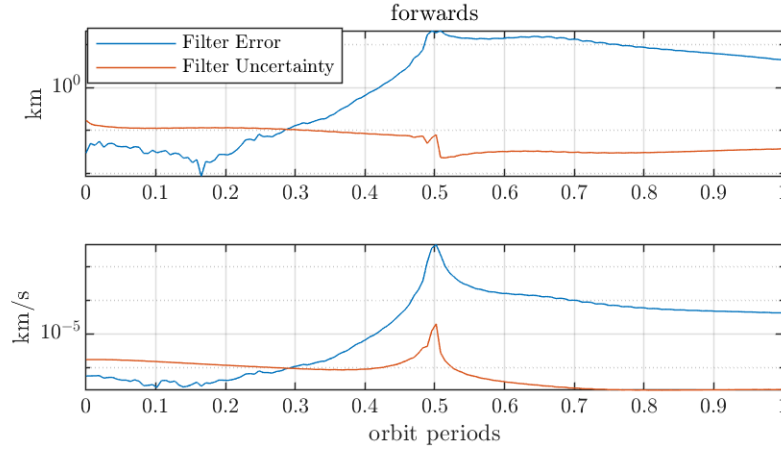
The objective of the observation strategy analysis is to determine the requirements for an optical measurement scheme to provide robust and consistent state estimation for the given reference orbits using the linearized filters. To begin, a simulation with no mismodeling or maneuvers is instantiated to determine if a measurement scheme provides feasible estimation. These simulations are started at apoapsis to allow the filter to decrease its uncertainty before periapsis. Decreasing uncertainty prior to periapsis is desirable because it is the region most sensitive to perturbations and therefore the most likely location for filter divergence.

First, a simple angle only measurement approach is examined. Given a single L2 observer and varied measurement frequencies ranging from 1 minute to 12 hours, every strategy leads to filter divergence before periapsis for both reference orbits. In an attempt to constrain the system with additional geometric information the L2 observer was supplemented with a DRO observer and 3 lunar stations. Even with this additional information the unstable and 9:2 NRHO filter nearly always fail. In an all-out attempt to enable filter convergence, both the measurement accuracy was increased and initial covariance reduced by an order of magnitude. The filter estimate still diverged from the true solution.

An example of the 9:2 NRHO filter divergence is shown in Figure C.3. This example includes all 5 observers taking measurements every hour, as well as reduced measurements and initial uncertainty by an order of magnitude. The figures display the true error in blue and the 3-sigma filter uncertainty in red. The plots are also split into position and velocity space by taking the norm error and uncertainty respectively, with the top being position and bottom being velocity. A filter that produces a consistent solution has the true error bounded below the filter uncertainty. Normally the smoothed solution is also shown, but since the forwards filter fails the smoothed solution is

completely erroneous.

Figure C.3: Typical 9:2 NRHO filter divergence with angle only measurements from L2, DRO, and lunar surface observers. This particular example has improved measurements and initial uncertainty, and the filter estimate still diverged from the true solution.



Next, angle and angle rate measurements are examined. This time a single L2 observer with measurement frequencies from 1 minute to 12 hours lead to consistent filtering solutions in every scenario. As expected higher measurement frequencies lead to greater filter certainty over a single orbit. Final uncertainty levels given various measurement rates on both reference orbits are given in Table C.4.

As anticipated, the stable NRHO filter produces qualitatively similar but quantitatively better results than the 9:2 NRHO filter. All rates produce consistent filters that properly track the true state so the measurement rate can be used to set the desired filter accuracy. A typical successful 9:2 NRHO filter result is plotted in Figure C.4. Since only 1 observer is necessary to produce a working system the other observers are ignored.

These results exemplify the fact that severe non-linearity needs to be addressed for linearized filtering in the CR3BP. This is most notably apparent by the fact taking angular measurements at a high enough frequency should lead to equivalent information ingested into the filter as taking angular rate measurements, but as found in the results it still led to filter failure. This is because the issue is not insufficient information content. The issue is the fact that the 1st-order linearized filter

Table C.4: Idealistic final filter covariance levels for various measurement rates given no mismodeling.

| Measurement Rate | 9:2 NRHO | | Stable NRHO | |
|------------------|-----------------------|--------------------------|-----------------------|--------------------------|
| | Position σ (m) | Velocity σ (mm/s) | Position σ (m) | Velocity σ (mm/s) |
| 1 min | 30 | 0.2 | 20 | 0.1 |
| 30 min | 160 | 0.8 | 100 | 0.4 |
| 1 hr | 230 | 1.1 | 140 | 0.6 |
| 2 hr | 320 | 1.5 | 200 | 0.9 |
| 3 hr | 390 | 1.9 | 240 | 1.1 |
| 6 hr | 540 | 2.6 | 340 | 1.5 |
| 12 hr | 740 | 3.7 | 460 | 2.0 |

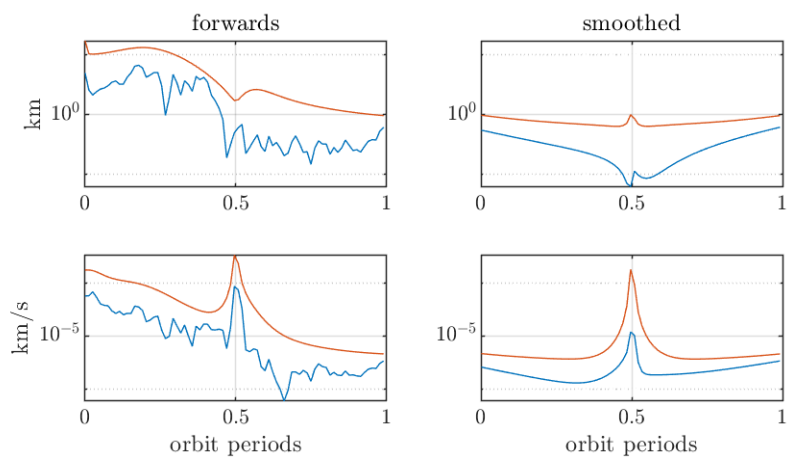
cannot accurately represent the strongly non-linear dynamics. Adding angular rate measurements assuages this issue enough to produce working linearized filters. This negates the need for higher order filters and maintains the speed and simplicity of a linearized filter.

Given these results, the chosen observation strategy used for the rest of the paper is a single L2 observer with angle and angular rate measurements taken every 2 hours. This rate is chosen as a low demand requirement that balances filter accuracy and algorithm run-time, and appears to be a sweet spot for the 9:2 NRHO. With the observation strategy outlined the filter method of maneuver processing can now be addressed.

C.3.2 Filter Tuning Strategy

The OCBE handles mismodeled dynamics by including a fictitious control policy that absorbs errors and represents the errors as an optimal control policy. This control policy is similar to a process noise term and needs to be tuned in a similar manner. A primary difference between the control policy and process noise is that the control policy is constrained to follow system dynamics and is calculated via the STM. While this removes the need to explicitly define a process noise term, it does not remove the tuning process since the control uncertainty term determines the OCBE performance.


Figure C.4: A typical working 9:2 NRHO filter with angular and angular rate measurements with measurements every 2 hours from a single L2 observer. The forwards filter is on the left and the smoothed filter is on the right.



C.3.2.1 Adaptive OCBE Tuning


When picking a tuning method for the OCBE there are several options. The most attractive option is the entirely autonomous adaptive OCBE algorithm, which identifies and tunes filter uncertainty automatically. The principle behind the adaptive OCBE is that it defines a statistical metric derived from the estimated control, measurement, and a priori uncertainty to determine if a maneuver occurs. If the metric remains above a probable limit, set to 99% for this paper, for 3 consecutive epochs then a maneuver is detected at the first epoch. If a maneuver is detected, then the filter solves for a state covariance that sets the metric at that epoch equal to its mean value. This algorithm produces near optimal filtering results for impulsive system with large maneuvers as it automatically adjusts the control uncertainty to reflect the true dynamics.

Figure C.5: Adaptive tuning state estimate performance given a 3 m/s impulse at apoapsis. The filter automatically increases its uncertainty during the maneuver in the forwards process, which is at the 1/2 orbit period, to maintain a correct filter solution.



To demonstrate the adaptive OCBE performance on a sufficiently larger maneuver a simulation was created with a 3 m/s maneuver at apoapsis on the 9:2 NRHO reference orbit. The simulation is started at periapsis to allow the filter to reduce its uncertainty before the station keeping maneuver. All other simulation parameters remain the same as previously defined. The state estimation results for the forwards and smoothing process are displayed in Figure C.5. The

Figure C.6: Filter control estimates are plotted on top, and OCBE metrics on the bottom for the adaptively tuned system. Again, the forwards filter is on the left and the smoothed filter is on the right.



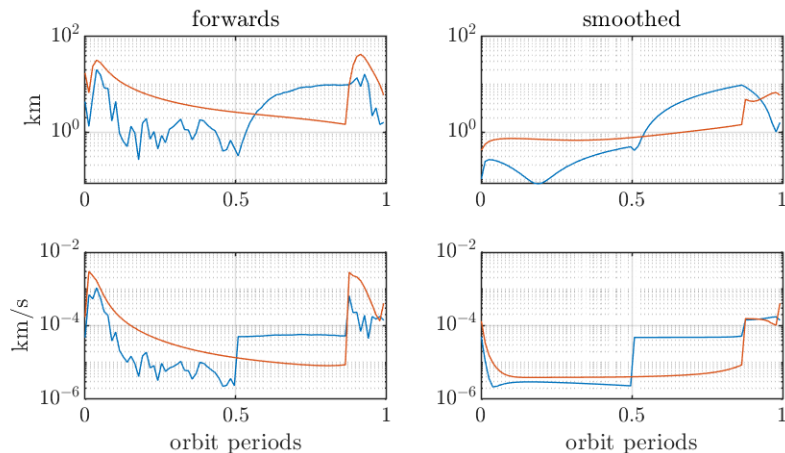
norm of the estimated control policy and the normalized OCBE metric are displayed in Figure C.6. The normalized OCBE metric is the OCBE metric divided by the maneuver detection limit. The OCBE metric figures plot the calculated metric in blue, and the statistical metric mean in red.

The filter was able to autonomously identify the maneuver and adjust its uncertainty to reflect the control uncertainty necessary to reduce the OCBE metric. This appears as an instantaneous increase in filter uncertainty in the forwards solution, and spike in control estimates in the forwards and smoothed solutions. It also appears as an outlier in the forwards metrics, which is subsequently removed in the smoothed metrics. This is the preferred tuning method given maneuvers over 1 m/s, but unfortunately this method doesn't perform as well for the smaller station keeping maneuvers.

C.3.2.2 Constant OCBE Tuning


As the maneuver magnitude decreases the detection becomes delayed as it takes time for the information to transition into the metric space. This is demonstrated with a simulation that replaced the 3 m/s with a 50 mm/s maneuver. Filter results given this reduced maneuver are displayed in Figures C.7 and C.8.

Figure C.7: Filter performance for adaptively tuned system given a 50 mm/s impulse at apoapsis. The filter doesn't catch the deviation until nearly periapsis where it adjusts its uncertainty to reconverge to the true solution.



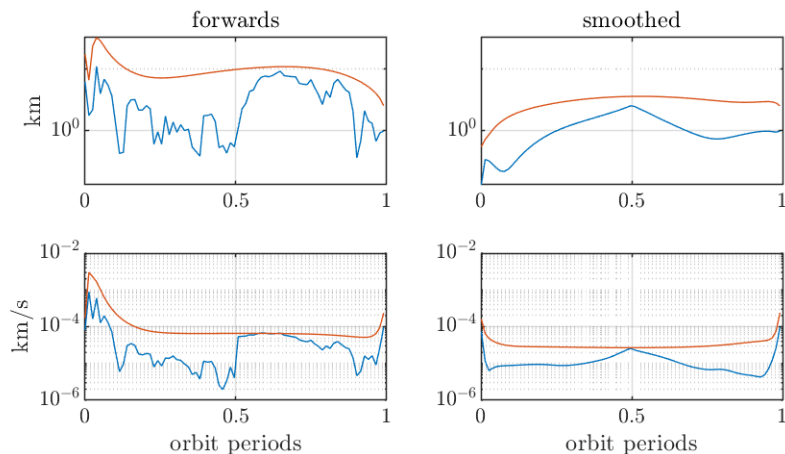
The filter is unable to immediately detect the small change in trajectory, and error builds as

Figure C.8: Filter control estimates and metrics for adaptive system given a 50 mm/s maneuver which goes undetected for a long time.



the true system drifts from the estimate. Eventually the error builds to a point where the filter can detect a modeling error and makes a significant correction to its uncertainty bringing the filter back to the true solution. This does fix the filter for the later epochs but leaves a large portion of the estimated trajectory to be erroneous. If the user only cares about a forwards solution at the final time this may be acceptable, but if navigation during that arc was necessary then this leaves much to be desired.

Figure C.9: Filter performance with increased control uncertainty given a 50 mm/s maneuver. Both the forwards and smoothed filters produce consistent, but less certain, solutions.



Instead of adaptively tuning the filter the system can simply contain a larger control uncertainty that would account for the maneuver. This produces a filter with a greater overall uncertainty but also results in a consistent state estimate for the entire observation window. An example of taking this approach is displayed in Figures C.9 and C.10.

Given this approach the event is obscured in the forwards solution, but it is still visible as a control effort spike in the smoothed solution. This control spike does not appear if there is no maneuver, and thus control policy can be used to distinguish maneuvers. This is the basis for the maneuver detection methodology proposed in the next section.




Figure C.10: Filter control estimates and metrics with increased control uncertainty given a 50 mm/s maneuver. The maneuver appears as a spike in estimated control only in the smoothed system.

C.3.2.3 Proposed OCBE Tuning

For this study the filter mismodeling is limited to the station keeping maneuvers, and thus it makes sense to limit the control uncertainty to around apoapsis. Therefore, the final tuning strategy is to increase the control uncertainty for an 8 hour window around apoapsis to account for the unknown station keeping maneuver. A typical filter result given this approach is displayed in Figures C.11 and C.12. This produces a consistent filter with a smoothed control spike around apoapsis that is indicative of a maneuver.

This limited control uncertainty approach exemplifies the estimation strategy for all future simulations. An Earth-based radar estimate is handed off to the optical observer at periapsis, where the system is observed for a single orbit period by a L2 observer with angular and angular-rate measurements. A measurement cadence of 1 observation per 2 hours is adopted, and control uncertainty is increased over an 8 hour window at apoapsis to account for the expected station keeping maneuver. An observer can use this strategy to maintain custody of the craft for an indefinite time as it performs station keeping maneuvers on a NRHO.

Figure C.11: Filter performance for appropriately tuned system given a small maneuver which still performs well and identifies maneuver.

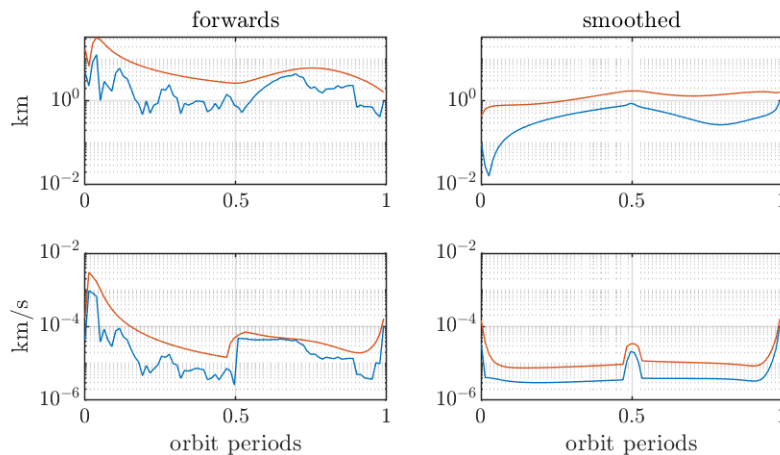
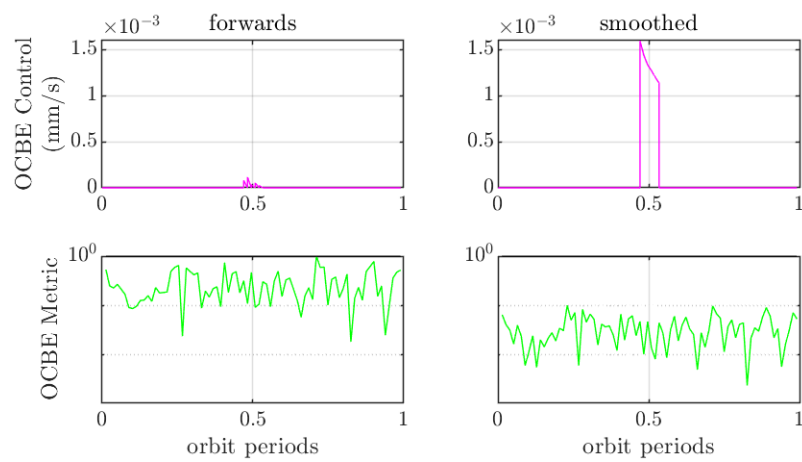


Figure C.12: Filter control estimates and metrics for appropriately tuned system given a small maneuver which still performs well and identifies maneuver.



C.4 Maneuver Detection

The methodology proposed for maneuver detection is a binary hypothesis test on the integral of OCBE control policy, which represents the total estimated control on the system. This approach is chosen because the sensitivity to accelerations of the control estimate makes it an ideal indicator of mismodeling, while the OCBE's filtering properties act to minimize the effects of noisy measurements and initial errors to provide low variance results. The overall process for maneuver detection is as follows: apply the OCBE to the estimation problem, integrate the estimated control policy from the OCBE, and perform a hypothesis test to determine if a maneuver has occurred.

C.4.1 Hypothesis Testing Using Optimal Control Policies

The null hypothesis (\mathcal{H}_0) is that the observed system did not maneuver, and the alternative (\mathcal{H}_1) is that it did maneuver. The test statistic, the integral of the OCBE control profile, is a random value denoted by \mathcal{Z} . A realization of the random variable, z , is calculated from the OCBE control policy by Equation C.12.

$$z(u) = \int_{t_0}^{t_f} \|\bar{\mathbf{u}}(\tau)\| \partial\tau \quad (\text{C.12})$$

There are many approaches to binary hypothesis tests, but the approach taken here is maximum likelihood decision rule. This rule is a conservative approach that assumes an equivalent cost for incorrect decisions and equally likely a-priori hypothesis probabilities. Taking this approach yields the decision function Equation C.13. The derivation for this approach is found in many texts, but it is cleanly derived from a Bayesian minimum error sense in [94].

$$D(z) = \begin{cases} \mathcal{H}_0 & \text{if } PDF_{\mathcal{H}_0}(z) \geq PDF_{\mathcal{H}_1}(z) \\ \mathcal{H}_1 & \text{otherwise} \end{cases} \quad (\text{C.13})$$

In order to perform this hypothesis test the distributions for \mathcal{Z} need to be determined or numerically approximated via an uncertainty quantification method. The uncertainty quantification

method of choice for this paper is a Monte-Carlo analysis. Beyond being necessary for the hypothesis test, approximating these distributions also allows for a better understanding of OCBE control policy which is often non-intuitive.

C.4.2 Control Policy Distributions

The test statistic \mathcal{Z} is random variable dependent on the stochastic inputs to the filter which are: a station keeping maneuver, measurements, and initial filter error. For the case where the system does not maneuver then it is only dependent on the measurements and initial error. Thus, when numerically approximating the non-maneuvering PDF the input dimensionality is lower and requires less samples to return the same level of accuracy. As a result, the non-maneuvering Monte-Carlo analysis only contains 500 sample simulations, while maneuvering case contains 3000 sample simulations.

Beginning with the active station keeping policy with a mean maneuver size of 1 m/s and standard deviation of 0.3 m/s, it is clear that non-maneuvering systems produce statistically smaller control policies than maneuvering systems. The histograms of the control policy integrals for the 9:2 NRHO Monte-Carlo are displayed in Figure C.13. PDFs for these distributions are numerically calculated and plotted in Figure C.14. The maneuvering system appears to have an approximately Gaussian shape which is reassuring since the station keeping magnitude comes from a truncated Gaussian.

Figure C.13: Histogram of Monte-Carlo results for 1 m/s mean non-maneuvering and maneuvering case.





Figure C.14: Numerical PDF of Monte-Carlo results for 1 m/s mean non-maneuvering and maneuvering case.



From examining these figures it is clear that $PDF_{\mathcal{H}_0}(z)$ only crosses with $PDF_{\mathcal{H}_1}(z)$ once. Therefore, the maximum likelihood rule can be simplified to Equation C.14. For the active case presented here $z_{lim} = 0.4106$ m/s. These PDFs and z_{lim} values are used to calculate the probability that \mathcal{H}_0 is incorrectly rejected or that \mathcal{H}_0 is incorrectly accepted. This is done by integrating the PDF of \mathcal{H}_0 above z_{lim} , and \mathcal{H}_1 below z_{lim} . These rates are more commonly called false positive and missed detection rates respectively. The expected hypothesis test results for the active policy are given in Table C.5.

$$D(z) = \begin{cases} \mathcal{H}_0 & \text{if } z(u) \leq z_{lim} \\ \mathcal{H}_1 & \text{otherwise} \end{cases} \quad (\text{C.14})$$


Table C.5: Expected 9:2 NRHO maneuver detection probability for the active station keeping policy based on Monte-Carlo results.

| | No Maneuver | Maneuver |
|------------------------|-------------|----------|
| No Maneuver Identified | 0.9913 | 0.0123 |
| Maneuver Identified | 0.0087 | 0.9877 |

To further investigate the control policy's sensitivity to maneuvers, the Monte-Carlo results are plotted against the station keeping maneuver parameters in Figure C.15. The results show a bi-modal distribution with peaks corresponding to large maneuvers pointing in the $\pm\hat{j}$ direction, which corresponds to an azimuth of ± 90 deg. This makes sense since maneuvers in this direction

are aligned with the measurement sensitivity matrix at apoapsis. Given these results it may be possible to roughly estimate the maneuver parameters that result in a given z value, but that is future work. The stable NRHO produced similar results to the 9:2 NRHO, so the plots are omitted.

Figure C.15: Estimated control integral plotted with the true station keeping pointing parameters from the Monte-Carlo analysis given the active policy.



The control policy estimate was also plotted against the station keeping maneuver magnitude and the results are in Figure C.16. It can be observed from the figures that the optimal control policy returns an appropriate representation of the true mismatching and could be used to estimate maneuver magnitude because of the apparently linear dependence.

The quiet station keeping policy with a mean maneuver size of 50 mm/s and standard deviation of 15 mm/s produced similar distributions, excepting the correlation between azimuth and elevation of the station keeping policy appears less bi-modal and more uniform. The histograms of the control policy integrals for the 9:2 NRHO are in Figure C.17 and the numerical PDFs are in Figure C.18. Given these results $z_{lim} = 17.22$ mm/s and the expected hypothesis test results are given in Table C.6.

A scatter plot of the maneuvering Monte-Carlo results is displayed in Figure C.19 to show how the control policy varies with maneuver parameters, and Figure C.20 shows just the relation between magnitudes. Once again, the stable NRHO produced similar results to the 9:2 NRHO so the plots are omitted.

Figure C.16: Estimated control integral plotted with true station keeping magnitude from the Monte-Carlo analysis given the active policy.




Table C.6: Expected 9:2 NRHO maneuver detection probability for the quiet station keeping policy based on Monte-Carlo results.

| | No Maneuver | Maneuver |
|------------------------|-------------|----------|
| No Maneuver Identified | 0.9960 | 0.0160 |
| Maneuver Identified | 0.0040 | 0.9840 |

Figure C.17: Histogram of Monte-Carlo results for 50 mm/s mean non-maneuvering and maneuvering case.




Figure C.18: Numerical PDF of Monte-Carlo results for 50 mm/s mean non-maneuvering and maneuvering case.




Figure C.19: Estimated control integral plotted with the true station keeping pointing parameters from the Monte-Carlo analysis given the quiet policy.





Figure C.20: Estimated control integral plotted with true station keeping magnitude from the Monte-Carlo analysis given the quiet policy.



It is important to note that 2 different distributions are examined here to demonstrate the flexibility of the algorithm on 'larger' and 'smaller' expected maneuvers. When applying this to a real system the level of station keeping should be readily available since there are numerous studies on station keeping for NRHOs which detail expected maneuver sizes [15, 16, 93, 95, 96]. If, for some reason, this information is not readily available then one simply needs to run a bank of filters with various levels of expected maneuvers. The filter that performs the best should then be selected. Performance can be based on standard metrics such as residual levels, normalized innovation squared testing, or by selecting the filter that returns the most likely control policy.

C.4.3 Maneuver Detection

With the OCBE control policy distributions approximated, test data with unknown maneuvers is now examined. 100 maneuvering and 100 non-maneuvering data sets were generated for every policy and reference orbit, and the integrated control policy from the unknown data was used to test for maneuvers given Equation C.14.

For the 9:2 NRHO with active station keeping, the maneuver tests returned an accuracy of 100%, and with quiet station keeping the maneuver tests returned an accuracy of 97%. The full results are in Table C.7. The stable NRHO had a marginally better performance and the results for the stable NRHO are in Table C.8.

Table C.7: 9:2 NRHO Maneuver Detection Results

| | Active ($\mu = 1m/s, \sigma = 0.3m/s$) | | Quiet ($\mu = 50mm/s, \sigma = 15mm/s$) | |
|------------------------|---|----------|--|----------|
| | No Maneuver | Maneuver | No Maneuver | Maneuver |
| No Maneuver Identified | 100 | 0 | 97 | 3 |
| Maneuver Identified | 0 | 100 | 3 | 97 |

Table C.8: Stable NRHO Maneuver Detection Results

| | Active ($\mu = 1m/s, \sigma = 0.3m/s$) | | Quiet ($\mu = 50mm/s, \sigma = 15mm/s$) | |
|------------------------|---|----------|--|----------|
| | No Maneuver | Maneuver | No Maneuver | Maneuver |
| No Maneuver Identified | 99 | 1 | 100 | 0 |
| Maneuver Identified | 0 | 100 | 0 | 100 |

C.4.4 Control as a Health Metric

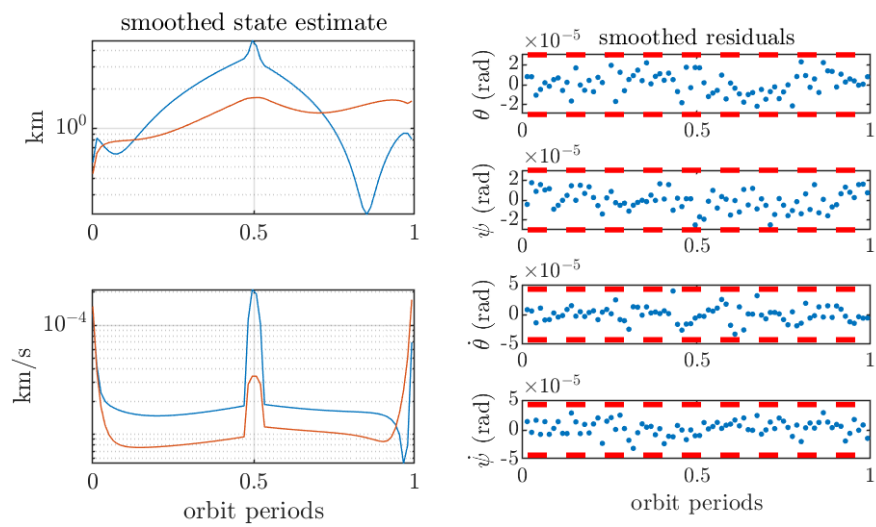
In addition to being able to identify, and potentially characterize, stochastic maneuvers the integral of the control policy can be utilized as a fast acting health metric for the OCBE. The fast acting nature of this metric is due to the control metrics sensitivity to mismodeling, which often shows significant deviations before measurement residuals do. This behavior is demonstrated in the following simulation where the OCBE is tuned for the quiet station keeping policy, but the observed system undergoes a larger maneuver of 0.5 m/s. The smoothed state estimate and measurement residuals are plotted in Figure C.21.

Because the system is not tuned to compensate for such a large maneuver the filter diverges, but this divergence does not appear evident in the measurement residuals. Most common filters rely on measurement residuals to indicate filter deterioration, and therefore would miss this filter divergence. However, when the control policy is examined for this system $z = 0.4761$ m/s. This is approximately at 30σ of the maneuvering distribution of \mathcal{Z} , thus is strongly indicating mismodeling. Not only is it indicative of mismodeling, but is also produces an appropriate estimate of the maneuver size. This demonstrates the value of control policy as a fast acting health metric for filtering.

C.5 Conclusion

Cislunar SSA is crucial to maintain the safety of spacecraft going to the Moon and beyond. This task comes with a variety of challenges and is exacerbated by limited Earth-based resources and chaotic dynamics. First, the results in Section C.3 reformulated the ballistic OCBE into SRI

Figure C.21: Filter performance given a larger maneuver than expected. The filter estimate diverges (left) but the measurement residuals appear benign (right) missing the maneuver. The optimal control integral z was approximately 30σ though which does catch the maneuver.



space to help reduce numerical instabilities related to filtering. Then the results went on to show that including angular rate information as a measurement in linearized filters is critical to the success of optical cislunar observation and, given angular rate measurements, only a single observer is necessary to provide accurate state estimation.

Next, to ensure an accurate solution for cislunar vehicles then some form of maneuver detection or compensation needs to be implemented to handle the necessary station keeping maneuvers for cislunar vehicles. In this paper this is accounted for via the newly presented hypothesis testing utilizing the OCBE's control policy. Results from Section C.4 show that maneuvers with a mean of 50 mm/s and a 15 mm/s standard distribution were correctly identified 194 out of 200 times for the 9:2 NRHO and 200 out of 200 for the stable NRHO. The value of OCBE control as a health metric is also demonstrated as it experiences large deviations to its expected value before measurement residuals do.

This work opens the door for various future developments. First, the demonstrated methodology should be implemented in a high fidelity model subject to additional dynamic, observer state, and measurement uncertainty to verify the results remain as accurate. Second, the Monte-Carlo approximation of the control policy distribution can be used to reconstruct the unknown maneuver. Lastly, the filter and maneuver detection performance should be examined given multiple unmodeled events in a single orbit period. Expanding on the presented work in this way would lead to improved tracking performance for cislunar vehicles.