
Doctoral Dissertations

Student Theses and Dissertations

Summer 2020

Applications of information theory in filtering and sensor management

Matthew James Gualdoni

Follow this and additional works at: https://scholarsmine.mst.edu/doctoral_dissertations



Part of the [Aerospace Engineering Commons](#)

Department: Mechanical and Aerospace Engineering

Recommended Citation

Gualdoni, Matthew James, "Applications of information theory in filtering and sensor management" (2020). *Doctoral Dissertations*. 2914.

https://scholarsmine.mst.edu/doctoral_dissertations/2914

This thesis is brought to you by Scholars' Mine, a service of the Missouri S&T Library and Learning Resources. This work is protected by U. S. Copyright Law. Unauthorized use including reproduction for redistribution requires the permission of the copyright holder. For more information, please contact scholarsmine@mst.edu.

APPLICATIONS OF INFORMATION THEORY IN FILTERING AND SENSOR
MANAGEMENT

by

MATTHEW JAMES GUALDONI

A DISSERTATION

Presented to the Graduate Faculty of the

MISSOURI UNIVERSITY OF SCIENCE AND TECHNOLOGY

In Partial Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

in

AEROSPACE ENGINEERING

2020

Approved by:

Kyle J. DeMars, Advisor

Christopher D' Souza

Carolin Frueh

Serhat Hosder

Henry Pernicka

Jagannathan Sarangapani

Copyright 2020

MATTHEW JAMES GUALDONI

All Rights Reserved

ABSTRACT

A classical sensor tasking methodology is analyzed in the context of generating sensor schedules for monitoring resident space objects (RSOs). This approach, namely maximizing the expected Kullback-Leibler divergence in a measurement update, is evaluated from a probabilistic perspective to determine the accuracy of the conventional approach. In this investigation, a new divergence-based approach is proposed to circumvent the myopic nature of the measure, forecasting the potential information contribution to a time of interest and leveraging the system dynamics and measurement model to do so. The forecasted objective exploits properties of a batch measurement update to frequently exhibit faster optimization times when compared to an accumulation of the conventional myopic employment. The forecasting approach additionally affords the ability to emphasize tracking performance at the point in time to which the information is mapped.

The forecasted divergence is lifted into the multitarget domain and combined with a collision entropy objective. The addition of the collision consideration assists the tasking policy in avoiding scenarios in which determining the origin of a measurement is difficult, ameliorating issues when executing the sensor schedule. The properties of the divergence-based and collision entropy-based objectives are explored to determine appropriate optimization schemes that can enable their use in real-time application. It is demonstrated through a single-target tasking simulation that the forecasted measure successfully outperforms traditional approaches with regard to tracking performance at the forecasted time. This simulation is followed by a multitarget tasking scenario in which different optimization strategies are analyzed, illustrating the feasibility of the proposed tasking policy and evaluating the solution from both schedule quality and runtime perspectives.

ACKNOWLEDGMENTS

The work presented in this dissertation is not possible without the resources provided by my advisor, Dr. Kyle DeMars. Throughout my graduate school experience, he has been an incredible source of guidance, technical expertise, and patience, and I am grateful for everything he has provided. I would also like to thank the members of my committee, Drs. Henry Pernicka, Chris D'Souza, Carolin Frueh, Serhat Hosder, and Jagannathan Sarangapani. The perspectives they provided aided in elevating this work to its full potential.

I would like to thank everyone involved that enabled me to receive the Graduate Assistance in Areas of National Need (GAANN) fellowship.

Thank you to the past and present members of the AREUS lab, as well as several other notable individuals that provided me the professional and personal relationships that enabled me to live and enjoy my life. James McCabe, our conversations were both technically helpful and at times necessarily nonsensical. Christine Schmid, I am grateful that you and I were able to go through this process concurrently and commiserate. Matt Glascock, I couldn't have asked for a better roommate through those years. The list could go on, but those of you that have assisted me in this endeavor, you know who you are; please know that you are sincerely appreciated.

Thank you to Stephanie Martensen. Your counseling was pivotal throughout my time in Rolla. You were a much needed outlet, and have become a friend I will not forget.

Thank you to my closest friend and colleague, Kari Ward. You have been so patient and helpful, and provided a solid foothold in reality through times of calm and quarantine.

This work is dedicated to my brother, Michael, my mother, Janice, and my father, James. Their strength, courage, and perseverance has been a constant source of inspiration, and has allowed me to endure my graduate journey, while as a family we weathered our roughest. My praise and appreciation for them alone could fill the pages of this work and more. Thank you for all of the sacrifices you have made and all of the love you provide.

TABLE OF CONTENTS

	Page
ABSTRACT	iii
ACKNOWLEDGMENTS	iv
LIST OF ILLUSTRATIONS	ix
LIST OF TABLES	xi
 SECTION	
1. INTRODUCTION	1
1.1. MOTIVATION	1
1.2. HISTORICAL PERSPECTIVE AND PREVIOUS WORKS	3
1.3. CONTRIBUTIONS	9
1.4. ORGANIZATION OF THE DISSERTATION	9
1.5. SOME COMMENTS ON NOTATION	11
2. PROBABILITY AND INFORMATION THEORY	13
2.1. UNCERTAINTY REPRESENTATION	13
2.1.1. Probability Theory	14
2.1.2. Finite Set Statistics	17
2.2. ENTROPIES AND INFORMATION	22
2.2.1. Shannon Entropy	22
2.2.2. Rényi Entropy	26
2.3. INFORMATION DIVERGENCES	29

3.	DYNAMICS AND MEASUREMENT MODELING	33
3.1.	DYNAMICAL MODELING	33
3.1.1.	Clohessy-Wiltshire Model	34
3.1.2.	Two-Body Problem	35
3.2.	OBSERVER MODELING.....	36
3.2.1.	Field of View, Field of Regard, and Probability of Detection	37
3.2.2.	Sensor Types	38
3.2.2.1.	Range and range rate observations	38
3.2.2.2.	Right ascension and declination observations	39
3.3.	LINEARIZATION	40
4.	SINGLE TARGET SENSOR TASKING	41
4.1.	SINGLE TARGET FILTERING APPROACHES	42
4.1.1.	The Kalman Filter	42
4.1.1.1.	Predictor	43
4.1.1.2.	Corrector	44
4.1.2.	The Extended Kalman Filter	45
4.1.2.1.	The information update	47
4.1.2.2.	The batch update	47
4.2.	INFORMATION DIVERGENCE AS A TASKING OBJECTIVE.....	49
4.2.1.	Myopic Kullback-Leibler Divergence	49
4.2.1.1.	Statistical analysis	53
4.2.1.2.	Conventional employment	56
4.2.1.3.	Objective function formulation	59
4.2.2.	Forecasted Kullback-Leibler Divergence.....	64
4.2.2.1.	Development.....	65
4.2.2.2.	Objective function formulation	68

4.2.2.3.	Comparison to myopic	68
4.2.3.	Optimization Considerations.....	72
4.2.4.	Computation in Other Spaces	74
5.	MULTITARGET SENSOR TASKING	78
5.1.	MULTITARGET DIVERGENCE EXTENSION.....	79
5.2.	SUBMODULAR FUNCTION ANALYSIS	80
5.2.1.	Myopic Kullback-Leibler	80
5.2.2.	Forecasted Kullback-Leibler	84
5.3.	DATA ASSOCIATION	85
5.4.	ENTROPY CONSIDERATION.....	89
5.4.1.	Shannon Entropy	89
5.4.2.	Rényi Entropy.....	90
5.5.	OBJECTIVE FUNCTION FORMULATION	93
5.6.	OPTIMIZATION.....	97
5.6.1.	Convex Closure	98
5.6.2.	Greedy Approaches	101
5.6.2.1.	Greedy selection via observer	103
5.6.2.2.	Greedy selection via target.....	103
5.6.2.3.	Greedy selection via time	104
5.6.2.4.	Further consideration.....	105
6.	RESULTS.....	107
6.1.	SINGLE-TARGET SCENARIO.....	107
6.1.1.	Eccentric Orbit	107
6.1.2.	Nearly Circular Orbit.....	116
6.2.	MULTITARGET SCENARIO.....	121
6.2.1.	Continuous Optimizer Configurations.....	123

6.2.2. Optimization Runtimes	124
6.2.3. Collision Avoidance	127
6.2.4. Schedule Performance	129
6.2.5. Comments on Multiple Observers	133
7. CONCLUSIONS	138
7.1. RESEARCH SUMMARY	138
7.2. FUTURE RESEARCH DIRECTIONS	141
APPENDICES	
A. THE SHANNON ENTROPY FOR A POISSON RFS	143
B. THE RÉNYI ENTROPY FOR A POISSON RFS	147
REFERENCES	150
VITA	157

LIST OF ILLUSTRATIONS

Figure	Page
2.1. The 1-, 2-, and 3- σ intervals for a multitarget pdf.	18
4.1. A target following along its nominal trajectory (dashed black line) with two observation sets originating from two (potentially) different sensor locations; one consisting of three measurements (green) generated according to measurement model $h_A(\mathbf{x})$, the other consisting of four measurements (gold) and generated according to measurement model $h_B(\mathbf{x})$	60
4.2. An illustration of mapping the measurement set information to a reference time t_f for the Sensor A measurement set; note the arcs accompanying the state transition matrices illustrate their mappings from time t_k to their respective measurement times.	66
4.3. The myopic and forecasted KL divergence objective when scheduling a single measurement for a ground-based sensor, with the optimal measurement time denoted with an \times	69
4.4. The entropy of the state estimate pdf when the myopic and forecasted schedules are processed by an EKF.	70
4.5. The entropy of the state estimate pdf with the optimal forecasted schedule and the two mixed schedules.	71
4.6. An eccentric orbit sampled evenly across time (left) or across the universal anomaly (right) with units of Earth radii (ER).	74
5.1. The 1- σ intervals for the two targets centered on their respective means and plotted at each time step in the state space.	86
5.2. The squared Mahalanobis distance of the expected measurement generated by Target 2 with respect to the measurement density for Target 1.	88
5.3. A simple example of the interaction of three targets.	90
6.1. The resulting optimal schedules of varying sizes utilizing the forecasted objective (top) and the range and range-rate measurement profiles (bottom).	110
6.2. The sample states used in generating the initial guess for the optimizer.	111
6.3. The differing sensor tasking schedules (top) plotted alongside the range and range-rate measurement profiles (bottom).	113

6.4. Average entropy over the simulation for each sensor tasking objective for the eccentric orbit scenario.....	115
6.5. The resulting optimal schedules of varying sizes utilizing the forecasted objective (top) and the right ascension and declination measurement profiles (bottom).....	118
6.6. The sample states used in generating the initial guess for the interior-point optimization.....	119
6.7. The differing sensor tasking schedules (top) and the right ascension and declination measurement profiles (bottom).	119
6.8. The average entropy over the simulation for each sensor tasking objective for the nearly circular orbit scenario.....	120
6.9. The measurement profiles of the twelve targets.	123
6.10. The phase one and two schedules generated via the myopic convex closure objective function optimized target-by-target, plotted with the collision avoidance measure.	128
6.11. The phase one and two schedules generated via the forecasted convex closure objective function optimized target-by-target, plotted with the collision avoidance measure.....	128
6.12. The average collision measures at each measurement time for the different myopic schedules.	130
6.13. The average collision measures at each measurement time for the different forecasted schedules.	130
6.14. The resulting myopic and forecasted phase one schedules for each of the different submodular optimization schemes as compared to the continuous optimization of the first phase as well as the simultaneous optimization.	132

LIST OF TABLES

Table	Page
6.1. The position RSS values (in meters) at the end of the two observation windows for the eccentric orbit scenario.	114
6.2. The velocity RSS values (in millimeters per second) at the end of the two observation windows for the eccentric orbit scenario.	114
6.3. The position RSS values (in meters) at the end of the two observation windows for the nearly circular orbit scenario.	121
6.4. The velocity RSS values (in millimeters per second) at the end of the two observation windows for the nearly circular orbit scenario.	121
6.5. The initial states for the 12 targets, given in the Hill frame.	122
6.6. The myopic objective function runtimes for each of the optimization schemes normalized by the simultaneous optimization.	125
6.7. The forecasted objective function runtimes for each of the optimization schemes normalized by the simultaneous optimization.	125
6.8. The system entropy (in nats) at the end of the hour long time interval and the different phase two sensor schedules have been processed.	133
6.9. The phase one optimization runtimes for the greedy-in-observer (G.I.O.) and the greedy-in-observer and greedy-in-target (G.I.O. G.I.T.) convex closure objective functions, normalized against the G.I.O. convex closure runtime.	135
6.10. The system entropy (in nats) at the end of the hour long time interval for each convex closure optimization, optimizing the space-based sensor schedule first followed by the GDSCC.	135
6.11. The system entropy (in nats) at the end of the hour long time interval for each convex closure optimization, optimizing the GDSCC schedule first followed by the space-based sensor.	136

1. INTRODUCTION

1.1. MOTIVATION

Since the dawn of space exploration, attempts have been made to maintain a level of awareness of the objects in orbit about Earth. Catalogs were created and maintained to not only record knowledge of the location and trajectory of these resident space objects (RSOs), but also in an attempt to determine the intent of unknown RSOs. The launch of Sputnik I in 1957 placed the first artificial satellite in orbit about the Earth, and the population of RSOs has only increased since. Advances in technology over the decades have enabled more launches to occur and more satellites to be deployed, and, consequently, RSO catalogs have been forced to grow in magnitude.

The increase in RSOs is not solely a function of the new assets placed in orbit either; the risk of conjunction events in orbit is ever-present and can result in the generation of a significant amount of space debris with dynamics that are difficult to predict or model. In 2009, an unintentional collision occurred between the American Iridium 33 satellite and Russian Kosmos 2251 spacecraft, generating an array of debris of varying sizes [86]. In 2013, two similar events occurred: a collision between a Russian BLITS nano-satellite with debris from the destroyed Fengyun FY-1C satellite [46], and a collision between Ecuador's Pegasus cubesat with a debris cloud from a Tysklon-3 upper stage left over from a previous launch [66]. Each of these events can generate hundreds of new RSOs [45], a frustrating fact as each conjunction might have been avoided with sufficient knowledge of the objects present and their trajectories.

The problem of discovering these RSOs after their creation, be it through collisions, decommissions, or launches, is one of significant importance; however, as the population of RSOs continues to grow unbounded, the problem of improving knowledge and maintaining

awareness of previously detected objects becomes increasingly difficult. Further, the rate at which RSO catalogs grow will likely do nothing but increase, not only due to advances in space programs across the globe, making satellite deployment more accessible, but also due to the advances in ground-based sensors as the size of detectable objects decreases and the distances at which they can be detected increases.

As the number of RSOs has grown, so have the complications they induce. From a military perspective, a larger population of satellites provides more cover for nefarious RSOs to carry out their objectives. Maintaining awareness of these potential adversaries becomes compounded by the growing popularity of cubesats, a much smaller form factor spacecraft as compared to older satellites, providing an option to reduce the size and thus the ability to detect RSOs. Alternatively, the presence and density of RSOs pose potential threats to upcoming efforts for crewed space flight and, as such, possessing knowledge of their orbital characteristics at any given time is paramount to successful mission planning. Research in recursive filtering has been heavily leveraged in addressing this issue; given ample data, many different robust and reliable filtering approaches exist for propagating and updating target state estimates, allowing RSO catalogs to be maintained for reference when planning or executing missions or in determining the intent of particular targets. In order to maintain these state estimates, observations must be acquired through either ground- or space-based observer stations.

Diversity in geographic locations, measurement types, and the sensor fidelity of ground-based stations provides a variety of data available for a filter to process. For example, the United States Space Surveillance Network can provide radar or optical data on targets in deep space as well as lower Earth orbits, generating either range, range rate, right ascension, or declination measurements of the targets (or some combination thereof) from over thirty different ground-based station locations [87]. Data from space-based satellites, such as the U.S. Air Force's Space-Based Space Surveillance Block 10 system, can also be processed in tandem with the ground-based observations [29]. The differing perspectives

brought on by the different sensor locations and data types provide a diverse collection of information, enabling different facets of a target state to be honed and the overall state of the target to be determined to a reasonable precision. However, while tracking RSOs can be achieved with ample data, the true problem lies in the number of these objects, as they vastly outnumber the available sensor resources. It is simply not possible for ground-based sensors to continuously observe every RSO in the sky, necessitating a means of properly scheduling measurements of RSOs in an effective manner. These schedules should maintain target state estimates without neglecting other targets and without losing track custody (i.e., the ability to regain sight of the target), while utilizing the differing data types and perspectives provided by the available sensor resources. This is an example of the problem of sensor tasking, and this question has been posed and investigated under several different paradigms.

1.2. HISTORICAL PERSPECTIVE AND PREVIOUS WORKS

The initial conception of a sensor tasking policy with respect to space object tracking began with the launch of Sputnik 1. Due to the anticipation of the first launches of human-made space vehicles and the concerns experienced during the Cold War, much thought and consideration had been put into a system in which these artificial satellites would be tracked and observed. One individual in particular was Harvard astronomer and then director of the Smithsonian Astrophysical Observatory (SAO) Dr. Fred Whipple, who in 1955 defined a three phase plan for addressing the problem. Phase one would establish the Minitrack network, a network of sensors capable of observing satellites. Phase two would establish the SAO as the centralized computing station that would receive and process all of the observations received by the network. All of this performance would be dependent upon phase three, dubbed Operation Moonwatch, a program that would enlist the amateur scientists of the public to turn their cameras and telescopes upward in an attempt to identify potential craft as they traversed the sky [60].

The initial observations from Operation Moonwatch gave a rough idea of where objects of interest might be in the sky in order to provide an approximate pointing direction to the sensors located across the globe. In the end, it was this group of public servants, ranging from astronomy hobbyists to high school science groups, that was called into action as the launch of Sputnik I took the world, including the professionally staffed tracking stations, by surprise [60, 80]. Operation Moonwatch was discontinued in 1975, and the Minitrack network eventually became obsolete, being replaced by more advanced and more specialized sensor networks such as the Tracking and Data Relay Satellite System and the Deep Space Network. However, investigations into the problem of sensor tasking have continued to this day, with the complexity and sophistication of the proposed solutions growing proportionally to the technologies that the policies are scheduling.

The problem of sensor tasking is an intrinsic problem in the areas of orbit determination, space situational awareness (SSA), and space traffic management (STM), and the number of avenues that can be taken in addressing the problem is staggering. Depending upon the particular goal for a solution, referred to as a sensor tasking objective or a sensor management policy, the actual implementations or policies can vary drastically. Broadly speaking, space object sensor tasking objectives can be classified under two high-level categories: object discovery and catalog maintenance. The former is of primary concern in building catalogs of newly identified RSOs, whereas the latter is concerned with maintaining these catalogs and observing any potential changes to an RSO orbit brought on by factors such as maneuvers or orbit decay. The two objectives are not necessarily mutually exclusive, but comprise the broad goals that a sensor tasking policy intends to address in SSA and STM applications.

Object discovery and initial orbit determination are imperative to constructing RSO catalogs and enabling sensor networks to monitor and maintain knowledge of the positions and velocities of the RSOs. Due to the nature of the problem, however, determining sensor management strategies is difficult with a complete lack of knowledge of an object's

existence, let alone its trajectory. At best, surveying for new objects can be handled from a sensor operations perspective, scanning the sky (or a portion of interest) in some clever manner and capturing images in a specific way. Solutions in this domain focus on aspects such as varying exposure times of optical sensors to capture streaks of moving objects, masking techniques to eliminate known stars from collected images, and sensor trajectory designs such as scanning across fixed right ascension angles to generate composite images (or data tracklets) of a portion of the sky to search for any unknown objects [74].

Alternatively, these methods can also be used to revisit objects that have previously been discovered, utilizing heuristic principles and *a priori* information to anticipate a pass in designating a sensor action [24, 25]. This is an example of a policy providing dual functionality in both object discovery and catalog maintenance; the *a priori* knowledge allows additional observations to be acquired of the object, while swaths of the sky are captured and processed for new object detections. Simultaneously growing and maintaining an RSO catalog is efficient, but as the size of RSO catalogs continues to grow, so does the complexity of maintaining their entries. Consequently, sensor tasking policies that place an emphasis on maintaining custody of RSOs within a catalog have received much attention over the years, resulting in a plethora of different strategies covering a wide range of sophistication and complexity.

A common theme throughout many of these sensor tasking approaches is to place an emphasis on a particular facet of the multiobject (or multitarget) tracking problem. For instance, policies can be defined almost exclusively on the abilities of the sensor being tasked. One approach is to perform an observability analysis to determine at which point along the nominal target trajectory the object's position and velocity are the easiest to infer from the perspective of the employed sensor. This is achieved by utilizing the singular values of the observability Gramian as indicators or measures of observability for a given target [23]. Another strategy that is focused primarily on the measurements quantifies the amount of information a sensor can provide through the use of the Fisher information gain

(FIG) [44]. Alternatively, the dynamics governing the motion and evolution of the target state estimate can be examined through an assessment of the stability of the system; for instance, leveraging the largest Lyapunov exponent (LLE) estimate enables observations to be taken at points along the target trajectory that are more sensitive to the system dynamics as compared to the rest of the trajectory [89]. Alternatively, more sophisticated methods such as machine learning techniques can utilize training data or positive feedback to produce schedules; methods such as these have shown to be very useful and versatile, however their performance is dependent upon factors such as training data or tuning parameters that require careful consideration [54]. It is of interest to explore solutions that leverage knowledge of the employed models (such as the observability analysis, FIG-based, and LLE-based approaches) that more explicitly considers the state estimate density.

Some approaches that fall under this paradigm leverage advances in multitarget statistics and are built upon multitarget filters. The posterior expected number of targets (PENT) or the posterior expected number of targets of interest (PENTI) can be maximized to encourage a sensor platform to observe the most targets possible, or similarly the uncertainty in the number of targets present can be minimized [58]. These measures operate well as approximations to specific information theoretic measures and emphasize the importance of observing each target to ensure that they are not neglected, but again do not directly consider the impact an observation will have on the state estimates at hand. Other methods more explicitly handle the reduction in the uncertainty of the individual target state estimates, such as through the use of the reduction matrix [38] or more commonly through the use of information theoretic measures [2, 20, 27, 37, 50, 89].

Information theory, in this context, provides an array of tools enabling the uncertainty in a probability density function (pdf) to be quantified. One fundamental category of tools, known as information entropies or simply entropies, enables the uncertainty in a pdf to be quantified in a volumetric sense [75]. Further, entropies facilitate the development of measures that yield a directed “distance” between two pdfs; these distance-like measures

provide a non-negative scalar quantification of the difference between two pdfs, akin to a distance metric though lacking the symmetric property and failing to satisfy the triangle inequality. When coupled with techniques such as expectation hedging [58], these measures, known as information divergences [75], seem natural to use in a tasking framework, as they provide a mathematically rigorous means of quantifying the “strength” of the update that is realized in a filter. Consequently, attempting to maximize this distance is equivalent to maximizing the strength of the filter update in an information theoretic sense, where the strength of the update is the quantity of information ingested by the filter due to a particular observation. Additionally, working in the information space with regard to state estimates gives equal footing for comparison of different sensor types, whether that would be different types of measurements, different qualities of data, different sensor locations, different acquisition times, or some composite of the aforementioned elements, as the uncertainties (or reductions in uncertainties) are translated from potentially different vector spaces to a common information space.

This approach has been leveraged and tested as a sensor tasking objective in many works, ranging from DeMars and Jah’s comparison of employing the first moment of several different information divergences [20] to Hintz and McVey’s measure of global information flow [39]. While not a criticism of the approaches, many information theoretic methods consider only the expected value of the information divergence measures that they are utilizing, neglecting higher-order information and raising concerns that useful context is being discarded. Furthermore, these information theoretic formulations typically have a “myopic” construction, meaning specifically that the observation acquisition is only concerned with the impact of a measurement at its specific observation time. The result in the single-target domain when considering a set of observations is an objective function that accumulates the first moment of the divergence measure at discrete and unique measurement

times; this first accumulation is then subsequently extended to the multitarget domain by simply further accumulating the expected divergences across all targets. While the approach is not inappropriate, there are several questions that are raised in its application:

- *What context is omitted when using only the first moment of the divergence?* Is there useful knowledge that is being discarded when taking the first moment of the divergence rather than attempting to use its full underlying distribution?
- *Is it appropriate to accumulate divergence measures that are computed at different points in time?* Does this summation hold any physical interpretation? Is there a more appropriate way to address the problem of simultaneously considering multiple observation times?
- *Can other facets of the multitarget sensor tasking problem be addressed simultaneously in the same information space?* The accumulation of expected information divergences across targets is a sound solution in the presence of independent and sufficiently separated targets; however, in the case where multiple targets can (potentially) occupy a sensor's field of view simultaneously, does this approach put systems at risk of entering potential misassociation scenarios (i.e. situations in which measurements are associated to targets that did not generate them, leading to problems such as label switching or track coalescence)? Is there a way to also consider and mitigate these risks?

Each of these questions warrants investigation. This dissertation documents this venture, detailing the theories and the resulting experiments as each point is explored in search of an effective and tractable sensor tasking policy. The result is new insight into a conventional divergence-based sensor tasking objective, the development of new sensor tasking objectives, and a study of their behavior to provide an intelligent and tractable means of optimizing sensor schedules.

1.3. CONTRIBUTIONS

This dissertation seeks to add to the body of work in addressing the problem of sensor tasking specifically in the context of space object catalog maintenance, space traffic management, and space situational awareness. The key contributions of this work can be summarized as:

- performing a statistical characterization of conventional Kullback-Leibler sensor tasking, enabling quantification of the accuracy of its approximation;
- developing a novel generalization of expectation-hedged divergence-based sensor tasking objectives, extending information divergence use to non-myopic (referred to herein as forecasted) sensor management policies that provide support for multiple observations from multiple observers of multiple targets;
- carrying out a classification of the myopic and forecasted sensor tasking objectives, providing insight into appropriate selection of optimization solutions as well as a survey of some of the available solutions; and
- creating an entropic characterization of potential conjunction events, providing a sensor tasking objective formulation that affords a user information-rich sensor schedules while avoiding potentially difficult data association events.

1.4. ORGANIZATION OF THE DISSERTATION

The subsequent sections are organized as follows: Section 2 provides the necessary background in probability and information theory, beginning first with a discussion of uncertainty representations for both single- and multitarget states. This is followed by Section 2.2 with the introduction of entropies, information theoretic measures that enable

the scalar quantification of the uncertainty present in stochastic variables. The section concludes with the introduction of information divergences, defining the Kullback-Leibler divergence and its interpretation in the context of sensor tasking.

Section 3 provides the dynamical and observational modeling employed in this work. The different dynamic models are presented first, beginning with the general definition and continuing to define the two-body and Clohessy-Wiltshire models. Section 3.2 gives the general definition for the observer model and discusses different aspects to consider in a sensor tasking context, concluding the section with definitions of the measurement models that are employed in Section 6.

Section 4 presents a detailed discussion of single-target sensor tasking, beginning with the concept of filtering in Section 4.1. A brief discussion of the different types of filters leveraged in this work is included in this section, providing the various forms of the update that are used in later developments. Section 4.2 details the use of the Kullback-Leibler divergence as a sensor tasking objective. The conventional use of the divergence measure is defined first, providing an explanation of its myopic nature and a statistical analysis of the quantity. This analysis enables the appropriateness of the typical first moment approximation to be justified and provides useful insight that is leveraged when formulating a multipurpose objective function in Section 5. Section 4.2.2 provides a novel concept that generalizes the myopic form of divergence-based tasking measures, forecasting the quantity to a reference time of interest. The development of the forecasted measure is given in detail along with discussion of its behavior under different assumptions and its connection to the myopic objective. The section presents some optimization considerations in Section 4.2.3 and concludes with the computation of the measure in the measurement space in Section 4.2.4.

Section 5 states necessary aspects to consider in the multitarget sensor tasking problem. The section opens with the extension of the objectives presented in the single-target tasking section to the multitarget domain in Section 5.1, followed by a proof classifying

both the myopic and forecasted objectives belong to the class of submodular set functions in Section 5.2. The section continues with a discussion of the problem of data association and the situations in which it can be difficult to address in Section 5.3. Section 5.4 proposes the use of information entropy to avoid difficult association scenarios and details its use as an objective simultaneously with the myopic and forecasted divergence objectives in Section 4. Section 5 goes on to formally construct the optimization schemes applied to the sensor schedule optimization in Section 5.6 and motivates a two phase optimization method in which the divergence- and entropy-based quantities are optimized separately, enabling submodular optimization schemes to be employed in the more computationally demanding optimization.

Section 6 exhibits two simulation scenarios. The first is intended to analyze the performance of the novel divergence-based objectives against other, more conventional information theoretic tasking policies. In this investigation, the behavior of the measures can be studied from a ground-based tracking scenario. The second simulation focuses on the results developed in Section 5, studying the performance of the proposed objective measures from both a schedule quality and runtime perspective in an increasingly demanding scenario, beginning with linear dynamics and a single observer and culminating with nonlinear dynamics and two observers.

Section 7 concludes the dissertation, providing final remarks on the investigations performed, additional questions the research has raised, and further directions for investigations to pursue.

1.5. SOME COMMENTS ON NOTATION

In the interest of clarity, this dissertation attempts to adhere to a consistent notation and terminology. In all of the following mathematical developments, lower case boldface symbols, such as \mathbf{x} and $\boldsymbol{\delta}$, indicate that the quantity is a vector, whereas non-bold lowercase symbols, such as t or δ , indicate scalar values. Similarly, boldface functions, such as $\mathbf{f}(\cdot)$,

indicate vector outputs while non-bold functions (and functionals) indicate scalar outputs, such as $M_{KL}(\cdot)$ (and $D_{KL}[\cdot]$). Capital boldface symbols, such as \mathbf{A} and $\mathbf{\Phi}$, represent matrices.

It is frequently the case that a subscript on a variable indicates the time index at which the variable exists, for instance target state \mathbf{x}_k at time t_k . When moving to the multitarget domain, a superscript in parentheses is added to indicate the particular target, i.e. $\mathbf{x}_k^{(i)}$ is the state of target i at time t_k . Lastly, in the event that there are also multiple observers, an additional subscript will precede the time index to indicate the current sensor, i.e. $\mathbf{z}_{s,k}^{(i)}$ denotes a measurement of target i at time t_k acquired by sensor s . Due to the nature of some of the developments, the indices (and the symbols representing them) may change, but this is explicitly noted within the text.

2. PROBABILITY AND INFORMATION THEORY

Given perfect mathematical models, absence of noise in sensor data, and full knowledge of any maneuvers, tracking the position and velocity of a space object would be greatly simplified, as the motion of the objects would be deterministic in nature. In practice, such a scenario is not realizable. Sensors inherently have noise, such as thermal noise for electrical sensors, interference power received by radar receivers, or errors accumulated in image processing for optical sensors. Models are rarely perfect, and errors due to model mismatch are generally present in both the dynamic and measurement models. Consequently, the processes that are being monitored are stochastic in nature, and in attempting to track the state of an RSO, it is useful to model the target state as a probabilistic quantity. This section covers the necessary background in probability and information theory to provide the foundation and context for the tools utilized throughout this work. Section 2.1 details modeling the target state; Section 2.2 discusses the roots of information theory and some of the fundamental tools that came from its development; and Section 2.3 provides the definition of information divergences and explains their initial use as well as the employment of the Kullback-Leibler divergence in the context of this work.

2.1. UNCERTAINTY REPRESENTATION

In order to account for the uncertainty present in tracking a target brought on by discrepancies between mathematical models and true behaviors as well as noisy sensors, a probabilistic handling of the target state and the measurement model is necessary. Through this treatment, the uncertainty in the true state can be incorporated in the model and noisy measurement data can be processed appropriately. In the single-target scenario, this is readily handled with basic probability theory. In the multitarget tracking case, the issue of maintaining target labels or target identities arises. This is ameliorated through an

extension of probability theory to reduce the problem dependence on the specific labels, namely finite set statistics (FISST). This section will detail the basics of probability theory and its generalization to FISST.

2.1.1. Probability Theory. The recursive filtering and the information theoretic tools (namely entropy and divergence) employed in this work require the random mapping of a variable to be defined along with the corresponding event space. Given a continuous random vector, \mathbf{x} , the event space for the random vector (or its support) can be defined as \mathcal{X} such that $\mathbf{x} \in \mathcal{X}$. For example, if the random vector of interest is defined to be the n_x -dimensional target state, the event space is defined as the space of all real numbers of dimension n_x , or $\mathcal{X} = \mathbb{R}^{n_x}$. The definition of the random vector and the event space enables the definition of a *probability density function* (pdf). A pdf is a scalar mapping of the random vector, \mathbf{x} , to a relative likelihood of a specific realization of the random quantity. More explicitly, the pdf $f : \mathcal{X} \rightarrow \mathbb{R} \geq 0$, $f(\cdot)$ exhibits the following properties [85]:

$$f(\mathbf{x}) \geq 0,$$

$$\int_{\mathcal{X}} f(\mathbf{x}) d\mathbf{x} = 1.$$

Possession of an analytic expression for the pdf of a random variable enables different properties of the random variable to be obtained. For instance, a fundamental definition the pdf affords is the *expectation* operator, a function that returns the value the stochastic input is most likely (on average) to take on; taking the expected value of a random variable, \mathbf{x} , with respect to pdf $f(\cdot)$ is defined as [61]

$$\begin{aligned} \mathbb{E}\{\mathbf{x}\} &\triangleq \mathbb{E}_{f(\mathbf{x})}\{\mathbf{x}\} \\ &= \int_{\mathcal{X}} \mathbf{x} f(\mathbf{x}) d\mathbf{x}, \end{aligned} \tag{2.1}$$

where the subscript $f(\mathbf{x})$ denotes the expectation is taken with respect to pdf, $f(\cdot)$. Note that throughout this work the pdf argument is often omitted with the understanding that the expectation is taken with respect to the pdf of the argument unless otherwise specified.

The variance-covariance matrix (herein denoted simply as the covariance matrix) is another useful piece in characterizing the behavior of the random vector, \mathbf{x} , as it contains not only the variances for the n_x individual components of \mathbf{x} (located on the diagonal of the covariance matrix), but also the covariances between the components (located on the off-diagonals of the covariance matrix). The covariance matrix of a random vector is defined to be the second central moment of the distribution, or more explicitly as

$$\text{Cov}\{\mathbf{x}\} = \mathbb{E}\{(\mathbf{x} - \mathbb{E}\{\mathbf{x}\})(\mathbf{x} - \mathbb{E}\{\mathbf{x}\})^T\}.$$

It is important to note that a covariance matrix, $\mathbf{\Pi}$, is necessarily symmetric, positive semi-definite (denoted $\mathbf{\Pi} \geq 0$), i.e. $\mathbf{\Pi} = \mathbf{\Pi}^T$ and

$$\mathbf{x}^T \mathbf{\Pi} \mathbf{x} \geq 0.$$

For a scalar random variable, the standard deviation is the square root of the variance, or

$$\sigma = \sqrt{\mathbb{E}\{(x - \mathbb{E}\{x\})^2\}}.$$

Higher-order central moments for the random variable can also be computed to more completely characterize the pdf. For these higher-order moments, it is typical to normalize them by the standard deviation, σ . For a scalar random variable x , the k^{th} standardized central moment is computed according to

$$\mu^{(k)} = \mathbb{E}\left\{\left(\frac{x - \mathbb{E}\{x\}}{\sigma}\right)^k\right\},$$

where $k = 3$ produces the skewness and $k = 4$ gives the kurtosis. Higher order moments can be computed indefinitely, yielding more information on the behavior of the random event. This may be useful in data analysis when working with experimental data, but in many filtering applications it is often beneficial to model the sources of stochasticity (e.g. target state, measurement noise, process noise) through the selection of an analytic pdf. Different random events warrant different pdf models, but perhaps the most commonly employed model for a pdf is the Gaussian distribution due to its prevalence in nature and its usefulness in modeling events such as measurement noise.

Given random vector $\mathbf{x} \in \mathcal{X}$, let $\boldsymbol{\mu} \in \mathcal{X}$ denote the mean of the distribution, and let $\boldsymbol{\Pi} \in \{\mathcal{X} \times \mathcal{X}\}$ be the covariance. The Gaussian distribution is then defined as

$$p_g(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Pi}) = |2\pi\boldsymbol{\Pi}|^{-1/2} \exp \left\{ -\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Pi}^{-1}(\mathbf{x} - \boldsymbol{\mu}) \right\}, \quad (2.2)$$

where $|\cdot|$ denotes the matrix determinant. Note that the Gaussian distribution is entirely characterized by its first two moments, that is to say that all that is needed to fully define the pdf is the mean and the covariance. This, as well as a handful of other useful properties exhibited by the Gaussian distribution, motivates its employment throughout this work. *Note that the assumption of Gaussianity is not necessary for the approaches presented here;* its selection is simply a function of its simplicity in implementation and analysis. Due to this fact as well as the prevalence of Gaussian densities in literature and application, the assumption is taken on throughout the duration of this dissertation.

The collection of these probability concepts lay the foundation for a probabilistic modeling of a random state or event, facilitating the modeling of the state of a single RSO through an expected state and associating an uncertainty in that state with a covariance matrix. Additionally, a similar treatment can be executed in modeling noisy sensors and

the uncertainties brought on by the imperfect dynamic and measurement models. However, in order to perform this modeling for a multitarget state, new facets must be considered; a more appropriate choice lies in FISST.

2.1.2. Finite Set Statistics. The probability theory concepts from the previous section can be extended to the multitarget domain in the case that the target labels are known with absolute confidence. For example, consider two scalar states labeled as states 1 and 2, x_1 and x_2 , that are modeled as Gaussian random variables with means $\mu_1 = 2$ and $\mu_2 = 1$, respectively, with associated variances $\sigma_1^2 = 4$ and $\sigma_2^2 = 1$. The resulting multitarget pdf can be modeled simply as the *joint pdf* of x_1 and x_2 , illustrated in Figure 2.1 with the stated labeling scheme depicted in the left plot. However, if the labels are flipped, the statistics would swap labels as well (i.e. $\mu_1 = 1$, $\mu_2 = 2$, $\sigma_1^2 = 1$ and $\sigma_2^2 = 4$), and the corresponding multitarget pdf is given on the right. What if the target labels are not known with any level of confidence? Rather, what if the problem is not concerned with target labels, but rather just the presence of targets in the space of interest? Some combination of the two pdfs in Figure 2.1 would have to be considered.

It is apparent from the previous example that when considering the multitarget problem, modeling the multiple states via random vectors is no longer appropriate if there is no method of maintaining awareness of the individual target identities. However, the field of finite set statistics (FISST) provides the framework to handle multitarget distributions without requiring the maintenance of target labels. Suppose L targets simultaneously exist in the state space \mathcal{X} with states $\mathbf{x}^{(1)}$, $\mathbf{x}^{(2)}$, \dots , and $\mathbf{x}^{(L)}$. Define an unordered set, \mathbf{X} , such that

$$\mathbf{X} = \left\{ \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(L)} \right\},$$

where there is no explicit order, the elements simply exist within the set, and permuting them does not change the set. This is referred to as a *random finite set* (RFS), which provides a mathematically rigorous approach to modeling a multitarget state in which the number of targets is not deterministic and target labels are not necessarily pertinent to the problem at

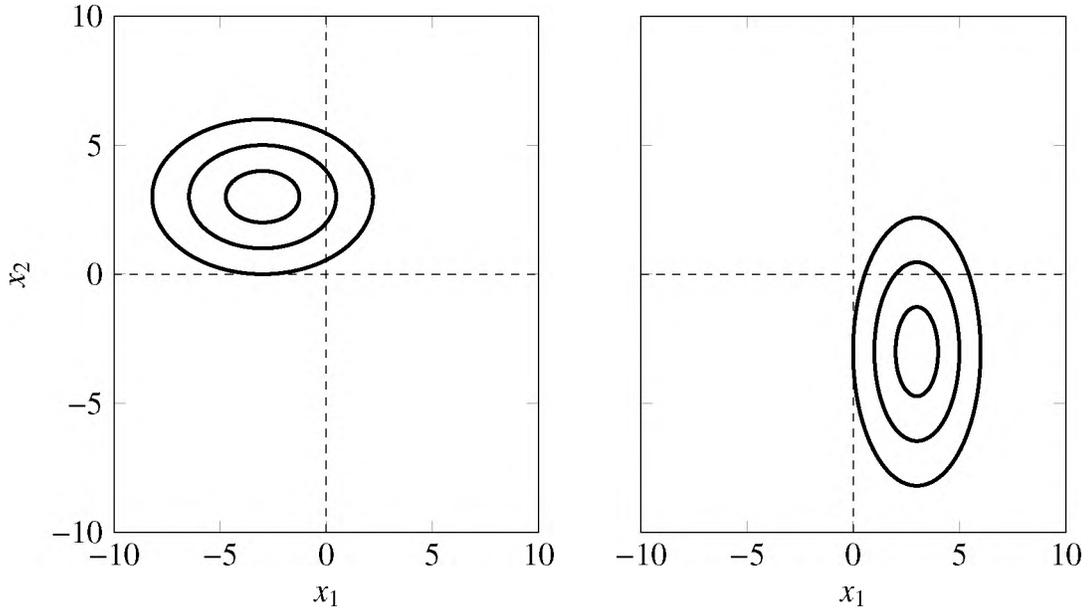


Figure 2.1. The 1-, 2-, and 3- σ intervals for a multitarget pdf. (left) The multitarget pdf with the correct label association. (right) The multitarget pdf with the labels exchanged.

hand. More formally, an RFS, Ψ , is a random variable that is realized as instantiations in the hyperspace of all finite subsets (including the null set) of some underlying space [57]. Casting the multitarget state (as well as the multitarget observations) as an RFS enables the target labels to be removed from the problem, allowing their collective existence in the state space to be maintained. This provides an option to either expend computational effort carrying the labels of the individual targets or to reserve computational resources by omitting them.

An RFS can also be described statistically with a density in a similar manner to the single-target problem, providing a solution to the situation in Figure 2.1 in which the labels of the targets are not known. Consider L single-target states, each with a corresponding single-target pdf $s(\mathbf{x}^{(k)})$. If these targets are said to be *statistically independent* (or just simply *independent*), the joint pdf is simply the product of the single-target (or *marginal*) pdfs, as

$$p(\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(L)}) = \prod_{k=1}^L s(\mathbf{x}^{(k)}). \quad (2.3)$$

Consequently, if it is assumed that the RFS, \mathbf{X} , is an independent and identically distributed (i.i.d.) cluster process, the multitarget pdf can be represented as [57]

$$f(\mathbf{X}) = L! \rho(L) \prod_{k=1}^L s(\mathbf{x}^{(k)}), \quad (2.4)$$

where L is the *cardinality* of the RFS (e.g. the number of targets in the multitarget tracking example), and $\rho(\cdot)$ is the cardinality density. In contrast to $p(\mathbf{x})$ in Eq. (2.3), by operating on an RFS, $f(\mathbf{X})$ in Eq. (2.4) does not directly consider the specific labels for the individual targets.

The factorial present in Eq. (2.4) accounts for all of the potential labeling events. Considering the example illustrated in Figure 2.1, it is clear to see that the multitarget pdf for the RFS consisting of the two targets is simply the sum of the two pdfs in Figure 2.1, multiplied by the cardinality as well as the cardinality density. A common choice is to model the cardinality density as a Poisson distribution with the corresponding pmf given as

$$\rho(L) = \frac{1}{L!} \lambda^L e^{-\lambda}, \quad (2.5)$$

where λ is the mean and the variance of the distribution, which is commonly referred to as the rate parameter. Substituting Eq. (2.5) into Eq. (2.4) yields

$$f(\mathbf{X}) = \lambda^L e^{-\lambda} \prod_{k=1}^L s(\mathbf{x}^{(k)}). \quad (2.6)$$

Equation (2.6) is the multitarget pdf for an i.i.d. cluster process that exhibits a cardinality following a Poisson distribution, referred to as a Poisson RFS.

In following the development of the RFS, it becomes clear that this is an elegant approach to the multitarget problem, providing generalizations to several statistical tools, e.g. probability generating functionals as opposed to probability generating functions. Consequently, the first moment of the multitarget density in Eq. (2.4) for an i.i.d. cluster

process can be generated similarly to the first moment of a single-target density via

$$v(\mathbf{x}) \triangleq \int f(\{\mathbf{x}\} \cup \mathbf{X}) \delta \mathbf{X}, \quad (2.7)$$

where $\delta \mathbf{X}$ denotes a set integral, which is defined as

$$\int f(\mathbf{x}) \delta \mathbf{X} = \sum_{L=0}^{\infty} \frac{1}{L!} \int_{\mathcal{X}} \cdots \int_{\mathcal{X}} s(\mathbf{x}) d\mathbf{x}^{(1)} \cdots d\mathbf{x}^{(L)}. \quad (2.8)$$

Equation (2.7) is the statistical first moment of the multitarget density and is referred to as the *probability hypothesis density* (PHD), the *intensity function*, or just simply the *intensity* of the RFS. Note that the intensity is *not* a pdf, as it is defined such that integrating over the entire support of the intensity function resolves to the expected cardinality of the set, or

$$\lambda = \int_{\mathcal{X}} v(\mathbf{x}) d\mathbf{x}. \quad (2.9)$$

That is, integrating the intensity function over its support does not integrate to unity. Nonetheless, it is a useful representation of the multitarget state, as it provides a relative measure of the likelihood of a target inhabiting a particular space and provides a means of collecting estimates for the individual target states. For example, let \hat{n} be the nearest integer to λ , and let $\boldsymbol{\mu}^{(1)}, \boldsymbol{\mu}^{(2)}, \dots, \boldsymbol{\mu}^{(\hat{n})}$ be the \hat{n} highest peaks in the intensity function. The collection of these vectors $\hat{\mathbf{X}} = \{\boldsymbol{\mu}^{(1)}, \boldsymbol{\mu}^{(2)}, \dots, \boldsymbol{\mu}^{(\hat{n})}\}$ is then a multitarget state estimate.

For the i.i.d. Poisson cluster process in Eq. (2.6), the intensity function can be found to be

$$v(\mathbf{x}) = \lambda s(\mathbf{x}). \quad (2.10)$$

Similar to the single-target scenario, there are several advantages to leveraging the Gaussian distribution in modeling or approximating the intensity function for a multitarget state density. This is achieved via the Gaussian mixture (GM) model, as [57, 84]

$$v(\mathbf{x}) = \sum_{k=1}^n w_k p_g(\mathbf{x}^{(k)}; \boldsymbol{\mu}^{(k)}, \boldsymbol{\Pi}^{(k)}), \quad (2.11)$$

where $\boldsymbol{\mu}^{(k)}$ and $\boldsymbol{\Pi}^{(k)}$ are the mean and covariance of the k^{th} component, respectively, and $w_k > 0$ is a weighting term for each Gaussian component. It is important to note that the mixture in Eq. (2.11) is slightly different than its employment in [4, 78] and other subsequent works in GM filtering methods. The difference is in the restrictions on the function being modeled; conventional GM representations are employed to approximate a non-Gaussian pdf, and as such the weights of the individual Gaussian components must sum to unity. However, it follows from Eqs. (2.9) and (2.11) that the weights will, in fact, sum to the cardinality of the set.

It is important to note, however, that the approaches detailed in this work do not preclude the use of GMs in approximating non-Gaussian densities. In the presented developments, each Gaussian corresponds to a single target, and consequently each of the weight coefficients are unity. However, the presence of the w_k term enables non-Gaussian densities to be modeled with multiple Gaussian functions, affording a means of extending the tasking methodologies to non-Gaussian problems. This is not the only means of considering densities outside the normal pdf, but in the developments to follow, it is straightforward to see how one might apply this work to a Gaussian mixture model, simply providing an avenue of thought. Regardless, in order to preserve a concise presentation of results, throughout the multitarget developments in this work, the single-target state densities are modeled with a single Gaussian, and all of the weights in the mixture are set to unity.

2.2. ENTROPIES AND INFORMATION

The probability theory and FISST concepts laid out in the previous section provide a means of associating a level of uncertainty with a state estimate through a probabilistic treatment. This section details select information theoretic concepts that can be leveraged with these probabilistic representations and provides a reference for the developments in the sensor tasking sections.

2.2.1. Shannon Entropy. The field of information theory was conceived in response to advances in communication methods, specifically in an attempt to generate a set of mathematical tools to appropriately handle signal processing. The field was first posited by Claude Shannon in his seminal work “A Mathematical Theory of Communication,” posing the fundamental problem of communication as one “of reproducing at one point either exactly or approximately a message selected at another point” [75]. This work was later expounded upon by Warren Weaver, unifying Shannon’s theory with previous works by esteemed individuals such as Ralph Hartley and Harry Nyquist and establishing the new field of mathematics known as *information theory*.

Shannon’s work approached the problem of communication by acknowledging that communication signals are stochastic in nature, and, as such, it is imperative to handle these signals as random events with some knowledge of the language or the alphabet that is employed in the communication. Consequently, endeavors in the field resulted in an array of tools capable of analyzing random events and their underlying distributions. The work began with a simple, fundamental element, namely the concept of information. This building block was derived with three intuitive properties in mind; for a random event x with n potential realizations x_1, x_2, \dots, x_n occurring with corresponding probabilities p_1, p_2, \dots, p_n , this notion of information should exhibit the following:

1. information should be a continuous function of the probability mass function (pmf; that is, the discrete equivalent to a pdf);

2. if each event is equally probable, i.e. $p_i = \frac{1}{n}$, the information should be a monotonically increasing function of n ; and
3. if a sequence of events can be decomposed into successive events, this decomposition should not affect the amount of information available.

The quantification of information that obeys these properties can be shown to be

$$H[p] = -\kappa \sum_{i=1}^n p_i \log p_i, \quad (2.12)$$

where κ is a positive constant that dictates units and p_i denotes the probability of event i . In this work, κ will be taken to be unity, and the natural logarithm will be employed resulting in units of nats (though other options are available, such as using a base two logarithm to yield units of bits). Due to the similarity to thermodynamic entropy, $H[\cdot]$ is referred to as *information entropy* in the sense of Shannon, or just simply Shannon entropy. This concept is extended to the continuous domain, yielding

$$H[p] = - \int_{\mathcal{X}} p(\mathbf{x}) \log p(\mathbf{x}) d\mathbf{x}, \quad (2.13)$$

where $p(\mathbf{x})$ is the pdf of \mathbf{x} . To differentiate between the discrete and continuous quantities in Eq. (2.12) and Eq. (2.13), the latter is sometimes referred to as *differential entropy*. The terms differential entropy and entropy will be used interchangeably throughout this work where there is no risk of ambiguity.

Equation (2.13) provides a means of mapping the full distribution of a random variable to a single scalar quantity that is representative of the spread of the distribution, a value that can be considered as the amount of information there is to gain on the random quantity or the amount of uncertainty in its outcome. In the context of communication and signal processing, this can be interpreted as a quantification of the amount of information communicated or a measure of certainty in a transmitted or received message. For instance,

when an alphabet consisting of two symbols is employed and a single symbol is transmitted, if the two symbols are equally likely to be transmitted, then a single bit of uncertainty is present, or a single bit of information has been communicated. The latter mentality is perhaps a more intuitive and straightforward line of reasoning, and upon first glance, the two perspectives may seem incompatible. It is useful to think of these as two sides of the same coin, with one paradigm applicable to a transmitter and the other to a receiver, or one frame of mind concerned with the amount of information communicated versus another concerned with how much information there is left to gain on an event. Once the event is known deterministically (i.e. the pmf or pdf is the Kronecker or Dirac delta), there is no longer any information to be gained on the subject, and the corresponding entropy is at its minimum (zero for discrete random variables or $-\infty$ for continuous random variables).

When \mathbf{x} is a target state and the pdf $p(\mathbf{x})$ describes the state density, the entropy indicates the volume of uncertainty (realized as a quantity proportional to the determinant of the covariance) associated with the target state estimate. This is an attractive measure to consider in the context of target tracking, as it gives a rough idea as to how useful an observation may be at a moment in time. Further, if the target state estimate is assumed Gaussian, a compact, analytic, closed-form solution is available, as

$$H[p] = \frac{1}{2} \log |2\pi e \mathbf{\Pi}|,$$

where $\mathbf{\Pi}$ is the covariance matrix of the pdf $p(\cdot)$. It is important to note that unlike the entropy for discrete probability spaces that quantify the randomness of a variable in an absolute way, differential entropy is not scale invariant. This means that the measure is relative to the selected coordinate system and thus computing the differential entropy in different coordinate systems can result in disparate quantifications [75]. For instance, given the linear transformation

$$\mathbf{y} = \mathbf{A}\mathbf{x},$$

where \mathbf{X} follows any arbitrary pdf $p(\mathbf{x})$ and \mathbf{A} is an invertible matrix, the resulting difference in entropy is [90]

$$H[q] - H[p] = \log |\mathbf{A}|,$$

where $q(\mathbf{y})$ is the pdf for random variable \mathbf{y} . This implies that, regardless of the distribution, a linear mapping of a random variable into any other coordinate system results in a constant difference in entropy, namely the log of the determinant of the transformation matrix. This lack of scale invariance is sometimes addressed with the inclusion of a multiplication of a unit hypersphere as

$$H[p(\mathbf{x})] = - \int_{\mathcal{X}} p(\mathbf{x}) \log (up(\mathbf{x})) d\mathbf{x},$$

where the density $p(\mathbf{x})$ has units of u^{-1} and u nullifies the units so as to avoid taking the logarithm of a unit [71]. However, this multiplication is usually implied, and will be herein.

Shannon entropy, as defined in Eq. (2.13), cannot be used directly to describe the uncertainty in a multitarget system, as it is the entropy obtained from a pdf representing the uncertainty of a single object. When the multitarget state is modeled as an RFS \mathbf{X} , Eq. (2.13) can be recast via the set integral, such that

$$H[f] = - \int_{\mathcal{X}} f(\mathbf{X}) \log \left(u^L f(\mathbf{X}) \right) \delta \mathbf{X}, \quad (2.14)$$

where u^{-L} are the units of the RFS density $f(\mathbf{X})$ and L is the cardinality of the RFS \mathbf{X} ; again, the inclusion of u is for mathematical completeness, but will be implied herein. In the event that the RFS is assumed to be a Poisson RFS, the Shannon entropy can be shown to be [16, 19]

$$H[f] = \lambda - \int_{\mathcal{X}} v(\mathbf{x}) \log \{v(\mathbf{x})\} d\mathbf{x}. \quad (2.15)$$

The derivation of Eq. (2.15) is provided in Appendix A.

Equation (2.15) is the Shannon entropy for an RFS under the assumption that it is distributed according to an i.i.d. cluster process, with the further stipulation that the cardinality distribution is Poisson. The result shows that the entropy is composed of a cardinality entropy term and a spatial entropy term, though it should be noted that the spatial entropy term in Eq. (2.15) still contains cardinality elements through the representation of the intensity function. The spatial entropy term is of the exact form of the single-target entropy given by Eq. (2.13), but with the multitarget intensity function in place of the single-target pdf. Thus, the spatial term will tend to exhibit the same characteristics observed with the single-target entropy, lending intuition to the analysis of multitarget entropy, and the cardinality term will cause the entropy to rise as the number of targets in the multitarget state increases.

Except in special cases of the intensity, such as an intensity that is Gaussian, the Shannon entropy of Eq. (2.15) cannot be found in closed-form. For instance, when the intensity is represented as a Gaussian mixture, no closed-form solution to the entropy of Eq. (2.15) can be found. In such situations, numerical solutions to the integral, such as those obtained through Monte Carlo integration, must be used.

2.2.2. Rényi Entropy. Shannon's notion of information entropy provides a useful scalar quantification of the amount of uncertainty or the amount of information there is to be gained on the target state(s). However, it is not the only option available for such a measure. The definition of entropy in Eq. (2.13) was later generalized by Alfred Rényi to produce a family of entropy measures [15, 69]. These measures, referred to as α entropies or Rényi entropy, are defined as

$$H^{(\alpha)}[p] = \frac{1}{1-\alpha} \log \left(\int_{\mathcal{X}} p^{\alpha}(\mathbf{x}) d\mathbf{x} \right), \quad (2.16)$$

where $\alpha \neq 1$ is a parameter defined by the user. In the limit as α approaches unity the Shannon entropy is recovered [69]. In addition to the Shannon entropy, other conventional entropy measures can be recovered [64], such as the Hartley entropy ($\alpha \rightarrow 0$), the min-entropy ($\alpha \rightarrow \infty$), or a measure that is referred to as the *collision entropy* ($\alpha \rightarrow 2$).

Similar to the concept of extending the Shannon entropy into the multitarget domain, the Rényi entropy can be applied to multitarget densities by employing the set integral in Eq. (2.8) to Eq. (2.16), yielding

$$H^{(\alpha)}[f] = \frac{1}{1-\alpha} \log \left\{ \int_{\mathcal{X}^n} f^\alpha(\mathbf{X}) \delta \mathbf{X} \right\}. \quad (2.17)$$

Just as in the single-target case, the multitarget Rényi can be shown to approach the multitarget Shannon entropy in Eq. (2.14) in the limit as α approaches unity; that is, the multitarget Rényi entropy is a generalization of the multitarget Shannon entropy. For the sake of mathematical completeness, it should be noted that the naive form given in Eq. (2.17) implies taking the logarithm of a unit; similar to the Shannon entropy in Eq. (2.14), Eq. (2.17) implies the multiplication of a unit hypersphere to address this issue. If \mathbf{X} is taken to be a Poisson RFS, Eq. (2.17) becomes [19]

$$H^{(\alpha)}[f] = -\frac{\alpha\lambda}{1-\alpha} + \frac{1}{1-\alpha} \int_{\mathcal{X}} v^\alpha(\mathbf{x}) d\mathbf{x}. \quad (2.18)$$

The derivation of Eq. (2.18) is detailed in Appendix B.

Equation (2.18) is the Rényi entropy of order α for an RFS under the assumption that it is distributed according to an i.i.d. cluster process, with the further stipulation that the cardinality distribution is Poisson. The result, much like the Shannon entropy, shows that the Rényi entropy is composed of a cardinality entropy term and a spatial entropy term, where it is worth noting that the spatial entropy term in Eq. (2.18) contains cardinality

elements through the representation of the intensity function. Unlike the Shannon entropy, however, the spatial element of the Rényi entropy does not take on the same form as the single-target Rényi entropy, which can be seen by comparing Eq. (2.16) and Eq. (2.18).

In contrast to the Shannon entropy of Eq. (2.15), the Rényi entropy of Eq. (2.18) can be found in closed-form for certain choices of the control parameter, α , when the intensity function is given by Eq. (2.11). For instance, in the case of collision entropy ($\alpha = 2$), the Rényi entropy for a Gaussian mixture is given as [32]

$$\begin{aligned} H^{(2)}[f] &= 2\lambda - \int_{\mathcal{X}} v^2(\mathbf{x})d\mathbf{x}, \\ &= 2 \sum_{\ell=1}^L w^{(\ell)} - \sum_{i=1}^L \sum_{j=1}^L \left[w^{(i)} w^{(j)} \Gamma(\boldsymbol{\mu}^{(i)} - \boldsymbol{\mu}^{(j)}, \boldsymbol{\Pi}^{(i)} + \boldsymbol{\Pi}^{(j)}) \right], \end{aligned} \quad (2.19)$$

where

$$\Gamma(\mathbf{a}, \mathbf{A}) = |2\pi\mathbf{A}|^{-1/2} \exp \left\{ -\frac{1}{2} \mathbf{a}^T \mathbf{A}^{-1} \mathbf{a} \right\}.$$

The first term in Eq. (2.19) is the cardinality entropy and is simply given as the sum of the weights of the GM representation of the intensity; for a fixed number of targets, this term is constant. The latter term describes the spatial entropy of the multitarget state, yielding a relative quantification of the concentration of entropy within the multitarget state; as the targets come in close proximity to one another, this term gets larger, and decreases as the targets move apart. In this sense, the spatial Rényi entropy can be considered a relative measure of the level of pdf coalescence between the single-target states, or the level of interaction between the state estimates. This is a fact that will be leveraged in Section 5.

2.3. INFORMATION DIVERGENCES

Generally speaking, an information divergence is a measure of the similarity (or dissimilarity) between two given pdfs [70]. A generic information divergence describes the directed distance between two pdfs $p(\mathbf{x})$ and $q(\mathbf{x})$ is denoted $D[p||q]$, and this “distance” is called a metric if [5]:

1. the quantity is non-negative, $D[p||q] \geq 0$;
2. $D[p||q] = 0$ necessitates $p \equiv q$;
3. the function exhibits the triangle inequality, $D[p||q] \leq D[p||r] + D[r||q]$; and
4. the function is symmetric, $D[p||q] = D[q||p]$.

Information divergences that only satisfy the first two conditions are not metrics and are referred to as asymmetric divergences. Satisfaction of the fourth condition necessarily removes the restriction of referring to the divergence as asymmetric.

This mapping of pdf distance to scalar values has proven to be useful in several different applications in the field of estimation, ranging from the characterization of collision events as in [19] to the minimization objective in the development of filters, such as the minimum divergence filter in [17]. The information divergence measure has been formulated in a variety of ways, each definition generally falling under the class of either f -divergences or Bregman divergences when working with probability spaces [5]. One unique formulation of information divergence solely inhabits the intersection of these two families of discriminations – the Kullback-Leibler divergence.

The Kullback-Leibler divergence came out of the work of Solomon Kullback and Richard Leibler in their efforts to expand upon the foundation of information theory set forth by Shannon and is a fundamental measure for describing a directed distance between two pdfs [51]. This measure provides a means of comparing two statistical populations in a mathematically rigorous manner, as it can be derived directly from the entropy definitions

laid out by Shannon. It is defined in terms of the information content (in the sense of Shannon) of two arbitrary pdfs that exhibit absolute continuity with respect to one another (or more specifically are defined over the same support) [52]. Just as before \mathbf{x} is taken to be the n_x -dimensional random target state vector whose elements belong to the target state space \mathcal{X} . The Kullback-Leibler divergence from $q(\mathbf{x})$ to $p(\mathbf{x})$ is defined as the expected value of the log ratio of the two distributions taken with respect to $q(\mathbf{x})$; specifically, the Kullback-Leibler divergence from $q(\mathbf{x})$ to $p(\mathbf{x})$ is defined as [51]

$$D_{KL}[q||p] = \int_{\mathcal{X}} q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x})} d\mathbf{x}. \quad (2.20)$$

Due to the fact that this divergence measure does not satisfy the triangle inequality nor does it exhibit symmetry, the Kullback-Leibler divergence is frequently referred to as a directed distance, even though it is not a proper distance metric.

Equation (2.20) gives the Kullback-Leibler divergence in its most general form, though it is useful to substitute analytic densities when they are known. As illustrated previously, the differential entropy for a Gaussian density has a closed-form analytic solution, and so similarly does the Kullback-Leibler divergence between two Gaussian densities. Let $p(\mathbf{x})$ and $q(\mathbf{x})$ be multivariate Gaussian distributions, defined as

$$p(\mathbf{x}) = p_g(\mathbf{x}; \mathbf{m}, \mathbf{P}) \quad (2.21a)$$

$$q(\mathbf{x}) = p_g(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Pi}). \quad (2.21b)$$

Substitution of $p(\mathbf{x})$ and $q(\mathbf{x})$ from Eq. (2.21) into the general Kullback-Leibler divergence in Eq. (2.20) yields

$$D_{KL}[q||p] = \frac{1}{2} \left[\log |\mathbf{P}\boldsymbol{\Pi}^{-1}| + \text{tr} \{ \mathbf{P}^{-1}\boldsymbol{\Pi} \} + (\boldsymbol{\mu} - \mathbf{m})^T \mathbf{P}^{-1} (\boldsymbol{\mu} - \mathbf{m}) - n_x \right], \quad (2.22)$$

where $\text{tr} \{ \cdot \}$ denotes the trace operator.

A point worth noting is the composition of the divergence in Eq. (2.22). The quantity consists of two different components: statistical distance due to the differences in size and shape of the covariances, and translational statistical distance due to differences in the first moment of the distributions (scaled by a covariance). The translational statistical distance measures how far the first moment of $q(\mathbf{x})$ is with respect to $p(\mathbf{x})$ (which is in the form of the squared Mahalanobis distance). This result enables some intuition to be established on the measure. For instance, given two identical covariances ($\mathbf{P} = \mathbf{\Pi}$), the Kullback-Leibler divergence increases monotonically with the squared Mahalanobis distance (i.e. with the distance between the first moments \mathbf{m} and $\boldsymbol{\mu}$); conversely, if the first moments between the two distributions are identical ($\mathbf{m} = \boldsymbol{\mu}$), the Kullback-Leibler divergence characterizes the differences in the uncertainty present in the two systems. The result is a scalar measure that indicates how dissimilar two statistical populations or densities are, a task that can be difficult in spaces even as simple as three-dimensional Euclidean space and that only gets more difficult as spaces of higher dimensionality are considered.

One drawback to the Kullback-Leibler divergence is that it is an asymmetric divergence. This does not prevent its employment in characterizing the difference between two pdfs, but it does imply the direction in which it is computed impacts the resulting quantification, as $D[p||q]$ might be quantifying something entirely different than $D[q||p]$. For instance, the reverse Kullback-Leibler divergence, given by

$$D_{KL}[p||q] = \frac{1}{2} \left[\log |\mathbf{\Pi}\mathbf{P}^{-1}| + \text{tr} \{ \mathbf{\Pi}^{-1}\mathbf{P} \} + (\mathbf{m} - \boldsymbol{\mu})^T \mathbf{\Pi}^{-1} (\mathbf{m} - \boldsymbol{\mu}) - n_x \right], \quad (2.23)$$

is clearly not equivalent to the expression in Eq. (2.22). This issue can be circumvented via the symmetrized Kullback-Leibler divergence, which is defined as

$$D_S[q||p] = \frac{1}{2} (D_{KL}[q||p] + D_{KL}[p||q]). \quad (2.24)$$

Substitution of the Kullback-Leibler divergence between two Gaussians in Eq. (2.22) and the reverse Kullback-Leibler in Eq. (2.23) into Eq. (2.24) yields

$$D_s[q||p] = \frac{1}{4} \left[\text{tr} \{ \mathbf{P}\mathbf{\Pi}^{-1} + \mathbf{\Pi}\mathbf{P}^{-1} \} + (\mathbf{m} - \boldsymbol{\mu})^T (\mathbf{P}^{-1} + \mathbf{\Pi}^{-1}) (\mathbf{m} - \boldsymbol{\mu}) - 2n_x \right]. \quad (2.25)$$

An interesting case for the symmetric Kullback-Leibler divergence is in the instance when the two covariances are identical. Let $\mathbf{P} = \mathbf{\Pi} = \mathbf{\Xi}$; then, it can be shown that the symmetric Kullback-Leibler divergence in Eq. (2.25) becomes

$$D_s[p||q] = \frac{1}{2} (\mathbf{m} - \boldsymbol{\mu})^T \mathbf{\Xi}^{-1} (\mathbf{m} - \boldsymbol{\mu}).$$

Interestingly, this is one half of the squared Mahalanobis distance between either mean with respect to the other distribution. This is a very specialized case, and will not occur when using this measure as a sensor tasking objective (as the *a priori* and *a posteriori* covariances will never be equivalent); however, it is an illustrative case that provides some insight into the nature of the symmetrized Kullback-Leibler divergence.

3. DYNAMICS AND MEASUREMENT MODELING

The previous section provides the tools for appropriately modeling a stochastic event, such as the position and velocity of an RSO, probabilistically. By modeling the target state with a pdf, the uncertainty in the true state is captured through the spread of the distribution, or through the covariance in a Gaussian distribution. The information entropy enables this spread to be quantified with a scalar measure, and the divergences enable the difference between different stochastic events to be measured. This section describes the dynamics model that govern the temporal evolution of the state estimate densities, as well as the sensor models used to observe the process.

3.1. DYNAMICAL MODELING

This work considers the determination of sensor tasking policies to estimate the state of an object that evolves according to the continuous-time dynamical system

$$\dot{\mathbf{x}}^{(i)}(t_k) = \mathbf{f}(\mathbf{x}^{(i)}(t_k)), \quad (3.1)$$

where $i \in \{1, 2, \dots, L\}$ are the indices for the L targets, $\mathbf{x}^{(i)}(t_k) \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$ is the state (i.e. position $\mathbf{r}_k^{(i)}$ and velocity $\mathbf{v}_k^{(i)}$) of the i^{th} target at time t_k , and $\mathbf{f} : \mathcal{X} \rightarrow \mathcal{X}$ represents the (potentially) nonlinear dynamics of the system and is assumed to describe the dynamics accurately, i.e. the equation does not account for mismodeling effects. Consequently, Eq. (3.1) omits process noise. While this is not an appropriate assumption for every application, its impact on RSO tracking is relatively small. For instance, when considering state independent process noise, targets over common time intervals accumulate the same

amount of process noise, affecting tasking decisions minimally. As a result, it is assumed that sufficient knowledge of the system dynamics is possessed and Eq. (3.1) is employed to consider their effects.

For conciseness, a simplification of notation will be utilized such that $\mathbf{x}_k^{(i)} \triangleq \mathbf{x}^{(i)}(t_k)$. The following subsections give the dynamic models leveraged in this work, but it is important to note that the sensor tasking policies developed in Sections 4 and 5 are not limited to these dynamics; the following discussions are given primarily for reference for later developments presented in Sections 4 and 5 and for the simulations presented in Section 6.

3.1.1. Clohessy-Wiltshire Model. In the event that the sensor to be tasked is a space-based sensor monitoring RSOs residing in similar orbits, the relative motion of the targets can be tracked as opposed to their inertial positions directly. This enables relative motion models to be employed in the tasking scheme. When the orbit of the space-based sensor is circular, a commonly employed relative motion model is the Clohessy-Wiltshire equations [88]. For a state defined as the relative position and velocity of the form $\mathbf{x}_k^{(i)} = [(\mathbf{r}_k^{(i)})^T, (\mathbf{v}_k^{(i)})^T]^T$, the dynamics of the state are given by the linear, discrete-time, noiseless system

$$\mathbf{x}_k^{(i)} = \mathbf{\Phi}(t_k, t_{k-1})\mathbf{x}_{k-1}^{(i)}, \quad (3.2)$$

where $\mathbf{\Phi}(t_k, t_{k-1})$ is the state transition matrix of the Clohessy-Wiltshire model, which is given as

$$\mathbf{\Phi}(t_k, t_{k-1}) = \begin{bmatrix} \mathbf{\Phi}_{rr} & \mathbf{\Phi}_{rv} \\ \mathbf{\Phi}_{vr} & \mathbf{\Phi}_{vv} \end{bmatrix}, \quad (3.3a)$$

where

$$\mathbf{\Phi}_{rr} = \begin{bmatrix} 4 - 3 \cos \psi & 0 & 0 \\ 6(\sin \psi - \psi) & 1 & 0 \\ 0 & 0 & \cos \psi \end{bmatrix} \quad (3.3b)$$

$$\Phi_{rv} = \begin{bmatrix} \frac{1}{n} \sin \psi & \frac{2}{n}(1 - \cos \psi) & 0 \\ \frac{2}{n}(\cos \psi - 1) & \frac{1}{n}(4 \sin \psi - 3\psi) & 0 \\ 0 & 0 & \frac{1}{n} \sin \psi \end{bmatrix} \quad (3.3c)$$

$$\Phi_{vr} = \begin{bmatrix} 3n \sin \psi & 0 & 0 \\ 6n(\cos \psi - 1) & 0 & 0 \\ 0 & 0 & -n \sin \psi \end{bmatrix} \quad (3.3d)$$

$$\Phi_{vv} = \begin{bmatrix} \cos \psi & 2 \sin \psi & 0 \\ -2 \sin \psi & -3 + 4 \cos \psi & 0 \\ 0 & 0 & \cos \psi \end{bmatrix}, \quad (3.3e)$$

and $\psi = n(t_k - t_{k-1})$, where n is the mean motion of the sensor's orbit. Note that the Clohessy-Wiltshire model is only one of many available relative motion models, and is only selected for its linear nature, enabling simple state propagation and facilitating faster generation of results and thus affording accessible analysis across more optimization solutions. Other options for relative motion propagation are available [28, 41, 53, 79] and would apply in this work as well.

3.1.2. Two-Body Problem. The Clohessy-Wiltshire relative motion model provides a simple set of linear dynamics that assist in evaluating the performance of the proposed tasking methods; the linearity enables high-dimensional optimization problems (e.g. numerous targets and/or sensors) to be explored to not only assist in analyzing the performance of the generated sensor schedules, but also in determining the computational feasibility of the developed solutions. However, the more common (and more applicable) scenario is in generating sensor schedules for ground-based observer stations to acquire data on RSOs in orbit about the Earth. As such, it is necessary to consider the motion of the targets with respect to the Earth. For the purposes of this work, the two-body model is

considered here, with the equations of motion given as

$$\mathbf{f}(\mathbf{x}_k^{(i)}) = \begin{bmatrix} \mathbf{r}_k^{(i)} \\ \mathbf{v}_k^{(i)} \end{bmatrix} = \begin{bmatrix} \mathbf{v}_k^{(i)} \\ -\frac{\mu \mathbf{r}_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^3} \end{bmatrix}, \quad (3.4)$$

with the corresponding Jacobian

$$\mathbf{F}_{k-1} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \\ \mathbf{G}_{k-1} & \mathbf{0}_{3 \times 3} \end{bmatrix}$$

where $\mathbf{I}_{3 \times 3}$ is a 3×3 identity matrix, $\mathbf{0}_{3 \times 3}$ is a 3×3 matrix of zeros, and

$$\mathbf{G}_{k-1} = -\frac{\mu}{\|\mathbf{r}_k^{(i)}\|^3} \begin{bmatrix} 1 - \frac{3(x_k^{(i)})^2}{\|\mathbf{r}_k^{(i)}\|^2} & -\frac{3x_k^{(i)}y_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^2} & -\frac{3x_k^{(i)}z_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^2} \\ -\frac{3x_k^{(i)}y_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^2} & 1 - \frac{3(y_k^{(i)})^2}{\|\mathbf{r}_k^{(i)}\|^2} & -\frac{3y_k^{(i)}z_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^2} \\ -\frac{3x_k^{(i)}z_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^2} & -\frac{3y_k^{(i)}z_k^{(i)}}{\|\mathbf{r}_k^{(i)}\|^2} & 1 - \frac{3(z_k^{(i)})^2}{\|\mathbf{r}_k^{(i)}\|^2} \end{bmatrix},$$

where μ is the gravitational constant and the scalar $x_k^{(i)}$, $y_k^{(i)}$, and $z_k^{(i)}$ are the individual position components for the i^{th} target at time t_k . Higher fidelity models can and have been leveraged with this work; however, employing the two-body model enables the use of analytic state transition matrices, such as those developed by Goodyear [30], Battin [9], and Der [21].

3.2. OBSERVER MODELING

In addition to modeling the dynamics, the process of generating measurements must be defined and the different facets of the measurement space, $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$, need to be considered. In this work, a measurement of target i at time t_k is generated according to the discrete-time process

$$\mathbf{z}_k^{(i)} = \mathbf{h}(\mathbf{x}_k^{(i)}) + \mathbf{v}_k^{(i)}, \quad (3.5)$$

where $\mathbf{v}_k^{(i)} \in \mathcal{Z}$ represents additive measurement noise, which is assumed to be zero mean, uncorrelated with the target state, and white with positive definite covariance \mathbf{R}_k ; and $\mathbf{h}(\cdot) : \mathcal{X} \rightarrow \mathcal{Z}$ is the (potentially) nonlinear measurement function.

3.2.1. Field of View, Field of Regard, and Probability of Detection. In many cases, the space in which a sensor can actually observe is a subset of the measurement space due to the geometry of the problem; this observable region is constrained further when considering limitations of the sensor. For instance, consider a ground-based observer station. The measurement space for the sensor is likely defined on the n_z -dimensional set of all numbers, but only a subset that correspond to states that can be observed by the sensor, defined as the sensor's *field of regard*. An even smaller portion of the field of regard is going to be observed at any given point in time, dictated by what the sensor can “see” at once, denoted as the sensor's *field of view*. In this ground-based observer scenario, the space above the horizon is referred to as the field of regard, with the field of view being the portion of the field of regard that the sensor covers in a single observation. Since this work is focused on scheduling and anticipating sensor actions, it is assumed that the field of regard and field of view are equivalent. This is an appropriate assumption so long as constraints are placed on the sensor schedule optimization that account for this, such as requiring enough time in between measurements that the sensor can be redirected in the correct pointing direction or allowing ample time for a long exposure to be taken.

In addition to the field of view and field of regard, there are other aspects of the actual measurement acquisition to address. Sensors are inherently imperfect and operate in environments that can impact the sensor's ability to detect a target or fail to guarantee that all measurements are target-generated. These issues are modeled by defining a probability of detection that indicates the probability a target will be observed in its current state, and a clutter model that accounts for spurious measurements that are not target generated. A common approach to modeling the probability of detection is to employ the “cookie cutter”

field of view model as [58]

$$p_D(\mathbf{x}_k^{(i)}) = \begin{cases} 1 & \text{target is within field of view} \\ 0 & \text{otherwise} \end{cases}$$

or in other words the target is detectable with unity probability when it lies within the field of view of the sensor and is undetectable otherwise. Measurement clutter is frequently modeled as uniformly distributed across the sensor field of view, with the number of returns following a Poisson distribution. However, for the purposes of this work, the clutter model does not add much in terms of the analysis of a sensor tasking policy; if one assumes a uniform clutter model, it has no impact on the tasking decision as each potential task is subject to the same clutter characteristics, and the influence other clutter model selections is entirely a function of the chosen model. As the focus of this work is on the objective measures themselves and not an exploration of modeling clutter, a clutter-free predicted ideal measurement set (PIMS) is assumed to be returned by the sensor. Lastly, this work also takes on the common assumption that at an instant in time each target can only generate a single measurement for a given observer, and that each measurement is generated by a single target.

3.2.2. Sensor Types. The aim of the analyses presented in Section 6 is not to only observe the decisions the sensor tasking policies make, but to determine if they are appropriately utilizing the measurement models and dynamics in executing the schedule optimization. In an effort to provide some semblance of intuition to the results, two different measurement models are employed in this work and are defined here for reference. Note that in these definitions, the target state $\mathbf{x}_k^{(i)}$ is assumed to be the target state relative to the observer.

3.2.2.1. Range and range rate observations. Range and range rate data can be acquired on a target from a variety of ways, one of the more common being through two-way radar range-Doppler measurements. In this implementation, a known signal is transmitted

and reflected off of an object back to the transmitter, enabling the range to be computed from round-trip light-time and the range-rate determined by the frequency shift present in the received signal [62]. This provides range data with a level of precision dependent upon the size of the smallest detectable phase differences in the signal, e.g. two clock cycles if a code-ranging technique is implemented. The measurement model is defined as

$$\mathbf{h}(\mathbf{x}_k^{(i)}) = \begin{bmatrix} \rho_k^{(i)} \\ \dot{\rho}_k^{(i)} \end{bmatrix} = \begin{bmatrix} \sqrt{(\mathbf{r}_k^{(i)})^T \mathbf{r}_k^{(i)}} \\ \frac{(\mathbf{v}_k^{(i)})^T \mathbf{r}_k^{(i)}}{\rho} \end{bmatrix}, \quad (3.6)$$

with the corresponding Jacobian

$$\mathbf{H}_k^{(i)} = \begin{bmatrix} \frac{(\mathbf{r}_k^{(i)})^T}{\rho_k^{(i)}} & \mathbf{0}_3 \\ \frac{(\rho_k^{(i)}(\mathbf{v}_k^{(i)})^T - \dot{\rho}_k^{(i)}(\mathbf{r}_k^{(i)})^T)}{\rho_k^{(i)}} & \frac{\mathbf{r}_k^{(i)T}}{\rho_k^{(i)}} \end{bmatrix}. \quad (3.7)$$

3.2.2.2. Right ascension and declination observations. Right ascension and declination data can be acquired from optical sensors, with a precision dependent upon the accuracy of the pointing direction of the sensor and the applied image processing techniques. The measurement model is defined as

$$\mathbf{h}(\mathbf{x}_k^{(i)}) = \begin{bmatrix} \alpha_k^{(i)} \\ \delta_k^{(i)} \end{bmatrix} = \begin{bmatrix} \text{atan2}(y_k^{(i)}, x_k^{(i)}) \\ \text{atan2}\left(z_k^{(i)}, \sqrt{(x_k^{(i)})^2 + (y_k^{(i)})^2}\right) \end{bmatrix}, \quad (3.8)$$

with corresponding Jacobian

$$\mathbf{H}_k^{(i)} = \begin{bmatrix} -\frac{y_k^{(i)}}{r_{xy}^2} & \frac{x_k^{(i)}}{r_{xy}^2} & 0 & 0 & 0 & 0 \\ -\frac{x_k^{(i)} z_k^{(i)}}{(\rho_k^{(i)})^2 r_{xy}} & -\frac{y_k^{(i)} z_k^{(i)}}{(\rho_k^{(i)})^2 r_{xy}} & -\frac{x_k^{(i)} y_k^{(i)}}{(\rho_k^{(i)})^2 r_{xy}} & 0 & 0 & 0 \end{bmatrix}, \quad (3.9)$$

where

$$r_{xy} = \sqrt{(x_k^{(i)})^2 + (y_k^{(i)})^2}.$$

3.3. LINEARIZATION

In this work, the ubiquitous extended Kalman filter is employed when executing the sensor tasking policies [43]. This filter requires that the nonlinear models be expanded via a first order Taylor series in order to perform the prediction and correction stages, requiring the Jacobians of the nonlinear dynamics and nonlinear measurement models. In particular, let

$$\mathbf{F}_k^{(i)} \triangleq \mathbf{F}(\mathbf{x}_k^{(i)}) = \left. \frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_k^{(i)}} \quad (3.10a)$$

$$\mathbf{H}_k^{(i)} \triangleq \mathbf{H}(\mathbf{x}_k^{(i)}) = \left. \frac{\partial \mathbf{h}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}_k^{(i)}}, \quad (3.10b)$$

where $\mathbf{F}_k^{(i)}$ is the Jacobian of the dynamics model evaluated at the target state, $\mathbf{x}_k^{(i)}$, and $\mathbf{H}_k^{(i)}$ is the measurement model Jacobian evaluated at the target state relative to the observer. The Jacobians for the dynamics and measurement models in the preceding sections are necessary for this work. It is worth noting that the work presented in this dissertation does not preclude the use of other linearization methods; the choice of utilizing the first order Taylor series expansion is in the interest of computational burden as well the compact analytic solutions the approach affords.

4. SINGLE TARGET SENSOR TASKING

The preceding section has laid out some of the basic principles to approach the problem of sensor tasking. The probability theory provides a framework in which states can be probabilistically represented, while the information theory developments provide an array of tools for analyzing the state estimate densities in a filter and, as it will be shown in this section, assist in quantifying the potential information contribution of candidate observation times. Before the full multitarget sensor tasking problem is considered, it behooves one to study the single-target scenario.

On the surface, sensor tasking when attempting to track a single target may seem unnecessary. If there is only a single object to observe, there is not much else for a sensor to be tasked with. There are a few reasons that make this problem pertinent. First, in handling the single-target sensor tasking problem appropriately, the object can be tracked sufficiently with the fewest number of measurements possible, freeing up the sensor(s) to perform other tasks, such as searching for new, uncatalogued objects. Additionally, from an academic perspective, studying a sensor tasking policy in the presence of one target allows the methodology to be developed more completely and enables the behavior of a particular approach or objective to be evaluated, providing insight as to what segments of the observable trajectory the solution deems information-rich and warrants occupation of the limited sensor resources. However, in order to begin examining this measure, it is necessary to have an understanding of how the measurements are processed.

In the interest of providing context to the use of information theory in sensor tasking, this section opens with the celebrated Kalman filter in Section 4.1, as well as a few of its variants that are leveraged in the subsequent developments. Section 4.2 goes on to detail the use of the Kullback-Leibler (KL) divergence in sensor tasking. The conventional employment is discussed first in Section 4.2.1, presented with a statistical analysis of the

true KL divergence to investigate the assumptions that conventional KL sensor tasking requires, as well as the construction and analysis of an objective function that enables its use in considering sets of measurements. This is followed by the novel generalization of the myopic measure to a forecasted KL divergence in Section 4.2.2 along with a simple example illustrating the connection between the myopic and forecasted KL divergences as well as the differences in their performance. The section concludes with some comments on the optimization of the measures as well as their computation in other vector spaces.

4.1. SINGLE TARGET FILTERING APPROACHES

Consider a stochastic target state $\mathbf{x}_k \in \mathcal{X} \subseteq \mathbb{R}^{n_x}$ composed of a position vector \mathbf{r}_k and a velocity vector \mathbf{v}_k defined on the state space with temporal evolution governed by the linear dynamical system in Eq. (3.2). In addition to the dynamical process, measurements are generated according to Eq. (3.5). As this section is focused on the single-target scenario, the superscripts denoting the target index that were maintained in the previous section are dropped here to simplify notation. In order to initialize the filter, it is assumed that the target state is given an initial mean and covariance as

$$\mathbf{m}_0 = \mathbb{E} \{ \mathbf{x}_0 \} , \quad (4.1a)$$

$$\mathbf{P}_0 = \mathbb{E} \{ (\mathbf{x}_0 - \mathbf{m}_0)(\mathbf{x}_0 - \mathbf{m}_0)^T \} . \quad (4.1b)$$

4.1.1. The Kalman Filter. The minimum mean square error (MMSE) filter can be divided into two stages: a prediction stage in which the first two moments of the state density are propagated through time according to the dynamics, and a correction stage in which a noisy measurement is processed and the first two moments of the state density are updated accordingly. In the case of linear dynamics and a linear measurement model, the resulting predictor/corrector framework is the Kalman filter [42, 43].

4.1.1.1. Predictor. The first two moments of the target state density can be propagated forward in time without the processing of observation data to generate an *a priori* or a prior state estimate density. These two moments are generated by simply taking the expected value of Eq. (3.2) and by computing the second central moment, as

$$\mathbf{m}_k = \mathbb{E} \{ \mathbf{x}_k \} \quad (4.2a)$$

$$\mathbf{P}_k = \mathbb{E} \{ (\mathbf{x}_k - \mathbf{m}_k)(\mathbf{x}_k - \mathbf{m}_k)^T \} . \quad (4.2b)$$

Due to the linear dynamics, the transition matrix $\Phi(t_k, t_{k-1})$ is deterministic; the computation of the *a priori* mean yields

$$\begin{aligned} \mathbf{m}_k &= \mathbb{E} \{ \Phi(t_k, t_{k-1}) \mathbf{x}_{k-1} \} \\ &= \Phi(t_k, t_{k-1}) \boldsymbol{\mu}_{k-1}, \end{aligned} \quad (4.3)$$

where $\boldsymbol{\mu}_{k-1}$ is the *a posteriori* mean at the previous time step. In order to determine the *a priori* covariance, it is useful to define the state error at time t_k to be

$$\begin{aligned} \mathbf{e}_k &= \mathbf{x}_k - \mathbf{m}_k, \\ &= \Phi(t_k, t_{k-1}) \mathbf{x}_{k-1} - \Phi(t_k, t_{k-1}) \mathbf{m}_{k-1}. \end{aligned} \quad (4.4)$$

Again the linear system dynamics enables a useful simplification that allows the error at time t_k to be represented as a propagated error from the previous time step, as

$$\mathbf{e}_k = \Phi(t_k, t_{k-1}) \mathbf{e}_{k-1}. \quad (4.5)$$

Equation (4.5) can be substituted into Eq. (4.2b)

$$\begin{aligned} \mathbf{P}_k &= \mathbb{E} \{ \mathbf{e}_k \mathbf{e}_k^T \} , \\ &= \mathbb{E} \{ \mathbf{\Phi}(t_k, t_{k-1}) \mathbf{e}_{k-1} \mathbf{e}_{k-1}^T \mathbf{\Phi}^T(t_k, t_{k-1}) \} . \end{aligned}$$

Let $\mathbf{\Pi}_{k-1}$ be the *a posteriori* covariance at time t_{k-1} ; the covariance propagation becomes

$$\mathbf{P}_k = \mathbf{\Phi}(t_k, t_{k-1}) \mathbf{\Pi}_{k-1} \mathbf{\Phi}^T(t_k, t_{k-1}) . \quad (4.6)$$

Equations (4.3) and (4.6) together comprise the prediction stage of the Kalman filter, enabling the propagation of the first two moments of the state density according to the linear system dynamics.

4.1.1.2. Corrector. The prediction stage can be used recursively to continuously propagate the mean and covariance in the absence of new information. In order to process observations to refine the state estimate, a corrector stage must be implemented. Given an *a priori* state density with mean \mathbf{m}_k and covariance \mathbf{P}_k , consider an update that is a linear function of the measurement data \mathbf{z}_k , as

$$\boldsymbol{\mu}_k = \mathbf{a}_k + \mathbf{K}_k \mathbf{z}_k , \quad (4.7)$$

where $\boldsymbol{\mu}_k$ is defined to be the *a posteriori* mean, \mathbf{a}_k is a vector to be computed, and \mathbf{K}_k is a matrix to be computed. Define the posterior state estimation error to be

$$\boldsymbol{\varepsilon}_k \triangleq \mathbf{x}_k - \boldsymbol{\mu}_k .$$

In the presence of linear dynamics and additive process and measurement noise, minimization of the mean square error of the posterior state estimate is achieved via the Kalman update equations, given as

$$\boldsymbol{\mu}_k = \mathbf{m}_k + \mathbf{K}_k(\mathbf{z}_k - \hat{\mathbf{z}}_k) \quad (4.8a)$$

$$\boldsymbol{\Pi}_k = (\mathbf{I}_{n_x} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_k \quad (4.8b)$$

where $\hat{\mathbf{z}}_k$ is the expected measurement, defined as

$$\hat{\mathbf{z}}_k = \mathbf{h}(\mathbf{m}_k);$$

\mathbf{K}_k is the deterministic Kalman gain, which is defined to be

$$\begin{aligned} \mathbf{K}_k &= \mathbf{P}_k\mathbf{H}_k^T(\mathbf{H}_k\mathbf{P}_k\mathbf{H}_k^T + \mathbf{R}_k)^{-1} \\ &= \mathbf{P}_k\mathbf{H}_k^T(\mathbf{W}_k)^{-1}; \end{aligned} \quad (4.9)$$

and \mathbf{W}_k is the innovation (or residual) covariance. The linear gain given in Eq. (4.9) is optimal in the sense of minimizing the mean square error of the state estimate. Note that Eq. (4.8b) requires the optimal gain; as this work employs the optimal gain, this form of the covariance update is used in certain developments as it provides compact, simplified expressions. However, this does not preclude other forms of computing the update, so long as the optimal gain is used.

4.1.2. The Extended Kalman Filter. Due to the nonlinear dynamics, the Kalman filter is not directly applicable to the problem at hand when considering two-body motion. It is necessary to linearize the system dynamics and employ an extended Kalman filter (EKF) in much of this work. The predictor of the Kalman filter in Eq. (4.8) can then be extended

to the EKF predictor as [42]

$$\mathbf{m}_k = \mathbf{f}(\boldsymbol{\mu}_{k-1}) \quad (4.10a)$$

$$\mathbf{P}_k = \boldsymbol{\Phi}(t_k, t_{k-1}) \boldsymbol{\Pi}_{k-1} \boldsymbol{\Phi}^T(t_k, t_{k-1}), \quad (4.10b)$$

where $\boldsymbol{\Phi}(t_k, t_{k-1})$ is the state transition matrix that can be either approximated analytically if an approximation exists or propagated numerically via

$$\boldsymbol{\Phi}(t, t_{k-1}) = \int_{t_{k-1}}^{t_k} \mathbf{F}(\mathbf{x}^*(t)) \boldsymbol{\Phi}(t, t_{k-1}) dt, \quad (4.11)$$

where $\mathbf{F}(\cdot)$ is the Jacobian of the dynamics as defined in Eq. (3.10a), $\mathbf{x}^*(t)$ is a variable denoting *a priori* mean propagated to time t , and with the initial condition

$$\boldsymbol{\Phi}(t_{k-1}, t_{k-1}) = \mathbf{I}_{n_x \times n_x}.$$

Similar to the prediction stage, the update stage in Eq. (4.8) becomes [42]

$$\boldsymbol{\mu}_k = \mathbf{m}_k + \mathbf{K}_k(\mathbf{z}_k - \hat{\mathbf{z}}_k) \quad (4.12a)$$

$$\boldsymbol{\Pi}_k = \mathbf{P}_k - \mathbf{P}_k \mathbf{H}_k^T(\mathbf{m}_k) \mathbf{K}_k^T - \mathbf{K}_k \mathbf{H}_k(\mathbf{m}_k) \mathbf{P}_k + \mathbf{K}_k \mathbf{W}_k \mathbf{K}_k^T, \quad (4.12b)$$

where $\hat{\mathbf{z}}_k$ is the expected measurement, $\mathbf{H}(\cdot)$ is the Jacobian of the measurement model as defined in Eq. (3.10b), and \mathbf{K}_k is the (now random) Kalman gain, defined similarly to Eq. (4.9) as

$$\mathbf{K}_k = \mathbf{P}_k \mathbf{H}(\mathbf{m}_k)^T (\mathbf{W}_k)^{-1}. \quad (4.13)$$

The update given in Eq. (4.12b) is the covariance update for any linear gain, with the Kalman gain being the linear gain that (approximately) minimizes the mean square error of the posterior state estimate. As alluded to previously, other forms of the covariance update exist and afford different ways of representing the correction stage. Due to their usefulness in later developments, a brief discussion of some alternatives is warranted.

4.1.2.1. The information update. In the developments to follow, it is useful to consider the accumulation of candidate observation times and the information that they provide. A form of the Kalman update that is more conducive to this accumulation comes from the information filter and is referred to as the information form of the covariance update. This update is given as [81]

$$\mathbf{\Pi}_k^{-1} = \mathbf{P}_k^{-1} + \mathbf{H}_k^T \mathbf{R}_k^{-1} \mathbf{H}_k. \quad (4.14)$$

Note that the argument for the Jacobian $\mathbf{H}_k(\cdot)$ has been dropped here to simplify notation, and its inclusion will be implied herein. For reasons that will be evident later in this section, this form of the covariance update is particularly useful when considering measurements from multiple sources, be it from different sensors or different points in time.

4.1.2.2. The batch update. The aforementioned forms of the Kalman filter are successful in recursively updating the mean and covariance as new information is acquired and consequently, in the presence of ample data, is robust against changes in trajectory due to maneuvers or perturbations as the estimate is updated in real-time. However, in order to better consider the impact of multiple measurements, a batch processor must be considered.

Given an observation at time t_k , define a mapped measurement model Jacobian to be

$$\tilde{\mathbf{H}}_k = \mathbf{H}_k \mathbf{\Phi}(t_k, t_{\text{ref}}),$$

where t_{ref} is a reference time at which a measurement taken at time t_k is processed, $\tilde{\mathbf{H}}_k$ denotes that \mathbf{H}_k is mapped in time for a single measurement, and $\Phi(t_k, t_{\text{ref}})$ is the state transition matrix from time t_{ref} to time t_k that is defined either analytically or numerically propagated similar to Eq. (4.11). This allows an observation at any time to be processed at an arbitrary time of interest, providing a means of accounting for the influence measurements at different points in time have on a state estimate at time t_{ref} . Further, it enables these influences to be considered simultaneously. Given m_z measurements at times t_1, t_2, \dots, t_{m_z} , construct concatenated $\tilde{\mathbf{H}}$ and $\bar{\mathbf{R}}$ matrices as

$$\tilde{\mathbf{H}} = \begin{bmatrix} \tilde{\mathbf{H}}_1 \\ \tilde{\mathbf{H}}_2 \\ \vdots \\ \tilde{\mathbf{H}}_{m_z} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_1 \Phi(t_1, t_{\text{ref}}) \\ \mathbf{H}_2 \Phi(t_2, t_{\text{ref}}) \\ \vdots \\ \mathbf{H}_{m_z} \Phi(t_{m_z}, t_{\text{ref}}) \end{bmatrix} \quad \text{and} \quad \bar{\mathbf{R}} = \begin{bmatrix} \mathbf{R}_1 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_2 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{R}_{m_z} \end{bmatrix}, \quad (4.15)$$

where $\bar{\mathbf{R}}$ is simply a block diagonal consisting of the m_z individual measurement noise covariance matrices; note that it is block diagonal since the measurement noise is taken to be uncorrelated in time. Additionally, note that this formulation does not require \mathbf{H}_i to originate from the same measurement model, nor the measurement noise covariances \mathbf{R}_i to be constant or identical, allowing different observation sources and observation times to be jointly considered. The batch update for the covariance is given as [81]

$$\mathbf{\Pi}_{\text{ref}}^{-1} = \mathbf{P}_{\text{ref}}^{-1} + \tilde{\mathbf{H}}^T \bar{\mathbf{R}}^{-1} \tilde{\mathbf{H}}, \quad (4.16)$$

where \mathbf{P}_{ref} is the *a priori* covariance propagated to the reference time t_{ref} .

4.2. INFORMATION DIVERGENCE AS A TASKING OBJECTIVE

With a filtering framework established, the information theoretic measures presented in Section 2 can be leveraged in developing a sensor tasking methodology. A typical information-based sensor tasking approach evaluates candidate observation times by considering the potential information divergence a measurement would induce between the *a priori* state estimate density and an approximate *a posteriori* state estimate density, providing a relative measure of the impact an observation may generate. A common selection in these tasking policies is the Kullback-Leibler (KL) divergence.

4.2.1. Myopic Kullback-Leibler Divergence. The KL divergence in Eq. (2.22) provides the most general form of the KL divergence between two Gaussians, yielding a directed distance between the two densities. In a filtering framework, it is common to employ information divergences in quantifying the distance between the *a priori* and *a posteriori* state densities; in doing so, the amount of information ingested by the filter can be quantified. Assumptions are generally made in order to employ tractable filters to specific target tracking problems, specifically that the measurement noise is white, zero-mean, additive, and is not correlated with the state. Additionally, the prior covariance and the measurement model Jacobian are taken to be deterministic, enabling the use of the EKF update given in Eqs. (4.10) and (4.12). Taking on these assumptions and relating the two distributions through the EKF update can provide further simplifications to Eq. (2.22) and provide more insight to the divergence measure as a tasking objective.

Consider the case of an object with true state at time t_k defined as

$$\mathbf{x}_k = \begin{bmatrix} \mathbf{r}_k \\ \mathbf{v}_k \end{bmatrix},$$

where \mathbf{r}_k and \mathbf{v}_k are $\frac{n_x}{2}$ -dimensional vectors describing the target's position and velocity, respectively. The target state density is taken to be Gaussian, and let $p(\mathbf{x})$ and $q(\mathbf{x})$ denote the *a priori* and *a posteriori* densities, defined as

$$p(\mathbf{x}_k) = p_g(\mathbf{x}_k; \mathbf{m}_k, \mathbf{P}_k) \quad \text{and} \quad q(\mathbf{x}_k) = p_g(\mathbf{x}_k; \boldsymbol{\mu}_k, \boldsymbol{\Pi}_k). \quad (4.17)$$

The *a priori* mean and covariance can be partitioned as

$$\mathbf{m}_k = \begin{bmatrix} \hat{\mathbf{r}}_k \\ \hat{\mathbf{v}}_k \end{bmatrix} \quad \text{and} \quad \mathbf{P}_k = \begin{bmatrix} \mathbf{P}_{rr} & \mathbf{P}_{rv} \\ \mathbf{P}_{vr} & \mathbf{P}_{vv} \end{bmatrix},$$

where the $\hat{\cdot}$ notation indicates the estimated position ($\hat{\mathbf{r}}_k$) and velocity ($\hat{\mathbf{v}}_k$) at time t_k , \mathbf{P}_{rr} is the estimated position covariance, \mathbf{P}_{vv} is the estimated velocity covariance, and $\mathbf{P}_{rv} = \mathbf{P}_{vr}^T$ is the estimated position-velocity cross-covariance. Similarly, the *a posteriori* state density is also taken to be Gaussian with mean $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Pi}_k$, constructed similarly as \mathbf{m}_k and \mathbf{P}_k . The aforementioned assumptions enable the use of an EKF, relating the parameters of the two distributions in Eqs. (4.17) through the Kalman update equations in Eqs. (4.12).

Substitution of the densities in Eq. (4.17) and utilizing the EKF update equations in Eqs. (4.12), the KL divergence describing the strength of an update can be expressed in the form

$$D_{KL}[q||p] = \frac{1}{2} \left[\log |\mathbf{P}_k(\boldsymbol{\Pi}_k)^{-1}| + \text{tr} \{ (\mathbf{P}_k)^{-1} \boldsymbol{\Pi}_k \} + (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k))^T \mathbf{K}_k^T (\mathbf{P}_k)^{-1} \mathbf{K}_k (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k)) - n_x \right]. \quad (4.18)$$

This compact expression allows the impact of an observation on the state estimate density to be computed, simultaneously considering the measurement model and the uncertainty in the state estimate itself. Note that the first two terms are entirely concerned with the uncertainties present in both state estimate densities.

Consider the logarithmic term in Eq. (4.18); this term can be expressed instead as the difference of the logarithm of their respective determinants, as

$$\frac{1}{2} \log |\mathbf{P}_k(\mathbf{\Pi}_k)^{-1}| = \frac{1}{2} \log |\mathbf{P}_k| - \frac{1}{2} \log |\mathbf{\Pi}_k|,$$

or equivalently,

$$\begin{aligned} \frac{1}{2} \log |\mathbf{P}_k(\mathbf{\Pi}_k)^{-1}| &= \frac{1}{2} \log |2\pi e \mathbf{P}_k| - \frac{1}{2} \log |2\pi e \mathbf{\Pi}_k|, \\ &= H[p] - H[q]. \end{aligned}$$

Evidently, the first term in Eq. (4.18) accounts for the entropy reduction from the *a priori* density to the *a posteriori* density, a primary element of concern for the sensor tasking scheme. However, from Eq. (2.20), the KL divergence is constructed of more than the volumetric uncertainty reduction.

Consider the second term of Eq. (4.18) coupled with the subtraction of the dimension of the state; utilizing Eq. (4.12b) and the cyclic permutation property of the trace operator, this can be represented as

$$\begin{aligned} \frac{1}{2} \left(\text{tr} \{ \mathbf{P}_k^{-1} \mathbf{\Pi}_k \} - n_x \right) &= \frac{1}{2} \left(\text{tr} \{ \mathbf{P}_k^{-1} (\mathbf{I}_{n_x} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_k \} - n_x \right) \\ &= -\frac{1}{2} \text{tr} \{ \mathbf{K}_k \mathbf{H}_k \}. \end{aligned} \quad (4.19)$$

While the exact interpretation of Eq. (4.19) may be difficult to decipher, there are some observations to be made. By the definition of the KL divergence, this must be a portion of the cross-entropy between the two state estimate densities, specifically the portion that is independent of the translational shift induced by the observation \mathbf{z}_k (i.e. the third term in Eq. (4.18)). It is evident that the matrix remaining in the trace operator in Eq. (4.19) is the matrix responsible for manipulating the *a priori* uncertainty. This trace results in a non-negative value, as it is the sum of the eigenvalues of the product of the Kalman gain

and the measurement model Jacobian; these eigenvalues are necessarily non-negative, as the product of two matrices is positive semi-definite. Specifically, through the definition of the Kalman gain given in Eq. (4.9), this product can be expressed as

$$\mathbf{K}_k \mathbf{H}_k = \mathbf{P}_k \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k,$$

where \mathbf{W}_k is the innovation (or residual) covariance as in Eq. (4.9). If \mathbf{R}_k is assumed to be positive definite, then $\mathbf{W}_k > 0$ as is its inverse, and $\mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k$ is positive semi-definite (denoted $\mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \geq 0$). Since \mathbf{P}_k is taken to be positive definite, the product $\mathbf{K}_k \mathbf{H}_k$ will consist of non-negative eigenvalues and $\text{tr} \{ \mathbf{K}_k \mathbf{H}_k \} \geq 0$.

At this point it is useful to consider two reductive schematic scenarios for illustrative purposes: (i) the case in which the mean remains stationary after processing \mathbf{z}_k and no uncertainty contraction occurs and (ii) the case in which the mean remains stationary and the uncertainty is morphed or altered in some way other than a pure rotation that preserves its volume but changes its container. Case (i) is a very limited case, as

$$\mathbf{I}_{n_x} - \mathbf{K}_k \mathbf{H}_k = \mathbf{\Pi}_k \mathbf{P}_k^{-1},$$

and since the two covariance matrices are real positive definite matrices, their product is also be real positive definite [40]. If a rotation matrix is real and positive definite, it must correspond to a rotation that is an integer multiple of a complete rotation. Consequently, the trace is simply equal to the dimension n_x , which will result in a zero net contribution to the divergence, an intuitive result, as rotations of this nature will effectively leave the uncertainty unchanged.

As for Case (ii), a volume preserving transformation necessarily has a determinant equal to unity, but has no such constraint on the trace. As previously mentioned, the matrix within the trace in Eq. (4.19) has non-negative eigenvalues; consequently, Eq. (4.19) is less than or equal to zero, with equality only when the sum of the eigenvalues is zero. Since the

eigenvalues of the product of the Kalman gain and the measurement model Jacobian has non-negative eigenvalues, the equality case corresponds to the scenario in which $\mathbf{K}_k \mathbf{H}_k = \mathbf{0}$. This implies that, by the formal definition of information, in the absence of translational motion of the first moment of the densities *any* volume-preserving alterations to the uncertainty will result this term reducing the divergence.

4.2.1.1. Statistical analysis. The expression in Eq. (4.18) produces a quantification of the strength of an update, providing a complete and mathematically rigorous characterization of the information absorbed through the acquisition of an observation under the stated assumptions. By relating the *a priori* and *a posteriori* densities through a Kalman update, much of the KL divergence definition in Eq. (4.18) is deterministic, resulting in a term dealing explicitly with the covariance reduction and a leaving a single stochastic term, namely the translational element of the update that moves the mean. This statistical distance is quadratic in the measurement residual, defined as

$$\mathbf{y}_k = \mathbf{z}_k - \mathbf{h}(\mathbf{m}_k). \quad (4.20)$$

Considering the *a priori* pdf is taken to be Gaussian and the measurement is assumed to be Gaussian distributed about the truth, the measurement residual is also Gaussian distributed as

$$\mathbf{y}_k \sim p_g(\mathbf{y}_k; \mathbf{0}, \mathbf{W}_k).$$

This is a useful result, as it is well-known that the quadratic form of an n -dimensional zero-mean normal random variable \mathbf{a} is chi-square distributed, i.e.

$$\mathbf{a}^T \mathbf{A}^{-1} \mathbf{a} \sim p_{\chi^2}(\mathbf{a}^T \mathbf{A}^{-1} \mathbf{a}; n),$$

where \mathbf{A} is the covariance matrix of \mathbf{a} and the tilde notation denotes the quantity is distributed according to a chi-square density of n degrees of freedom. While the stochastic term in Eq. (4.18) is nearly in this form, the matrix $\mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k$ is not the residual covariance.

However, it is possible to apply a generalization of the theorem that a quadratic form of a set of normal random variables is distributed according to a chi-square distribution [7, 76]. To accomplish this, first define the stochastic quantity ξ_k as

$$\xi_k = \mathbf{y}_k^T \mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k \mathbf{y}_k. \quad (4.21)$$

Consider a $n_z \times n_z$ non-singular matrix \mathbf{M}_k that satisfies the conditions

$$\mathbf{M}_k^T \mathbf{W}_k \mathbf{M}_k = \mathbf{I}_{n_z} \quad \text{and} \quad \mathbf{M}_k^T \mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k \mathbf{M}_k = \mathbf{\Lambda}_k,$$

where $\mathbf{\Lambda}_k$ is a diagonal matrix consisting of the eigenvalues corresponding to the eigenvectors of \mathbf{M}_k ; more specifically, \mathbf{M}_k and $\mathbf{\Lambda}_k$ are constructed from the s distinct solutions (namely, the eigenvalues λ_i and the associated eigenvectors \mathbf{v}_i) of the equation

$$\mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k \mathbf{W}_k \mathbf{v}_i = \lambda_i \mathbf{v}_i \quad \text{for} \quad i = 1, 2, \dots, s. \quad (4.22)$$

By the definition of covariance matrices and the update provided by the EKF, it is guaranteed that the covariance matrices \mathbf{P}_k and \mathbf{W}_k are symmetric, positive definite; thus the matrix $\mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k$ is symmetric, positive semi-definite. Leveraging this fact, the product of the two matrices $\mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k$ and \mathbf{W}_k has the eigendecomposition

$$\mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k \mathbf{W}_k = \sum_{i=1}^s \lambda_i \mathbf{E}_i, \quad 0 \leq \lambda_i \in \mathbb{R}, \quad (4.23)$$

where the λ_i are the s distinct solutions to Eq. (4.22),

$$\mathbf{E}_i = \mathbf{M}_k^{-T} \mathbf{B}_i \mathbf{M}_k^T,$$

and \mathbf{B}_i is the $n_z \times n_z$ matrix consisting of a single nonzero entry of 1 where $\mathbf{\Lambda}_k$ has element λ_i . Using this eigendecomposition, the distribution of the stochastic quantity in Eq. (4.21) can be expressed as a sum of chi-square distributed random variables as [7]

$$p(\xi_k) = p\left(\sum_{i=1}^s \lambda_i \gamma_i\right), \quad (4.24)$$

where γ_i follows a chi-square distribution with r_i degrees of freedom and r_i is the rank of matrix \mathbf{E}_i . Specifically, the quantity ξ_k is distributed identically to a weighted sum of central chi-square random variables. It is important to note that the zero-mean measurement residual produces these central chi-square components; if the mean is nonzero (i.e. if there is a bias present in the measurement residuals), the γ_i are non-central chi-square distributed. Substituting Eq. (4.21) into Eq. (4.18) and leveraging the covariance update in Eq. (4.12b), it can be shown that

$$D_{KL}[q||p] = \frac{1}{2} \left[\xi_k - \log |\mathbf{I}_{n_x} - \mathbf{K}_k \mathbf{H}_k| - \text{tr} \{ \mathbf{K}_k \mathbf{H}_k \} \right], \quad (4.25)$$

i.e. the KL divergence is distributed as a mean-shifted weighted sum of chi-square random variables.

If the chi-square distributed random variables on the right-hand side of Eq. (4.24) are independent, their sum is also chi-square distributed. This requires the matrix on the left-hand side of Eq. (4.23) to be idempotent [31], which can be easily shown to be false by contradiction when \mathbf{K}_k is the Kalman gain; assume the *a priori* covariance and the measurement model to be identity ($\mathbf{P}_k = \mathbf{H}_k = \mathbf{I}_{n_x}$). The left-hand side of Eq. (4.23) becomes

$$\mathbf{K}_k^T \mathbf{P}_k^{-1} \mathbf{K}_k \mathbf{W}_k = (\mathbf{W}_k)^{-1}. \quad (4.26)$$

For the matrix in Eq. (4.26) to be idempotent, it must be case that multiplication with itself does not change the matrix, i.e. the product of the inverse innovation covariance matrix with itself yields the inverse innovation covariance. Clearly this is not the case, and while this is a contrived example, it is straightforward to verify numerically in other instances in which $\mathbf{P}_k, \mathbf{H}_k \neq \mathbf{I}_{n_x}$ that this holds true. Even if the γ_i are assumed independent, an analytic closed-form solution for the distribution of a weighted sum of chi-square random variables is not available, though if desirable, different numerical approaches exist for solving for the distribution [11, 13, 26, 63, 77]. However, in many instances, it may be desirable or more appropriate to instead approximate the value of the KL divergence in the Kalman filter rather than the full distribution itself. This is frequently done through expectation hedging, or computing the first moment of the divergence [58].

4.2.1.2. Conventional employment. The result in Eq. (4.25) illustrates that the directed distance between *a priori* and *a posteriori* pdfs can be decomposed into a sum of at most n_z chi-square random variables plus a constant offset [33], implying that for scalar measurements the quantity is necessarily non-central chi-square distributed. Further analysis illustrates utilization of the first moment of Eq. (4.18) can be sufficient in approximating the KL divergence for vector measurements in a Kalman filter [33]; use of this first moment has been accepted as the conventional approach when utilizing information divergences, and has been shown to be effective in sensor tasking applications [20, 27, 49, 50]. This first moment, oftentimes referred to as the *information gain*, can be computed from Eq. (4.18) by taking its expected value with respect to the measurement pdf, $p(\mathbf{z}_k)$, giving

$$\begin{aligned} M_{KL}(t_k) &= \int_{\mathcal{Z}} D_{KL}[q||p]p(\mathbf{z}_k)d\mathbf{z}_k \\ &= \frac{1}{2} \log |\mathbf{P}_k(\mathbf{\Pi}_k)^{-1}|, \end{aligned} \quad (4.27)$$

where $\mathcal{Z} \subset \mathbb{R}^{n_z}$ is the n_z -dimensional measurement domain and the argument t_k denotes the time at which measurement z_k is acquired and processed to update the prior covariance \mathbf{P}_k to obtain the posterior covariance $\mathbf{\Pi}_k$; note that while the time argument does not explicitly appear in Eq. (4.27), the notation is employed to clarify the mathematical developments in this work. Notice that this is identical to the first term in Eq. (4.18); specifically, Eq. (4.27) is exactly the difference in the entropies of the two densities. Additionally, note that this is equivalent to the mutual information [3].

Equation (4.27) gives the expected KL divergence between the *a priori* and *a posteriori* state densities in a Kalman filter assuming the state densities are modeled as Gaussians, providing a scalar measure of the impact an observation has on the target state estimate at the measurement time. Note that the first moment of the KL divergence in Eq. (4.18) averages out the information contribution brought on by the cross-entropy between the two distributions and leaves just a quantification of the entropic differences, a facet that increases its attractiveness as a sensor tasking objective. Due to the emphasis on uncertainty reduction at the observation time, this approach is referred to here as a “myopic” sensor tasking objective; however, the use of the word “myopic” is not intended to imply poor performance, but to signify the quantification of the uncertainty reduction at one moment in time rather than a downstream (or upstream) time. Use of the first moment of the KL divergence has been studied and shown to perform well as a maximization objective in sensor tasking routines [20, 33].

Similar to Eq. (4.18), the expression in Eq. (4.27) illustrates the benefits of utilization of the first moment of the KL divergence. The measure encapsulates the effects due to the employed measurement model of processing an observation at time t_k while remaining focused on uncertainty reduction, as it is defined based on the ratio of the *a priori* and *a posteriori* covariances. Consequently, larger reductions in the determinant of the state estimate covariance are realized as larger values of expected information divergence measures between the two pdfs. This formulation is particularly useful, as it is not necessary

to task a sensor and generate a measurement for its computation, enabling measurements of potentially differing type and quality to be compared on equal footing, even if they are acquired from different observer locations, without the need to actually task a sensor. This comparison is valid so long as the potential observations being compared occur at a common time t_k .

As previously stated, the first moment of the KL divergence has been utilized in many applications and has proven to be a successful sensor tasking objective. However, in the current formulation, there is no straightforward way to consider the impact of multiple measurements on the expected divergence; the individual contributions of m_z measurements can be accumulated, but this is an ad hoc approach – the information is not necessarily additive in that way, suggesting that the resulting quantity is a step removed from what it is intending to portray. Furthermore, there are instances of undesirable behavior or characteristics that should be addressed.

For instance, consider the case where the measurement noise covariance is independent of the target state and constant through time. For a non-conservative dynamical system in which the state estimate uncertainty grows in time, the result is that a larger KL divergence first moment is obtained the longer an observation acquisition is delayed. This tendency is referred to here as the “wait and see” mentality, and while it is not necessarily an incorrect policy under which to operate, it is undesirable for autonomous scheduling, as it will always suggest to delay tasking decisions until the final possible observation time. Similarly, in conservative systems, the entropy is preserved as the determinant of the covariance remains constant [83]. In some instances this results in the uncertainties present to move into different channels of the state in a periodic manner, yielding periodic optimal observation times. Again, while not necessarily incorrect, these times are maximizing the information expected to be gained at the instant of the observation and do not necessarily reflect how the observation might impact the uncertainty at a later time. This does not assist in scheduling other observations either before or after the fact. It is of interest to develop an

approach that avoids this trivial scheduling solution of waiting in some non-conservative systems, and enable the scheduling in both non-conservative and conservative systems in such a way that the impact each measurement in a set will have on one another is considered.

Additionally, computation of the first moment of the KL divergence at each potential measurement time prevents even-keeled comparisons between different observation options. Consider the scenario in which two observers have the ability to take measurements of an object; the first observer provides higher quality data (i.e. has a lower measurement noise covariance), but the second observer is can afford an additional measurement and can take these observations later in the trajectory, as is illustrated in Figure 4.1. If it is desirable to keep one of the observers tasked with the acquisition of other target observations, which sensor should be tasked with observing the current target? The former generally provides a better state estimate due to the preferable measurement noise characteristics, while the latter allows the uncertainties to propagate longer, potentially resulting in the aforementioned “wait and see” behavior. In short, comparing the potential strength of the updates provided by each sensor results in the divergences being computed at different observation times, thus the comparison does not accurately portray the relative impact each update has. In scenarios like these, previous work has suggested that selecting a reference time at which to compare information gains can help mitigate this while more explicitly considering the dynamics at play in the tasking problem [35].

4.2.1.3. Objective function formulation. Equation (4.27) provides a relative measure of the expected uncertainty reduction achieved in ingesting a measurement without the need to task the sensor, enabling its use as a sensor tasking objective for a single measurement time. In order to leverage this measure for observation sets, rather than consider the individual contributions of each measurement sequentially, the first moment of the KL divergence must be accumulated over time. Consider again a single-target tracking scenario and let $\zeta = \{\delta_1, \dots, \delta_m\}$ be a set of m strictly increasing measurement time indices (i.e. z_{δ_k} is the measurement acquired of the target state at time t_{δ_k}). A sensor tasking optimization

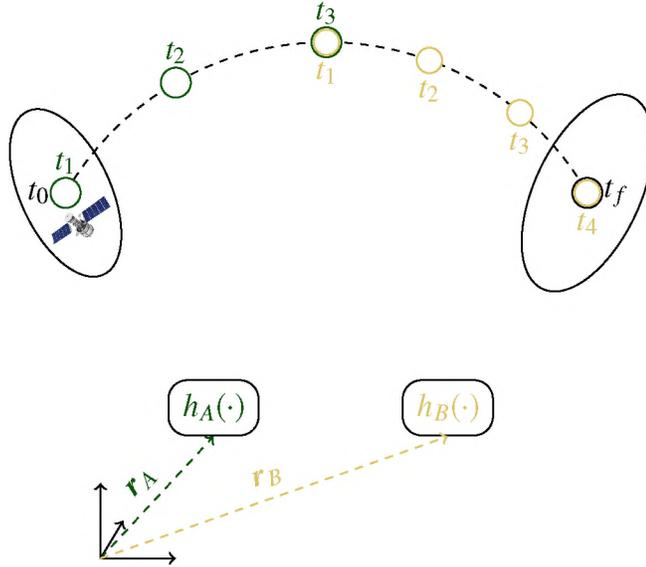


Figure 4.1. A target following along its nominal trajectory (dashed black line) with two observation sets originating from two (potentially) different sensor locations; one consisting of three measurements (green) generated according to measurement model $h_A(\mathbf{x})$, the other consisting of four measurements (gold) and generated according to measurement model $h_B(\mathbf{x})$.

leveraging the expected KL divergence given in Eq. (4.27) is defined as

$$\begin{aligned} \max_{\zeta} J(\zeta) &\triangleq \max_{\zeta} \sum_{\delta \in \zeta} M_{KL}(t_{\delta}) \\ &= \max_{\zeta} \frac{1}{2} \sum_{\delta \in \zeta} \log |\mathbf{P}_{\delta} \mathbf{\Pi}_{\delta}^{-1}|, \end{aligned} \quad (4.28)$$

where $J(\cdot)$ is the objective function to be maximized through the selection of the decision variables in ζ .¹

¹Note that in treating this as a conventional continuous optimization problem, this would be achieved through optimizing the individual measurement times t_{δ_i} . Representing these decision variables as measurement time indices aids in providing a consistent notation, as this formulation is leveraged in some mathematical developments in Section 5.

It is beneficial from both an analysis perspective as well as for developments presented in Section 5 to further analyze the myopic KL objective. From the properties of logarithms and determinants, the cost function appearing in Eq. (4.28) can be expressed as

$$\begin{aligned} J(\zeta) &= \frac{1}{2} \log \left(\prod_{\delta \in \zeta} |\mathbf{P}_\delta \mathbf{\Pi}_\delta^{-1}| \right) \\ &= \frac{1}{2} \log \left(|\mathbf{P}_{\delta_1}| \cdot |\mathbf{\Pi}_{\delta_1}|^{-1} \cdot |\mathbf{P}_{\delta_2}| \cdot |\mathbf{\Pi}_{\delta_2}|^{-1} \cdots |\mathbf{P}_{\delta_m}| \cdot |\mathbf{\Pi}_{\delta_m}|^{-1} \right). \end{aligned} \quad (4.29)$$

Expanding Eq. (4.28) as in Eq. (4.29) is useful in a couple of different ways. For instance, if state transition matrices are leveraged in the covariance propagation as in Eq. (4.10b) and assuming no process noise, each of the inner *a posteriori* and *a priori* covariance pairs in Eq. (4.29) can be represented as

$$\begin{aligned} |\mathbf{\Pi}_{\delta_k}|^{-1} \cdot |\mathbf{P}_{\delta_{k+1}}| &= |\mathbf{\Pi}_{\delta_k}|^{-1} \cdot |\mathbf{\Phi}(t_{\delta_{k+1}}, t_{\delta_k}) \mathbf{\Pi}_{\delta_k} \mathbf{\Phi}^T(t_{\delta_{k+1}}, t_{\delta_k})| \\ &= |\mathbf{\Phi}(t_{\delta_{k+1}}, t_{\delta_k})|^2. \end{aligned} \quad (4.30)$$

Equation (4.30) and the properties of state transition matrices can be used to simplify Eq. (4.29) as

$$J(\zeta) = \frac{1}{2} \log \left(|\mathbf{P}_{\delta_1}| \cdot |\mathbf{\Pi}_{\delta_m}|^{-1} \cdot |\mathbf{\Phi}(t_{\delta_m}, t_{\delta_1})|^2 \right). \quad (4.31)$$

Furthermore, if the dynamics governing the temporal evolution of the system are conservative, the volume of uncertainty is preserved as it propagates forward in time; this is due to the fact that state transition matrices for conservative dynamics are volume preserving or, more explicitly, exhibit a unity determinant as

$$|\mathbf{\Phi}(a, b)| = 1, \quad (4.32)$$

for all times a and b . Consequently, the determinant of two covariance matrices are equivalent so long as an update has not occurred. Since the set of observation time indices are monotonically increasing, it must be the case that $|\mathbf{\Pi}_k| = |\mathbf{P}_{k+1}|$ (and similarly $|\mathbf{P}_0| = |\mathbf{P}_{\delta_1}|$). This implies the objective given in Eq. (4.31) can be simplified even further to

$$J(\zeta) = \frac{1}{2} \log |\mathbf{P}_{\delta_1} \mathbf{\Pi}_{\delta_m}^{-1}|. \quad (4.33)$$

Under the assumptions of noiseless, conservative dynamics, and assuming the Kalman update in Eq. (4.12b) is employed, it can be shown that

$$|\mathbf{\Pi}_{\delta_m}| = |\mathbf{P}_0| \prod_{\delta \in \zeta} |\mathbf{I}_{n_x} - \mathbf{H}_\delta \mathbf{K}_\delta|. \quad (4.34)$$

Utilizing the constant volume of uncertainty in the propagation and Eq. (4.34), Eq. (4.33) can be expressed as

$$\begin{aligned} J(\zeta) &= -\frac{1}{2} \log \prod_{\delta \in \zeta} |\mathbf{I}_{n_x} - \mathbf{H}_\delta \mathbf{K}_\delta| \\ &= -\frac{1}{2} \sum_{\delta \in \zeta} \log |\mathbf{I}_{n_x} - \mathbf{H}_\delta \mathbf{K}_\delta|. \end{aligned} \quad (4.35)$$

Equation (4.35) provides some insight as to what the original objective function is quantifying; the original definition in Eq. (4.28) indicates that the measure is simply a sum of expected information divergences computed at different times (from potentially different observers), whereas Eq. (4.35) illustrates that it is in fact a sum of the log-determinants of the reduction matrices obtained from processing the observations taken at the times dictated by the indices in ζ . Unfortunately, Eq. (4.35) still depends on the intermediate *a priori* covariances \mathbf{P}_δ to compute the corresponding Kalman gains \mathbf{K}_δ . This can be alleviated and a more practical form is obtained if the information form of the covariance update in Eq. (4.14) is utilized.

Recursively substituting the information update of Eq. (4.14), the *a posteriori* covariance in Eq. (4.34) can be shown to be

$$\begin{aligned}\mathbf{\Pi}_{\delta_m}^{-1} &= \left(\mathbf{\Phi}(t_{\delta_m}, t_0) \mathbf{P}_0 \mathbf{\Phi}^T(t_{\delta_m}, t_0) \right)^{-1} + \sum_{\delta \in \zeta} \mathbf{\Phi}^T(t_\delta, t_{\delta_m}) \mathbf{H}_\delta^T \mathbf{R}_\delta^{-1} \mathbf{H}_\delta \mathbf{\Phi}(t_\delta, t_{\delta_m}) \\ &= \left(\mathbf{\Phi}(t_{\delta_m}, t_0) \mathbf{P}_0 \mathbf{\Phi}^T(t_{\delta_m}, t_0) \right)^{-1} + \sum_{\delta \in \zeta} \mathbf{\Xi}(t_\delta, t_{\delta_m}),\end{aligned}\quad (4.36)$$

where $\mathbf{\Xi}(t_\delta, t_{\delta_m})$ is a matrix mapping the information provided at time t_δ to the final measurement time t_{δ_m} . Equation (4.36) allows Eq. (4.33) to be represented as a function of the measurement noise covariance, measurement model Jacobian, state transition matrices evaluated along a nominal trajectory, and an initial covariance as

$$\begin{aligned}J(\zeta) &= \frac{1}{2} \left(\log |\mathbf{P}_0| \cdots \right. \\ &\quad \left. + \log \left| \left(\mathbf{\Phi}(t_{\delta_m}, t_0) \mathbf{P}_0 \mathbf{\Phi}^T(t_{\delta_m}, t_0) \right)^{-1} + \sum_{\delta \in \zeta} \mathbf{\Xi}(t_\delta, t_{\delta_m}) \right| \right).\end{aligned}\quad (4.37)$$

It is worth noting that the latter determinant in Eq. (4.37) is simply the inverse of the *a posteriori* covariance if a batch update is applied, processing the entire measurement set at the final measurement time. This shows the objective function in Eq. (4.37) is the difference in the volume of uncertainty at the initial time t_0 and the volume of uncertainty after the measurement set has been processed. As such, the accumulated myopic KL divergence approach provides an appropriate quantification of the information provided by the full observation set, since it is comprised of the overall reduction in uncertainty in the state density when processing all of the measurements as compared to just unobserved propagation. Lastly, leveraging the conservative dynamics enables the covariance in the first logarithmic term to be mapped to the final measurement time t_{δ_m} . Execution of this

mapping and through the properties of logarithms, Eq. (4.37) is simplified to

$$J(\zeta) = \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \tilde{\mathbf{P}}_{\delta_m} \sum_{\delta \in \zeta} \Xi(t_\delta, t_{\delta_m}) \right|, \quad (4.38)$$

where $\tilde{\mathbf{P}}_{\delta_m}$ is the initial covariance propagated to the final measurement time t_{δ_m} .

4.2.2. Forecasted Kullback-Leibler Divergence. As alluded to previously, one common issue that arises in the use of information divergences to measure the strength of measurements is that they tend to prioritize large prior uncertainties. Given observations of the same quality over time, the measurement with the largest strength tends to occur when the prior uncertainty is the largest. When the uncertainty grows over time, this leads to a “wait and see” structure that prioritizes measurements late in the trajectory that can possibly permit loss of track custody [33]. Additionally, this measure does not allow for a straightforward comparison when considering two potential observation sets. Consider again the simple schematic scenario in Figure 4.1, in which two sensor schedules are considered. One observer (Sensor A) generates measurements according to measurement model $h_A(\cdot)$ and is located at position \mathbf{r}_A , while the other (Sensor B) is located at \mathbf{r}_A and generates measurements according to a (potentially) different measurement model $h_B(\cdot)$. Sensor A can observe the target over the first half of the trajectory and can afford to take three observations (labeled in green) at the beginning, middle, and end of the first half of the plotted trajectory, while Sensor B can afford an additional measurement, but can only allocate them to the second half of the trajectory (labeled in gold). When employing the myopic KL divergence as an objective, how do these two sets compare? Is it fair to compare potential information gained at two different times? Is it worth considering that the sets consist of different numbers of measurements?

These questions are addressed and the “wait and see” structure is ameliorated by extending the information divergence approach to a set of multiple observations that can also be processed at a common, and more instructive, time [35]. In contrast to the

accumulated myopic measure in Eq. (4.28) in which a “reference time” is dictated by the final measurement time, an extension is carried out by selecting a time at which it is desirable to have a robust state estimate as a reference point at which to determine the impact a particular observation may have, e.g. before a target leaves the field of view or prior to the next maneuver in a trajectory.

4.2.2.1. Development. For the schematic toy problem illustrated in Figure 4.1, consider mapping the information contributions from Sensor A to a common reference time as in Figure 4.2. Computing the measure in this way provides a more direct comparison when considering measurement sets of differing qualities or different lengths. Recall the concatenated matrices constructed in Eq. (4.15) for use in the batch form of the covariance update in Eq. (4.16). By utilizing this form, it is possible to accumulate the potential information contribution from a set of measurements, which facilitates a means of “forecasting” the information to a time of interest. Leveraging this update, the expected KL divergence between the *a priori* and *a posteriori* densities at the reference time t_{ref} is found to be

$$\bar{M}_{KL}(\mathbf{t}) = \frac{1}{2} \log |\mathbf{P}_{\text{ref}} \mathbf{\Pi}_{\text{ref}}^{-1}|, \quad (4.39)$$

where the argument (\mathbf{t}) denotes the set of times in which the measurements that are processed in the update are acquired, \mathbf{P}_{ref} is the *a priori* state estimate covariance propagated to reference time t_{ref} in accordance with Eq. (4.10b), and $\mathbf{\Pi}_{\text{ref}}$ is the *a posteriori* state estimate covariance given the batch covariance update of Eq. (4.16). Note that the overbar notation in Eq. (4.39) denotes that the first moment of the KL divergence is mapped or “forecasted” to the reference time. An alternative for computing the *a posteriori* covariance is to use a form akin to the update in Eq. (4.8) as

$$\mathbf{\Pi}_{\text{ref}} = (\mathbf{I}_{n_x} - \bar{\mathbf{K}} \bar{\mathbf{H}}) \mathbf{P}_{\text{ref}}, \quad (4.40)$$

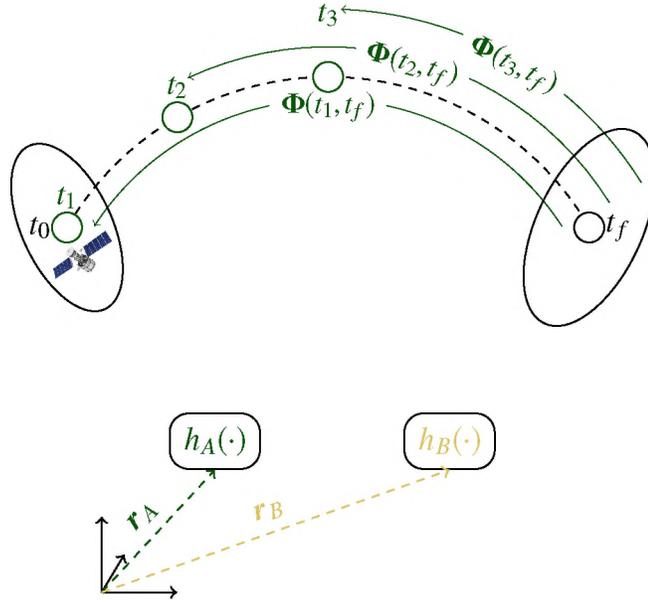


Figure 4.2. An illustration of mapping the measurement set information to a reference time t_f for the Sensor A measurement set; note the arcs accompanying the state transition matrices illustrate their mappings from time t_k to their respective measurement times.

where $\bar{\mathbf{K}}$ is the Kalman gain for the collected set of measurements, i.e.

$$\bar{\mathbf{K}} = \mathbf{P}_{\text{ref}} \bar{\mathbf{H}}^T (\bar{\mathbf{H}} \mathbf{P}_{\text{ref}} \bar{\mathbf{H}}^T + \bar{\mathbf{R}})^{-1}.$$

This proposed method is the exact formulation for batch processing measurements at a fixed time [81], but by taking this approach, the concatenated matrices in Eq. (4.15) can be used to evaluate the selected objective function for an arbitrary number of measurements, taken at arbitrary times, with any number of observation types exhibiting potentially different qualities or characteristics. Further, in the case of a single measurement, it is straightforward to show that in the absence of process noise the covariance in Eq. (4.40) can be expressed as a mapping of an *a posteriori* covariance computed at time t_k mapped to the reference time, i.e.

$$\mathbf{\Pi}_{\text{ref}} = \mathbf{\Phi}(t_k, t_{\text{ref}}) \mathbf{\Pi}_k \mathbf{\Phi}^T(t_k, t_{\text{ref}}). \quad (4.41)$$

Substitution of Eq. (4.41) into Eq. (4.39) yields

$$\begin{aligned}\bar{M}_{KL}(t_k) &= \frac{1}{2} \log |\mathbf{P}_{\text{ref}}(\mathbf{\Pi}_{\text{ref}})^{-1}|, \\ &= \frac{1}{2} \log \left| \mathbf{\Phi}(t_{\text{ref}}, t_k) \mathbf{P}_k \mathbf{\Phi}^T(t_{\text{ref}}, t_k) \left[\mathbf{\Phi}(t_{\text{ref}}, t_k) \mathbf{\Pi}_k \mathbf{\Phi}^T(t_{\text{ref}}, t_k) \right]^{-1} \right|,\end{aligned}\quad (4.42)$$

which collapses down to the myopic approach given in Eq. (4.27). This implies that given Eq. (4.40), the forecasted KL divergence for a single measurement in a linear (or linearized) system is essentially the mapping of the uncertainties at the measurement time to a reference time, an intuitive and comforting result, even though it implies the “wait and see” behavior persists in this “forecasted” approach in the case of single observation tasking. As a result, the forecasted objective in this scenario can be thought of as a generalization of the myopic approach and justifies the utilization of the first moment of the KL divergence when comparing single measurements, as it shows comparing two M_{KL} scores at different times is identical to comparing their corresponding \bar{M}_{KL} scores. However, the bulk of this formulation’s power lies in its ability to handle sets of observation.

The expression in Eq. (4.39) yields a tractable solution for the proposed sensor tasking objective; however, it requires the construction of the measurement model Jacobian, the reference *a priori* covariance, and the computation of the Kalman gain. Alternatively, a more implementation-friendly approach utilizes the information form of the update, as in Eq. (4.14), yielding a first moment of the KL divergence of

$$\bar{M}_{KL}(t) = \frac{1}{2} \log |\mathbf{I}_{n_x} + \mathbf{P}_{\text{ref}} \bar{\mathbf{H}}^T \bar{\mathbf{R}}^{-1} \bar{\mathbf{H}}|. \quad (4.43)$$

4.2.2.2. Objective function formulation. Similar to the myopic objective, the forecasted KL divergence objective function can be defined for the set of measurement time indices $\zeta = \{\delta_1, \dots, \delta_m\}$ to be [34, 35]

$$\max_{\zeta} \bar{J}(\zeta) \triangleq \max_{\zeta} \bar{M}_{KL}(\zeta) = \max_{\zeta} \frac{1}{2} \log |\mathbf{P}_{\text{ref}} \mathbf{\Pi}_{\text{ref}}^{-1}|, \quad (4.44)$$

where the overbar notation indicates the measure is forecasted to a reference time t_{ref} ; \mathbf{P}_{ref} is the *a priori* covariance propagated to the reference time, which is computed as

$$\mathbf{P}_{\text{ref}} = \mathbf{\Phi}(t_{\text{ref}}, t_0) \mathbf{P}_0 \mathbf{\Phi}^T(t_{\text{ref}}, t_0); \quad (4.45)$$

and $\mathbf{\Pi}_{\text{ref}}$ is the *a posteriori* state estimate covariance computed at time t_{ref} , with a batch information update as in Eq. (4.16).

The forecasted KL divergence objective function describes the expected KL divergence between the *a priori* and *a posteriori* state densities at an arbitrary reference time. The primary benefits of this formulation are natural support for multiple observations through simple concatenations of the matrices in Eq. (4.15) and the ability for a user to select the reference time of interest. Interestingly, by substituting Eqs. (4.16), (4.15), and (4.45) into Eq. (4.44), it can be shown that the forecasted KL measure becomes

$$\bar{J}(\zeta) = \frac{1}{2} \log |\mathbf{I}_{n_x} + \mathbf{P}_{\text{ref}} \bar{\mathbf{H}}^T \bar{\mathbf{R}}^{-1} \bar{\mathbf{H}}|, \quad (4.46)$$

which is equivalent to the myopic KL divergence objective when $t_{\text{ref}} = t_{\delta_m}$, connecting the forecasted KL to the conventional myopic KL objective and justifying the accumulation of the expected divergence measures over time.

4.2.2.3. Comparison to myopic. In order to observe the similarities and differences between the myopic and forecasted divergence objective functions, consider the scenario in which a ground-based observer is tasked with observing a single object in a circular

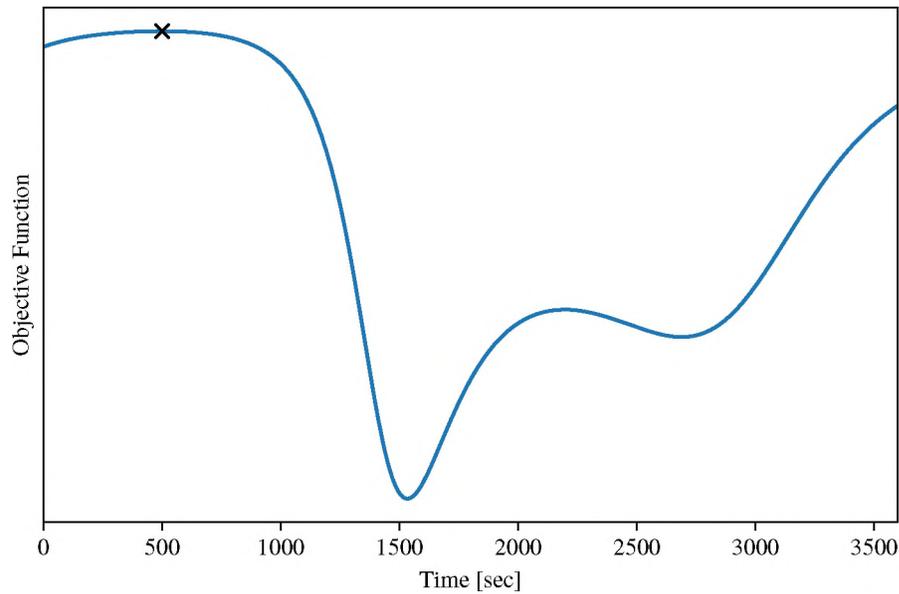


Figure 4.3. The myopic and forecasted KL divergence objective when scheduling a single measurement for a ground-based sensor, with the optimal measurement time denoted with an \times .

orbit. This sensor is capable of acquiring a range and range-rate measurement according to Eq. (3.6) at any point over an hour-long time interval. As shown previously, the myopic and forecasted measures are equivalent when considering a single measurement time, and consequently will agree on the optimal measurement time. This is observed in Figure 4.3, which depicts the objective functions (both the myopic and forecasted expected KL divergence in Eq. (4.38) and Eq. (4.46)) evaluated over the hour-long time interval (plotted in blue) with the optimal measurement time at approximately 500 seconds marked with an \times .

When a second measurement is considered, the purely myopic approach still highlights the same portion of the trajectory, scheduling the second observation immediately following the first. When defining the final observable time as the reference time, however, the forecasted measure determines that an observation at the beginning of the trajectory coupled with a measurement at roughly 1700 seconds yields a far more confident (i.e. less

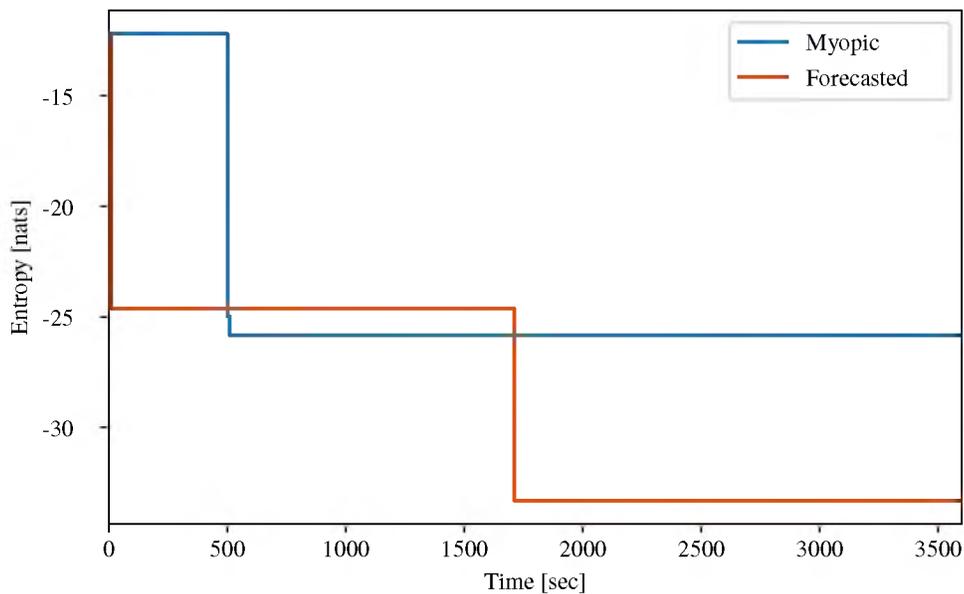


Figure 4.4. The entropy of the state estimate pdf when the myopic and forecasted schedules are processed by an EKF.

uncertain) state estimate at the end of the time interval. To evaluate the differences between the two schedules, an EKF is employed in tracking the target over the course of the hour; the filter is run once executing the myopic schedule, and a second time with the forecasted schedule. The resulting entropy in the state estimate density for the two schedules is shown in Figure 4.4. Note that the two successive measurements in the myopic schedule occur around 505 seconds and are successful in minimizing the uncertainty at that time step when compared to the forecasted schedule. However, the initial reduction in uncertainty brought on by the first measurement in the forecasted schedule enables a much more significant entropy reduction when the second measurement is processed as compared to the second myopic measurement, providing a much more confident (i.e. less uncertain) state estimate at the end of the time interval.

It may seem counterintuitive that mapping the information to the final reference time results in neglecting the measurement times that are suggested by the myopic measure. However this is the benefit of utilizing the forecasted measure; naturally accounting for the

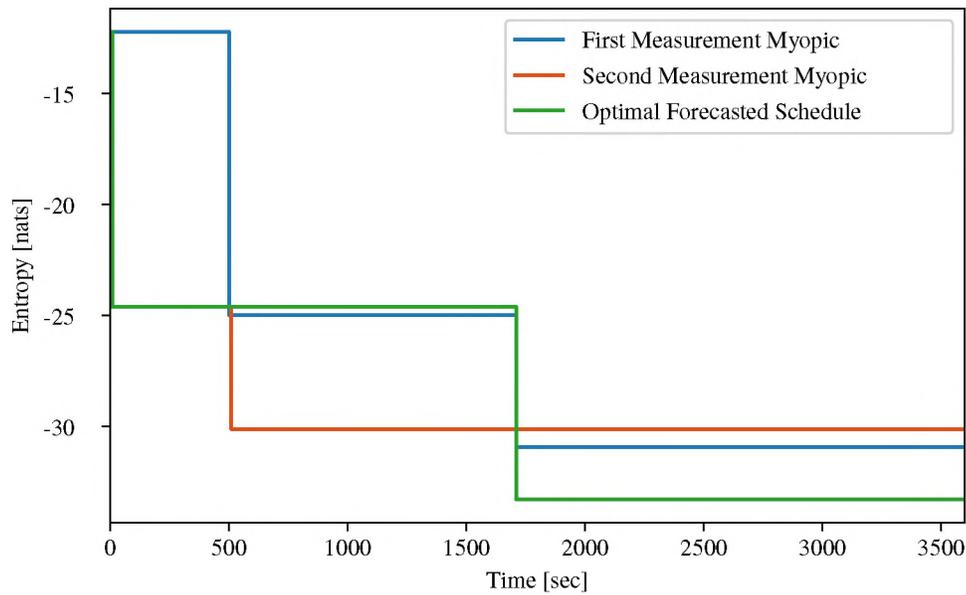


Figure 4.5. The entropy of the state estimate pdf with the optimal forecasted schedule and the two mixed schedules.

two measurement times simultaneously enables their respective information contributions to be coordinated, providing the state estimate with the smallest uncertainty at the final reference time. To better illustrate this, two new schedules are considered; one is composed of the first measurement time from the myopic schedule and the second measurement time of the forecasted, and the other composed of the first measurement time from the forecasted schedule and the second from the myopic schedule. These new schedules are constructed with the idea that if the trajectory is particularly information-rich around the 500 second mark, perhaps there are benefits to observing the target here along with a measurement time from the forecasted schedule. However, as illustrated by Figure 4.5, this is not the case; since the two non-optimal schedules do not benefit from coordinating the selection of the two measurement times simultaneously, the resulting entropy for the state estimate pdf falls short of matching the performance of the optimal forecasted schedule.

4.2.3. Optimization Considerations. The optimization of Eqs. (4.28) and (4.44) will be discussed in further detail in Section 5.6; however, it is useful at this point to address the obvious fact that the schedules generated from the optimization are heavily dependent upon the initialization of the optimization method, as several local maxima can exist on the surface (as can be seen, for example, in Figure 4.3). Different initialization approaches are available based on the desired performance, such as a multi-start approach performing m_{start} random initializations and further refining the best performing set, yielding a fast but not necessarily robust solution for small m_{start} or for a large number of measurements. A more robust (and more computationally taxing) approach can be achieved by computing a coarse approximation of the hypersurface generated by the objective function by constructing a collection of candidate observation sets that is representative of the different geometries an observation set can possess. Each of these candidate sets are then evaluated to determine which candidate provides the most information, and that set is then further refined through a more appropriate optimization technique.

This coarse approach better avoids lower quality solutions as opposed to generating random observation sets. Of course, this requires solving for the target state at a discrete set of candidate observation times (as well as the state transition matrix from these observation times to the reference times in the case of the forecasted measure) and evaluating the objective for each measurement combination, a task that can be very demanding from a computational standpoint. For instance, the optimization of a set of m_z measurements can be initialized using $m_{z,c}$ candidate observation times, but requires $m_{z,c}$ -choose- m_z function evaluations to determine the highest performing initialization. The number of candidate times can be increased in an effort to improve the initial candidate set; however, it is clear that the number of function evaluations increases combinatorially as more candidate observation times are considered. As such, it behooves a user to take steps to generate candidate observation sets in an appropriate manner such that they provide good coverage over the m_z -dimensional hypersurface. In many problems, this could be as simple as

generating $m_{z,c}$ candidate observation times that are evenly spaced over a time interval and then taking all possible combinations of the candidate times, subject to the constraints that $t_1 < t_2 < \dots < t_{m_z}$. The problem at hand, however, considers orbit tracking, resulting in dynamics that do not necessarily provide evenly-spaced target states when spacing them evenly in time and thus generating candidate times that will not yield a very representative set.

This issue of sampling points in time that are spaced by a constant time interval of Δt becomes apparent when the eccentricity of an orbit increases, as the object spends more time at apogee and generates more observation times, and conversely at perigee. To address this, it is useful to perform a Sundman transformation to enable integration with respect to an anomaly angle as opposed to time, given as [72]

$$\sqrt{\mu}dt = rd\theta,$$

where μ is the gravitational parameter, r is the distance of the object from the central body, and θ is referred to as a universal anomaly. Note that this is not the only form of the Sundman transformation, and in literature the gravitational parameter is often omitted to allow support for repulsive forces as well [73]. Since it is of interest here to have the universal anomaly be only a function of orbit geometry, this term is included here. This provides a simple way to generate a set of states separated by fixed angles on the ecliptic plane for an object undergoing Keplerian motion, resulting in a set much more representative of the geometry of the overall orbit. The Sundman based sampling is compared to the constant Δt sampling in Figure 4.6; note that the oversampling at apogee present with the constant Δt is reduced by applying the Sundman transformation and sampling with a constant $\Delta\theta$.

If the trajectory to be observed is sampled $m_{z,c}$ times (i.e. the target state and state transition matrix for $m_{z,c}$ times are computed) and m_z measurements are to be allocated over the observation window, there are $m_{z,c}$ -choose- m_z objective function evaluations performed

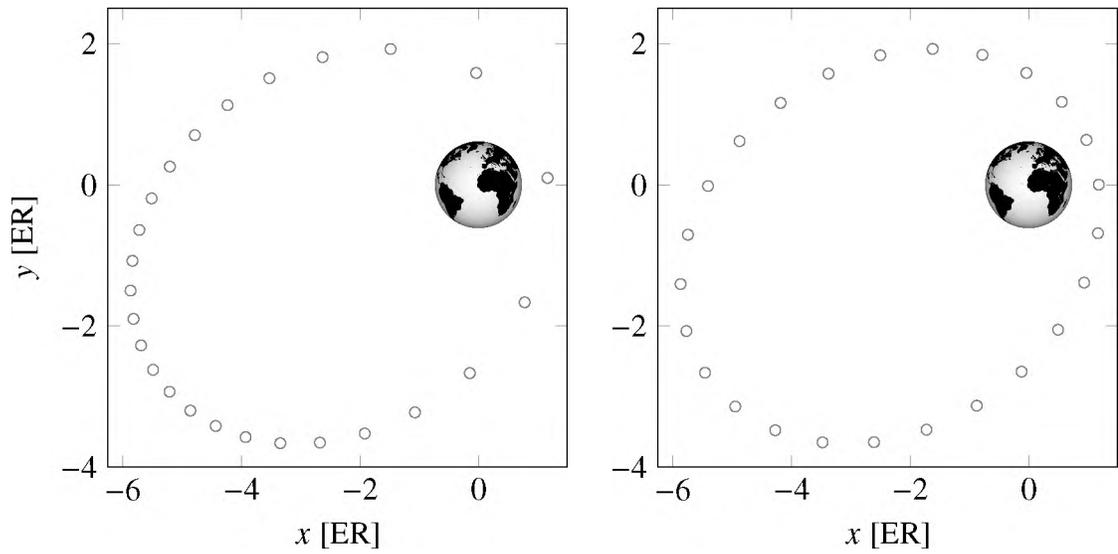


Figure 4.6. An eccentric orbit sampled evenly across time (left) or across the universal anomaly (right) with units of Earth radii (ER).

to obtain the coarse solution. From the resulting objective evaluations, the set that produces the maximal forecasted KL divergence measure is used to initialize a fast optimization routine in optimizing Eq. (4.44). Constraints on the individual measurements can also be specified (e.g. if the j^{th} element of the observation set \mathbf{t} is selected to be initialized as $t_j = \tau_k$, then the search region in that dimension can be limited to the interval between its neighboring candidate times (i.e. $\tau_{k-1} \leq t_j^* \leq \tau_{k+1}$) if the user chooses.

4.2.4. Computation in Other Spaces. A final note on the divergence measure in question concerns the vector space in which it is computed. All of the developments up to this point have defined the state space as simply position and velocity space, as the estimated state has thus been an estimated position and velocity. However, computation of the divergence-based cost function (both the myopic and forecasted implementations) is not limited to position and velocity estimates; applications in which updates provide information on other estimated parameters (e.g. sensor biases, misalignments, or other corruption parameters) can also benefit from these tasking policies. Further, although the form in Eq. (4.43) is defined on the *a priori* and *a posteriori* pdfs as they exist in the state

space, nothing is limiting its application to that space. It can easily be computed in any space deemed appropriate and of interest by the user. For instance, covariance scaling can be utilized to address condition number issues [12]; or the divergence can be computed in a different space such as the measurement domain, focusing more on the target's uncertainty in the measurement domain rather than the estimate as it exists in the state space.

This divergence can be computed in the measurement domain by revisiting Eq. (2.22) and utilizing the *a priori* and *a posteriori* measurement pdfs, as

$$p(\mathbf{z}_k) = p_g(\mathbf{z}_k; \mathbf{h}(\mathbf{m}_k), \mathbf{W}_k) \quad \text{and} \quad q(\mathbf{z}_k) = p_g(\mathbf{z}_k; \mathbf{h}(\boldsymbol{\mu}_k), \boldsymbol{\Omega}_k), \quad (4.47)$$

where \mathbf{W}_k and $\boldsymbol{\Omega}_k$ are the *a priori* innovation and *a posteriori* residual covariances. Computing the KL divergence between the pdfs in Eq. (4.47) yields

$$\begin{aligned} {}_z D_{KL}[q||p] = \frac{1}{2} & \left(\log |\mathbf{W}_k(\boldsymbol{\Omega}_k)^{-1}| + \text{tr} \{(\mathbf{W}_k)^{-1} \boldsymbol{\Omega}_k\} \cdots \right. \\ & \left. + (\mathbf{h}(\boldsymbol{\mu}_k) - \mathbf{h}(\mathbf{m}_k))^T (\mathbf{W}_k)^{-1} (\mathbf{h}(\boldsymbol{\mu}_k) - \mathbf{h}(\mathbf{m}_k)) - n_z \right), \end{aligned} \quad (4.48)$$

where the z pre-subscript on the left hand side of Eq. (4.48) indicates that the KL divergence is computed in the measurement space \mathcal{Z} . Recalling \mathbf{H}_k is the measurement model Jacobian evaluated at the mean \mathbf{m}_k allows $\mathbf{h}(\boldsymbol{\mu}_k)$ to be expanded in a first order Taylor series expansion as

$$\mathbf{h}(\boldsymbol{\mu}_k) \approx \mathbf{h}(\mathbf{m}_k) + \mathbf{H}_k(\boldsymbol{\mu}_k - \mathbf{m}_k). \quad (4.49)$$

Substitution of Eq. (4.49) into Eq. (4.48) yields

$${}_z D_{KL}[q||p] = \frac{1}{2} \left(\log |\mathbf{W}_k(\boldsymbol{\Omega}_k)^{-1}| + \text{tr} \{(\mathbf{W}_k)^{-1} \boldsymbol{\Omega}_k\} \cdots \right. \quad (4.50)$$

$$\left. + (\boldsymbol{\mu}_k - \mathbf{m}_k)^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k (\boldsymbol{\mu}_k - \mathbf{m}_k) - n_z \right), \quad (4.51)$$

The mean update of the EKF in Eq. (4.12a) can be used to substitute for the *a posteriori* mean $\boldsymbol{\mu}_k$ in Eq. (4.51), as

$$\begin{aligned} {}_z D_{KL}[q||p] = & \frac{1}{2} \left(\log |\mathbf{W}_k(\boldsymbol{\Omega}_k)^{-1}| + \text{tr} \{(\mathbf{W}_k)^{-1} \boldsymbol{\Omega}_k\} \dots \right. \\ & \left. + \text{tr} \left\{ (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k))^T \mathbf{K}_k^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{K}_k (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k)) \right\} - n_z \right), \end{aligned}$$

where the trace operator is introduced to leverage the fact that it remains constant under cyclic permutations; note that the term within the trace is a scalar, enabling the trace operator to be applied.

Just as before, taking the expected value with respect to the prior measurement likelihood $p(\mathbf{z}_k)$ will give the desired objective; assuming the prior covariance \mathbf{P}_k and the measurement model Jacobian \mathbf{H}_k are deterministic, this expectation can be expressed as

$$\begin{aligned} {}_z M_{KL} = & \frac{1}{2} \left(\log |\mathbf{W}_k(\boldsymbol{\Omega}_k)^{-1}| + \text{tr} \{(\mathbf{W}_k)^{-1} \boldsymbol{\Omega}_k\} + \text{tr} \left\{ \mathbf{K}_k^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{K}_k \dots \right. \right. \\ & \left. \left. \int_{\mathcal{Z}} (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k)) (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k))^T p(\mathbf{z}_k) d\mathbf{z}_k \right\} - n_z \right). \end{aligned} \quad (4.52)$$

Realizing the remaining integral is the *a priori* innovation covariance and leveraging the Kalman gain in Eq. (4.9), the third term in the expression can be manipulated to be

$$\begin{aligned} & \text{tr} \left\{ \mathbf{K}_k^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{K}_k \int_{\mathcal{Z}} (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k)) (\mathbf{z}_k - \mathbf{h}(\mathbf{m}_k))^T p(\mathbf{z}_k) d\mathbf{z}_k \right\} \\ & = \text{tr} \left\{ \mathbf{K}_k^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T \right\}. \end{aligned} \quad (4.53)$$

Similarly, the second trace term of Eq. (4.52) can be simplified through the definition of the *a posteriori* residual covariance $\boldsymbol{\Omega}_k$ and the covariance update in Eq. (4.12b), as

$$\begin{aligned} \text{tr} \{(\mathbf{W}_k)^{-1} \boldsymbol{\Omega}_k\} & = \text{tr} \{(\mathbf{W}_k)^{-1} (\mathbf{H}_k \boldsymbol{\Pi}_k \mathbf{H}_k^T + \mathbf{R}_k)\} \\ & = \text{tr} \{(\mathbf{W}_k)^{-1} (\mathbf{H}_k [\mathbf{I}_{n_x} - \mathbf{K}_k \mathbf{H}_k] \mathbf{P}_k \mathbf{H}_k^T + \mathbf{R}_k)\}. \end{aligned} \quad (4.54)$$

Substituting the definition of the Kalman gain from Eq. (4.9) into Eq. (4.54) yields

$$\begin{aligned} & \text{tr} \left\{ (\mathbf{W}_k)^{-1} (\mathbf{H}_k [\mathbf{I}_{n_x} - \mathbf{K}_k \mathbf{H}_k] \mathbf{P}_k \mathbf{H}_k^T + \mathbf{R}_k) \right\} \\ &= \text{tr} \left\{ (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T - (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T + (\mathbf{W}_k)^{-1} \mathbf{R}_k \right\}. \end{aligned} \quad (4.55)$$

Further, from the definition of the *a priori* residual covariance, $\mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T = \mathbf{W}_k - \mathbf{R}_k$.

Leveraging this fact and the Kalman gain in Eq. (4.9) once again in Eq. (4.55) gives

$$\begin{aligned} & \text{tr} \left\{ (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T - (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T + (\mathbf{W}_k)^{-1} \mathbf{R}_k \right\} \\ &= \text{tr} \left\{ \mathbf{I} - \mathbf{K}_k^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T \right\} \\ &= n_z - \text{tr} \left\{ \mathbf{K}_k^T \mathbf{H}_k^T (\mathbf{W}_k)^{-1} \mathbf{H}_k \mathbf{P}_k \mathbf{H}_k^T \right\}. \end{aligned} \quad (4.56)$$

Finally, substitution of Eqs. (4.53) and (4.56) into Eq. (4.52) results in

$${}_z M_{KL} = \frac{1}{2} \log |\mathbf{W}_k(\boldsymbol{\Omega}_k)^{-1}|, \quad (4.57)$$

a form akin to Eq. (4.27). This enables a new perspective on the utilization of information divergence by computing it in a different space and also permits the use of state transition matrices to generate an analogous forecasted measure ${}_z \bar{M}_{KL}$.

5. MULTITARGET SENSOR TASKING

In maximizing the expected KL divergence, either via the myopic or the forecasted approach, an information-rich measurement set is likely to be returned when optimizing over the objective function so long as appropriate constraints are defined and the initialization is handled in an appropriate manner (e.g. leveraging a multistart algorithm). The concepts from the previous section can be extended to the multitarget domain in a straightforward manner when the targets are independent and not coordinating, with most of the effort involved in handling the optimization (a process that is detailed further in Section 5.6). However, in approaching the problem of sensor tasking in the presence of multiple targets, other facets must be considered when arriving at a solution. One of the more glaring problems introduced in the multitarget tracking domain is that of data association, or the task of discerning the target from which a particular measurement originated. Failure to properly associate data prior to processing measurements can lead to a variety of behaviors that are detrimental to filter performance and can result in filter divergence, for example track coalescence or label swapping.

This section begins with a brief extension of the information divergence objectives given in Section 4 into the multitarget domain in Section 5.1 with a submodular analysis presented subsequently in Section 5.2. This is followed by a discussion of the problem of data association in Section 5.3 and a simple approach to performing data association when processing measurements. Section 5.4 concludes the section with a detailed employment of information entropies in characterizing potential collisions in the measurement space, illustrating their connection to difficult data association scenarios and enabling such situations to be avoided when the collision entropy is considered as a minimization objective.

5.1. MULTITARGET DIVERGENCE EXTENSION

Application of the divergence-based measures in Section 4 in a multitarget environment is straightforward when the observer(s) under consideration is (are) being tasked with maintaining i.i.d. target state estimates. Due to their independence, the observation of one target does not impact the estimate of another, and thus the divergences can be computed independently [36]. To be more explicit, the resulting multitarget extension of the myopic KL for L targets is defined as an accumulation of the single-target evaluations of Eq. (4.28) as

$$\mathcal{J}(\zeta) \cong \sum_{i=1}^L J(\zeta^{(i)}), \quad (5.1)$$

where the script notation $\mathcal{J}(\cdot)$ denotes the objective is computed for a multitarget state, $\zeta^{(i)}$ is the vector of measurement time indices for the i^{th} target, and the argument ζ is the collection of the L measurement time index vectors for each target.

A similar treatment can be given to the forecasted KL divergence. Leveraging the single-target objective in Eq. (4.44), the multitarget extension is defined as

$$\bar{\mathcal{J}}(\zeta) \cong \sum_{i=1}^L \bar{J}(\zeta^{(i)}). \quad (5.2)$$

Equations (5.1) and (5.2) provide myopic and forecasted approaches that can be used to maximizing the information provided by a set of observations over multiple targets. Maximization of these functions yields sensor schedules that exploit the advantages of the prescribed divergence measure and provides an appropriate objective when the targets are sufficiently separated. As the dimension of the input arguments increases (e.g. with the number of observers to consider, the number of targets to track, or the number of measurements to schedule), the optimization becomes cumbersome, warranting an exploration into how the optimization is approached.

5.2. SUBMODULAR FUNCTION ANALYSIS

The optimization of an arbitrary function is, in general, NP-hard [65], though certain classes of objective functions admit polynomial-time solutions, such as many classes of convex objective functions [68]. Unfortunately, objective measures based on the expectation-hedged KL divergence as continuous functions of time are not convex functions, in general. Alternatively, if the function is redefined as a set function that operates on a discrete set of observations, the measures exhibit diminishing returns, a property of *submodular set functions* (or simply submodular functions). A function $f(\zeta)$, $\zeta \subseteq \mathbb{Z}$, is said to be submodular if either of the following equivalent conditions holds true [67]:

$$\begin{aligned} f(\zeta_1) + f(\zeta_2) &\geq f(\zeta_1 \cup \zeta_2) + f(\zeta_1 \cap \zeta_2) && \forall \zeta_1, \zeta_2 \subseteq \mathbb{Z} \\ f(\zeta_1 \cup \{\tau\}) - f(\zeta_1) &\geq f(\zeta_2 \cup \{\tau\}) - f(\zeta_2) && \forall \zeta_1 \subseteq \zeta_2 \subseteq \mathbb{Z}, \forall \tau \in \mathbb{Z} \setminus \zeta_2. \end{aligned} \quad (5.3)$$

Similar to convex optimization, submodular optimization presents several useful properties that can be leveraged when selecting an optimizer [55].

5.2.1. Myopic Kullback-Leibler. In order to prove the single-target myopic objective in Eq. (4.28) is a submodular set function, it must be true that it exhibits the property of diminishing returns, as defined in Eq. (5.3). To illustrate this, define two measurement sets ζ_1 and ζ_2 consisting of monotonically increasing measurement time indices, such that $\zeta_1 \subseteq \zeta_2 \subseteq \mathbb{Z}$. Define τ as a measurement time index not contained in either measurement set, i.e. $\tau \in \mathbb{Z} \setminus \zeta_2$. The objective can be computed for the union of the measurement set ζ_1 and the additional measurement τ according to Eq. (4.37) as

$$J(\zeta_1 \cup \{\tau\}) = \frac{1}{2} \left(\log |\mathbf{P}_0| + \log \left| \tilde{\mathbf{P}}_{f_{\zeta_1 \cup \{\tau\}}}^{-1} + \Xi(t_\tau, t_{f_{\zeta_1 \cup \{\tau\}}}) + \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1 \cup \{\tau\}}}) \right| \right), \quad (5.4)$$

where $f_{\zeta_1 \cup \{\tau\}} = \max(\zeta_1 \cup \{\tau\})$ is the index corresponding to the last observation time. Solving the objective for solely the set ζ_1 as in Eq. (4.37) and subtracting from Eq. (5.4) yields the left hand side of Eq. (5.3) as

$$J(\zeta_1 \cup \{\tau\}) - J(\zeta_1) = \frac{1}{2} \left(\log \left| \tilde{\mathbf{P}}_{f_{\zeta_1 \cup \{\tau\}}}^{-1} + \Xi(t_\tau, t_{f_{\zeta_1 \cup \{\tau\}}}) + \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1 \cup \{\tau\}}}) \right| \cdots \right. \\ \left. - \log \left| \tilde{\mathbf{P}}_{f_{\zeta_1}}^{-1} + \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1}}) \right| \right), \quad (5.5)$$

where f_{ζ_1} is the greatest measurement time index in set ζ_1 . The right hand side of Eq. (5.3) can be computed following the same procedure for measurement set ζ_2 . It is immediately apparent that Eq. (5.5) is dependent upon when the additional measurement occurs relative to the end of the measurement set, as the propagation of the *a priori* covariance in the first determinant depends on this time. As a result, there are three cases to consider in order prove that the function is submodular: i) $t_\tau < t_{\delta_{f_{\zeta_1}}}$, ii) $t_{\delta_{f_{\zeta_2}}} > t_\tau > t_{\delta_{f_{\zeta_1}}}$, and iii) $t_\tau > t_{\delta_{f_{\zeta_2}}}$.

In Case i (i.e. the additional measurement t_τ occurs prior to the final measurement of the smaller observation set, $t_\tau < t_{\delta_{f_{\zeta_1}}}$), Eq. (5.5) becomes

$$J(\zeta_1 \cup \{\tau\}) - J(\zeta_1) = \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \Xi(t_\tau, t_{f_{\zeta_1}}) \left(\tilde{\mathbf{P}}_{f_{\zeta_1}}^{-1} + \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1}}) \right)^{-1} \right|, \quad (5.6)$$

and similarly for observation set ζ_2 . Assuming $\mathbf{R}_k > 0$, then it must be the case that $\Xi(t_\delta, t_{f_{\zeta_1}}), \Xi(t_\delta, t_{f_{\zeta_2}}) \geq 0$. It follows that

$$\left| \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1}}) \right| \leq \left| \sum_{\delta \in \zeta_2} \Xi(t_\delta, t_{f_{\zeta_2}}) \right|, \quad (5.7)$$

since the summation over ζ_2 will have at least as many elements as ζ_1 , and adding additional positive semi-definite matrices cannot lower the determinant of the resulting matrix. More rigorously, if matrices $\mathbf{A}, \mathbf{B} \geq 0$, then by Minkowski's inequality [59]

$$|\mathbf{A} + \mathbf{B}| \geq |\mathbf{A}| + |\mathbf{B}| \geq |\mathbf{A}|,$$

and since $\zeta_1 \subseteq \zeta_2$, Eq. (5.7) must be true. Following this same logic, it is easy to show that

$$\begin{aligned} & \left| \mathbf{P}_{f_{\zeta_1}} \Xi(t_\tau, t_{f_{\zeta_1}}) \left(\mathbf{I}_{n_x} + \mathbf{P}_{f_{\zeta_1}} \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1}}) \right)^{-1} \right| \geq \dots \\ & \left| \mathbf{P}_{f_{\zeta_2}} \Xi(t_\tau, t_{f_{\zeta_2}}) \left(\mathbf{I}_{n_x} + \mathbf{P}_{f_{\zeta_2}} \sum_{\delta \in \zeta_2} \Xi(t_\delta, t_{f_{\zeta_2}}) \right)^{-1} \right|. \end{aligned} \quad (5.8)$$

By Eq. (5.6), it is clear that to prove submodularity for Case i) it must be the case that

$$\begin{aligned} & \left| \mathbf{I}_{n_x} + \mathbf{P}_{f_{\zeta_1}} \Xi(t_\tau, t_{f_{\zeta_1}}) \left(\mathbf{I}_{n_x} + \mathbf{P}_{f_{\zeta_1}} \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1}}) \right)^{-1} \right| \geq \dots \\ & \left| \mathbf{I}_{n_x} + \mathbf{P}_{f_{\zeta_2}} \Xi(t_\tau, t_{f_{\zeta_2}}) \left(\mathbf{I}_{n_x} + \mathbf{P}_{f_{\zeta_2}} \sum_{\delta \in \zeta_2} \Xi(t_\delta, t_{f_{\zeta_2}}) \right)^{-1} \right|. \end{aligned} \quad (5.9)$$

Unfortunately, Minkowski's inequality is not applicable to Eq. (5.9). Instead, let $\lambda_i^{(1)}, i = 1, \dots, n_x$ be the n_x eigenvalues of the matrix within the determinant on the left hand side of Eq. (5.8), and $\lambda_i^{(2)}, i = 1, \dots, n_x$ be the n_x eigenvalues of matrix in the determinant on the right hand side of Eq. (5.9). For a matrix \mathbf{A} , if λ is an eigenvalue of \mathbf{A} and $f(x)$ is a polynomial, then $f(\lambda)$ is an eigenvalue of $f(\mathbf{A})$ [40]. As a result, it can be shown that $\lambda_i^{(1)} + 1, i = 1, \dots, n_x$ are the eigenvalues of the left hand side matrix of Eq. (5.9), and similarly $\lambda_i^{(2)} + 1, i = 1, \dots, n_x$ are the eigenvalues of the right hand side matrix. Consequently, Eq. (5.9) can be shown to be true by representing the determinants

in Eq. (5.8) as the product of their respective eigenvalues, or

$$\prod_{i=1}^{n_x} \lambda_i^{(1)} \geq \prod_{j=1}^{n_x} \lambda_j^{(2)} \quad (5.10)$$

where Eq. (5.10) is Eq. (5.9) represented as eigenvalue products. A similar representation of Eq. (5.10) is

$$\prod_{i=1}^{n_x} (\lambda_i^{(1)} + 1) \geq \prod_{j=1}^{n_x} (\lambda_j^{(2)} + 1), \quad (5.11)$$

Since Eq. (5.10) is true, Eq. (5.11) must also be true. Thus, the property of diminishing returns is true for the case in which the additional measurement occurs prior to the end of the observation set consisting of fewer measurements (i.e. before the end of set ζ_1).

In Case ii (i.e. the additional measurement t_τ occurs prior to the final measurement of the larger observation set but following the final measurement of the observation set consisting of fewer measurements ζ_1 such that $t_{\delta_{f_{\zeta_1}}} < t_\tau < t_{\delta_{f_{\zeta_2}}}$), it is the case that

$$f_{\zeta_1 \cup \{\tau\}} = \tau \quad \text{and} \quad f_{\zeta_2 \cup \{\tau\}} = f_{\zeta_2}. \quad (5.12)$$

Reflecting the relationships in Eq. (5.12), Eq. (5.5) can be expressed as

$$\begin{aligned} J(\zeta_1 \cup \{\tau\}) - J(\zeta_1) &= \frac{1}{2} \left(\log \left| \tilde{\mathbf{P}}_\tau^{-1} + \Xi(t_\tau, t_\tau) + \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_\tau) \right| \cdots \right. \\ &\quad \left. - \log \left| \tilde{\mathbf{P}}_{f_{\zeta_1}}^{-1} + \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{f_{\zeta_1}}) \right| \right) \end{aligned} \quad (5.13)$$

for the set consisting of fewer measurements ζ_1 and

$$J(\zeta_2 \cup \{\tau\}) - J(\zeta_2) = \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \Xi(t_\tau, t_{f_{\zeta_2}}) \left(\tilde{\mathbf{P}}_{f_{\zeta_2}}^{-1} + \sum_{\delta \in \zeta_2} \Xi(t_\delta, t_{f_{\zeta_2}}) \right)^{-1} \right|$$

for the set consisting of more measurements ζ_2 . The differences in Eq. (5.13) that prevent the expression from being simplified to Eq. (5.6) are subtle, but prevent the proof from being executed similarly to Case i. Due to the volume-preserving properties of conservative dynamics given in Eq. (4.32), time mappings can be employed to show that the property of diminishing returns holds for this case (as well as for Case iii), but the necessary application is not immediately apparent. To complete this proof, it is useful to consider the forecasted objective measure, as it reveals the mappings that must occur to prove the inequality.

5.2.2. Forecasted Kullback-Leibler. To prove the submodularity of the forecasted objective, it must be shown that the objective exhibits the property of diminishing returns, as defined in Eq. (5.3). As before, define two measurement sets ζ_1 and ζ_2 consisting of monotonically increasing measurement time indices, such that $\zeta_1 \subseteq \zeta_2 \subseteq \mathbb{Z}$, define τ as a measurement time index not contained in either measurement set, i.e. $\tau \in \mathbb{Z} \setminus \zeta_2$, and additionally select a reference time t_{ref} that occurs after all possible measurements, effectively dictating a finite time horizon. The objective can be computed for the union of the measurement set ζ_1 and the additional measurement τ according to Eq. (4.46) as

$$\begin{aligned} \bar{J}(\zeta_1 \cup \{\tau\}) &= \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \mathbf{P}_{\text{ref}} \sum_{\delta \in \zeta_1 \cup \{\tau\}} \Xi(t_\delta, t_{\text{ref}}) \right| \\ &= \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \mathbf{P}_{\text{ref}} \Xi(t_\tau, t_{\text{ref}}) + \mathbf{P}_{\text{ref}} \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{\text{ref}}) \right|. \end{aligned} \quad (5.14)$$

Just as in the myopic case, Eq. (4.46) can be solved for the measurement set ζ_1 and subtracted from Eq. (5.14) to yield the left hand side of Eq. (5.3) to be

$$\bar{J}(\zeta_1 \cup \{\tau\}) - \bar{J}(\zeta_1) = \frac{1}{2} \log \left| \mathbf{I}_{n_x} + \mathbf{P}_{\text{ref}} \Xi(t_\tau, t_{\text{ref}}) \left(\mathbf{I}_{n_x} + \mathbf{P}_{\text{ref}} \sum_{\delta \in \zeta_1} \Xi(t_\delta, t_{\text{ref}}) \right)^{-1} \right|, \quad (5.15)$$

and the right hand side of Eq. (5.3) can be solved similarly for ζ_2 . By selecting a reference time after all of the candidate observation times, the ambiguity present in the myopic objective function is lifted, and the only necessary condition to prove to ensure submod-

ularity is that Eq. (5.15) is greater than the same quantity evaluated for measurement set ζ_2 . Further, note that Eq. (5.15) is of the exact same form as Eq. (5.6) with a different time mapping. As a result, the same procedure as in Eqs. (5.7)–(5.11) can be followed to show that the forecasted objective and consequently its multitarget extension are submodular. Additionally, this same approach can be used to show that Cases ii and iii for the myopic objective function also hold true, and thus the myopic single- and multitarget formulations are submodular as well.

The fact that the objective functions in Eqs. (4.28), (4.44), (5.1), and (5.2) all belong to the class of submodular set functions makes several different submodular optimization schemes available when generating optimal schedules. These schemes afford approximate optimal results to be attained in feasible runtimes when considering problems of higher dimension, providing appropriate schedules in the event that the targets to be tracked are sufficiently separated. However, as the targets come into close proximity of one another, associating the acquired measurements becomes a more pressing issue. As such, the problem of data association warrants consideration in the sensor tasking policy.

5.3. DATA ASSOCIATION

The problem of data association in itself is a field of a substantial quantity of research, with methods ranging in complexity from solutions as simple as global nearest neighbors [8] to something as sophisticated as multi-hypothesis tracking [10]. In order to illustrate the issues that can arise when considering potentially imperfect data association, consider the simple scenario in which two targets are approaching each other in some space. It is of interest to be able to mathematically approximate when a conjunction may occur, as the sensor schedule should depend on it; for instance, in some applications it may be desirable to take observations of conjunction events if they occur in the state space, or it may be best to avoid the conjunction time if it is occurring in the measurement space to avoid misassociations. The $1\text{-}\sigma$ uncertainty intervals of two targets are plotted at each time

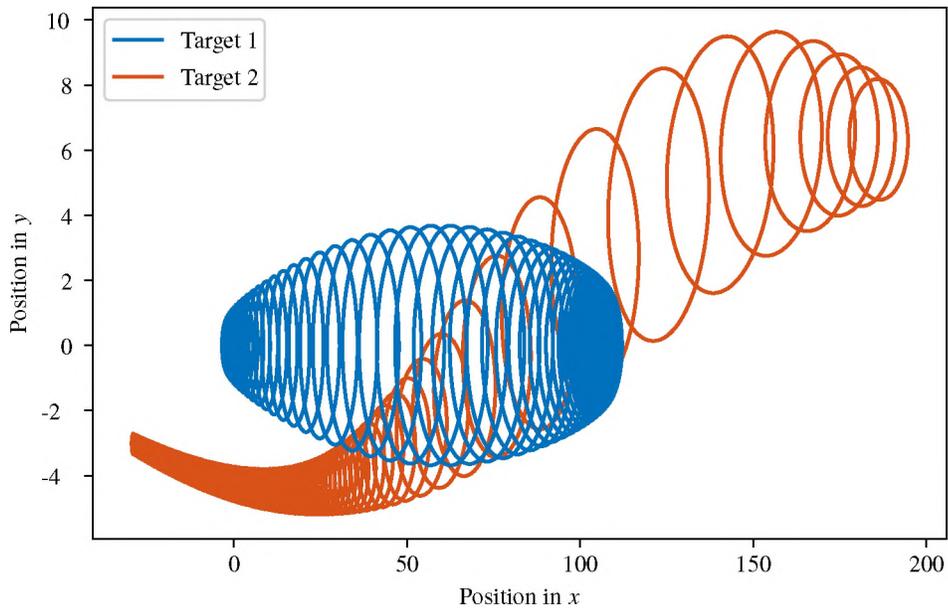


Figure 5.1. The $1\text{-}\sigma$ intervals for the two targets centered on their respective means and plotted at each time step in the state space.

step in the space of interest in Figure 5.1. Clearly, the two targets come into close proximity of one another, but when does the conjunction occur exactly? At what point do the two state estimate densities “collide,” and when does this “collision” cease?

A multitude of solutions is available in the field of conjunction analysis, and again these solutions exhibit an array of differing complexities. One approach in particular that has received some attention recently in conjunction analysis and has been frequently used in data association solutions is via gating the squared Mahalanobis distance [19]. The squared Mahalanobis distance of a target $\mathbf{x}_k^{(i)}$ with respect to its state estimate density $p(\mathbf{x}_k^{(i)})$ is then defined to be [56]

$$d^2 = \left(\mathbf{x}_k^{(i)} - \mathbf{m}_k^{(i)} \right)^T \left(\mathbf{P}_k^{(i)} \right)^{-1} \left(\mathbf{x}_k^{(i)} - \mathbf{m}_k^{(i)} \right), \quad (5.16)$$

where $\mathbf{m}_k^{(i)}$ and $\mathbf{P}_k^{(i)}$ are the mean and covariance of the i^{th} target state estimate density.

It is well known that when the distribution at hand is Gaussian, the squared Mahalanobis distance follows a chi-square distribution, characterized by a single parameter known as the degree-of-freedom. The degree-of-freedom in the squared Mahalanobis distance distribution is simply the dimension of the random vector. As such, it follows that the distribution of Eq. (5.16) is

$$p(d^2) = p_{\chi^2}(d^2; n_x).$$

Possession of the true density for this distance is particularly useful in data association, as it enables a probability gate to be defined, or a threshold with some statistical meaning. If a filter should accept P percent of measurements it receives on a particular target, an associated threshold γ can be determined and employed to either accept or reject processing a measurement.

For example, for the case of 2 degrees-of-freedom, a $P = 99.99\%$ acceptance probability would result in a squared Mahalanobis distance threshold of $\gamma \approx 18.421$. This threshold can be applied with a few different objectives in mind. If evaluating potential state space collisions is of interest, a new random variable can be defined to be the difference in the Cartesian position between the two states, enabling the squared Mahalanobis distance threshold to be applied to that random variable [19]. Employing this thresholding in data association applications requires the computation of the statistical distance between a received measurement and the distribution of a target projected into the measurement space; this statistical distance is then used to determine if the measurement should be associated to the target. In assessing scenarios in which data association might be difficult, it is useful to examine the statistical distance between the expected measurement of one target and the distribution of another target, or for the example given in Figure 5.1, the squared Mahalanobis distance between the expected measurement of Target 2 with respect to the distribution of Target 1.

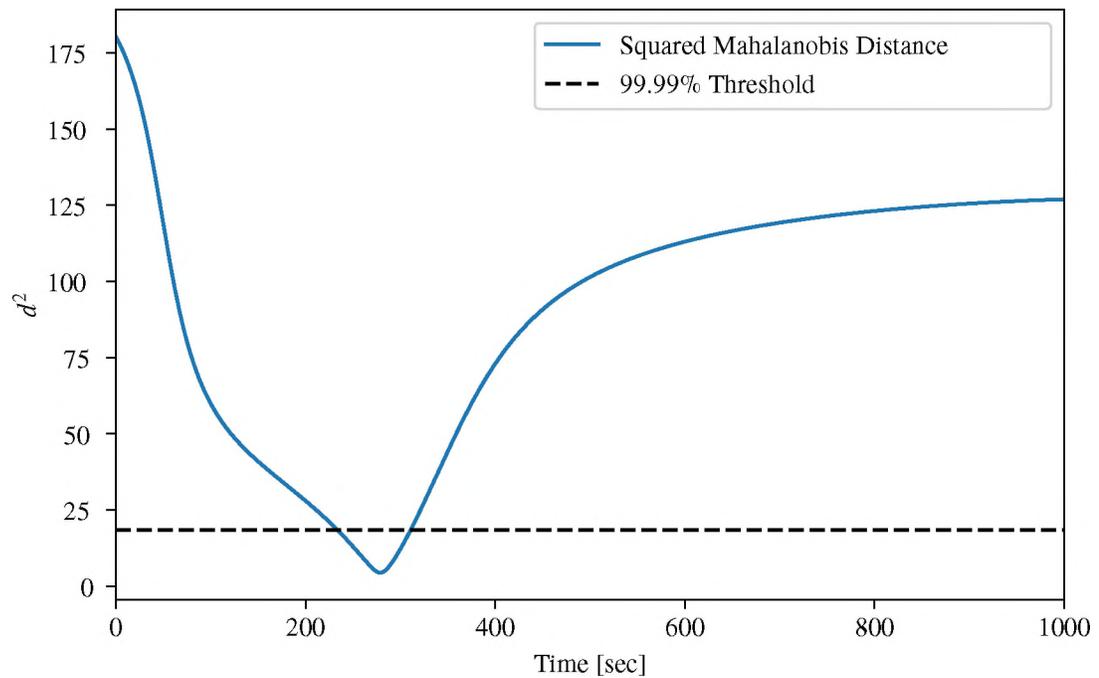


Figure 5.2. The squared Mahalanobis distance of the expected measurement generated by Target 2 with respect to the measurement density for Target 1.

This can be seen for the two targets depicted in Figure 5.2, with the squared Mahalanobis distance of the expected measurement generated by Target 2 with respect to the measurement density of Target 1 computed and shown in blue and the 99.99% threshold denoted by the black dashed line. The result illustrates that if an observation of Target 2 is taken around roughly 300 seconds, the data association is at risk of misassociating the observation, potentially processing the generated measurement as an observation of Target 1. While this visually illustrates when a measurement should be avoided, the measure becomes more difficult to employ in the presence of more targets without the necessity of considering each pairwise Mahalanobis distance. Fortunately, information entropies can provide an alternative approach to evaluating conjunctions.

5.4. ENTROPY CONSIDERATION

As illustrated by Figure 5.2, the use of a statistical distance in addressing the problem of data association becomes difficult in the presence of multiple targets, and issues can arise when the targets come in close proximity to one another. This section details the employment of information entropies in considering these issues when generating a sensor schedule.

5.4.1. Shannon Entropy. Recall the notion of entropy in the sense of Shannon presented in Section 2. Equation (2.13) provides a means of mapping the full distribution of a random vector to a single scalar quantity that is representative of the spread of the density, a value that can be considered the amount of information there is to gain on the random quantity or the volume of uncertainty in its outcome. In the context of communication and signal processing, entropy is interpreted as a quantification of the amount of information communicated or a measure of certainty in a transmitted or received message.

Another mentality for entropy measures is to think of it as a relative concentration of uncertainty for different target arrangements. More specifically, smaller entropies indicate a more highly concentrated or more highly localized uncertainty. To illustrate this, consider the schematic representation given in Figure 5.3, where the time history for one-, two- and three-sigma intervals for three objects are depicted as ellipses. Initially at time t_1 , the three objects are distinctly separated. As time progresses, differing levels of interactions occur between different combinations of the three objects. During these interactions, the concentration of uncertainty is increasing at these points of interaction, illustrating that the points of high concentration indicate pdf coalescence.

Revisiting the example given in Figure 5.1, the entropy of the entire multitarget system is at a minimum when the two object densities are at a maximal coalescence, providing a relative measure of a “level” of conjunction in their vector space. Unfortunately, Shannon entropy, as defined in Eq. (2.13), cannot be used to describe the level of interaction between multiple objects, as it is the entropy obtained from a pdf representing the uncertainty of a single object. Augmenting the state as a concatenation of the individual target states

also does not work; when the objects are independent, it is straightforward to show that the entropy of the distribution for the augmented state is equal to the sum of the entropies of the distribution of the individual objects. In essence, this extension to the multitarget domain omits the interactions between individual objects, entirely neglecting collisions.

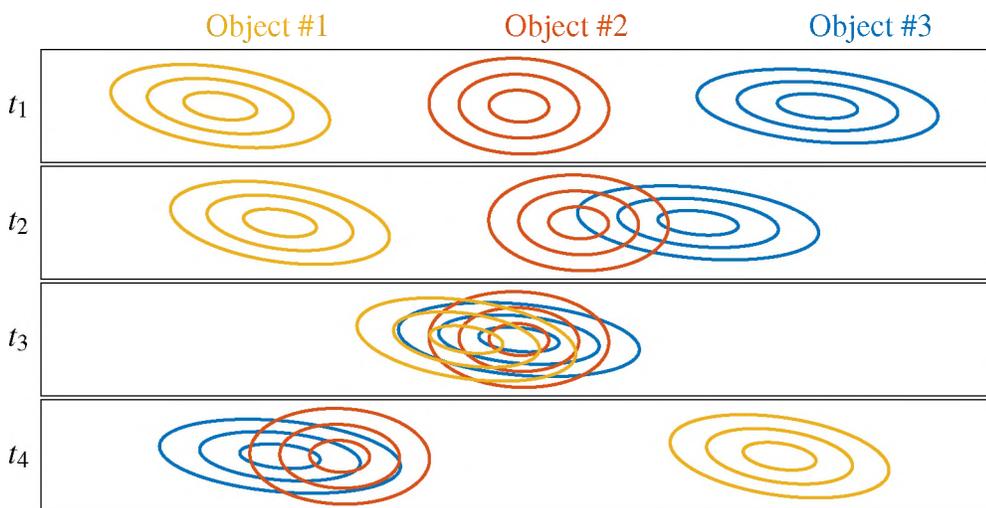


Figure 5.3. A simple example of the interaction of three targets.

When the multitarget state is taken to be an i.i.d. Poisson RFS X , with the intensity function modeled as a Gaussian mixture, Eq. (2.15) provides the multitarget differential entropy, realized as the difference between the expected number of targets and a spatial entropy term that is a function of its intensity function. As mentioned previously, this spatial entropy is identical to that of the single-target Shannon entropy but with the multitarget intensity in place of the single-target pdf, lending a sense of intuition for its behavior, whereas the cardinality term provides an offset for the entropy based on the expected number of targets in the system.

5.4.2. Rényi Entropy. The measure given in Eq. (2.15) cannot be computed in closed form when the system is modeled as a Gaussian mixture, but it can be approximated by omitting multitarget interactions and simply considering the entropy introduced by the cardinality of the system and the individual spatial entropies. This is useful for analyzing

performance in tracking multiple targets, but cannot be directly leveraged when analyzing collisions. Fortunately, this is not the case for the Shannon entropy generalization, namely the Rényi entropy.

Again, taking on the assumption that the multitarget system is an i.i.d. Poisson cluster process, the Rényi entropy can be expressed as a function of the rate parameter, λ , and the intensity function, $\nu(\mathbf{x})$, as in Eq. (2.17). Just as with the Shannon entropy computed under identical assumptions, the Rényi entropy in Eq. (2.17) exhibits a cardinality term and a spatial entropy term that is a function of the RFS intensity. Unlike the Shannon entropy, however, the spatial element of the Rényi entropy does not take on the same form as the single-target Rényi entropy, which can be seen by comparing Eq. (2.16) and Eq. (2.17).

Recall that, for an i.i.d. Poisson RFS with its intensity modeled as a Gaussian mixture, the collision entropy (i.e. $\alpha = 2$) affords the closed-form solution in Eq. (2.19). This expression illustrates that the collision entropy of the Gaussian mixture model is comprised of a cardinality entropy and a spatial component, much like the multitarget Shannon entropy. The cardinality term is given simply as the sum of the weights of the Gaussian components within the mixture; thus, for a known and constant number of targets, this portion of the entropy is constant. The spatial component, or the double summation in Eq. (2.19), contains information that describes the level of interaction between multiple targets, a facet of interest when attempting to avoid observing collisions in the measurement space. From Eq. (2.19), the spatial element of the Rényi entropy computed in the measurement space is given as [19]

$$H_{S,k}^{(2)}(\Theta) = - \sum_{i=1}^L \sum_{j=1}^L \left[w_k^{(i)} w_k^{(j)} \Gamma(\hat{\mathbf{z}}_k^{(i)} - \hat{\mathbf{m}}_k^{(j)}, \mathbf{W}_k^{(i)} + \mathbf{W}_k^{(j)}) \right], \quad (5.17)$$

where Θ is the collection of parameters that defines the GM representation of the intensity for which the spatial entropy is computed (i.e. $\Theta = \{w^{(i)}, \hat{\mathbf{z}}^{(i)}, \mathbf{W}^{(i)}\}_{i=1}^L$) and recall

$$\Gamma(\mathbf{a}, \mathbf{A}) = |2\pi\mathbf{A}|^{-1/2} \exp\left\{-\frac{1}{2}\mathbf{a}^T \mathbf{A}^{-1} \mathbf{a}\right\}.$$

Equation (5.17) provides a relative measure of pdf coalescence between all targets within the multitarget state. This quantity facilitates a comparison of different observation times from a data association standpoint, mitigating risky association scenarios by avoiding observations at times in which it is difficult to distinguish between targets. Equation (5.17) has been employed previously in analyzing potential conjunction events in the state space in an attempt to detect time intervals of physical collisions [19] and is used in a similar fashion here to avoid observation times that exhibit relatively difficult data association scenarios.

A further specialization to Eq. (5.17) can be made at a computational expense in the form of an increase in the number of optimizations required to obtain an optimal schedule. The coalescence between each individual target and the multitarget state can be represented through L function evaluations and omitting the outer sum in Eq. (5.17). Specifically, the interaction of target i with the rest of the multitarget state at time t_k can be computed as

$$H_{S,k,i}^{(2)}(\Theta) = -2 \sum_{j=1}^L \left[w_k^{(i)} w_k^{(j)} \Gamma(\hat{\mathbf{z}}_k^{(i)} - \hat{\mathbf{m}}_k^{(j)}, \mathbf{W}_k^{(i)} + \mathbf{W}_k^{(j)}) \right]. \quad (5.18)$$

The single-target spatial entropy given in Eq. (5.18) is attractive as it avoids the scenario in which a measurement time for target i is rescheduled (or not considered) due to collisions happening elsewhere in the multitarget state, though the appropriateness of its use is dependent upon the architecture of the sensor tasking framework. If one intends to simultaneously maximize the expected information gain and minimize collisions in the measurement domain, Eq. (5.18) is not applicable, as it would result in multiple outputs for the objective

function. However, if other solutions are considered that enable each of the L single-target spatial entropies to be considered, the measurement schedules for each individual target can be optimized against its own collision measure.

5.5. OBJECTIVE FUNCTION FORMULATION

The preceding sections have outlined two different measures that 1) quantify the expected amount of information provided by a measurement set and 2) provide a relative measure indicating measurement times that could (potentially) generate risks in data association. When used in tandem, special considerations must be taken since the two are measuring different facets of the tasking problem in different units. In order to simultaneously consider them in an optimizer, some attention must be given to formulating a formal objective function in a way such that the two measures are relatively evenly weighted.

Define an objective function as

$$O(\mathbf{t}) = \iota(\mathbf{t}) + \eta(\mathbf{t}), \quad (5.19)$$

where $\iota(\cdot)$ is a term related to the expected information gain and $\eta(\cdot)$ is related to the spatial entropy. An intuitive means of making sure neither the expected information gain nor the spatial entropy consistently dominates the other is to bound each of their contributions by 0 and 1. If upper and lower bounds on the terms can be acquired, this sort of normalization can be attained. With regard to the expected KL divergence, it is clear from Eq. (4.18) that the lower bound for the myopic KL divergence for a single target is

$$D_k^- = \frac{1}{2} \left(\log |\mathbf{P}_k \mathbf{\Pi}_k^{-1}| + \text{tr} \{ \mathbf{P}_k^{-1} \mathbf{\Pi}_k \} - n_x \right), \quad (5.20)$$

with the lower bound for the forecasted KL divergence being a straightforward extension. Additionally, as discussed in Section 4.2.1.1, a statistical distribution for the KL divergence can be acquired if the individual channels of the measurement are assumed to be uncorrelated, or alternatively if they are made to be independent [12] the measure can be expressed as the sum of a deterministic bias and a chi-square distributed random variable as in Eq. (4.25).

Equation (4.25) demonstrates that the KL divergence does not have a finite upper bound. Since the stochastic contributor to the quantity is chi-square distributed, it is possible to account for a P -quantile event to be the “upper bound” on the measure. For instance, similar to gating the squared Mahalanobis distance, if a measurement event resulting in 99.99% of the information that can be gained (i.e. continuously sampling the chi-square distribution would result in a smaller update 99.99% of the time), then the “upper bound” on the expected information gain on a two-dimensional measurement is dictated by a threshold of $\gamma \approx 18.421$, and the resulting threshold for the KL divergence is computed as

$$D_k^+ = \frac{1}{2} \left(\log |\mathbf{P}_k \mathbf{\Pi}_k^{-1}| + \text{tr} \{ \mathbf{P}_k^{-1} \mathbf{\Pi}_k \} + \gamma - n_x \right) . \quad (5.21)$$

The “bounds” in Eqs. (5.20) and (5.21) can then be used to “normalize” the expected information gain terms and accumulate them over the measurement set to yield

$$\iota(\zeta) = \sum_{\delta \in \zeta} \frac{M_{KL}(t_\delta) - D_\delta^-}{D_\delta^+ - D_\delta^-} . \quad (5.22)$$

Unfortunately, for the forecasted divergence measure the stochastic contributor to the measure is not, in general, exactly chi-square distributed; correlations are accumulated within this quantity as entire measurement sets are examined simultaneously. These correlations

tend to have small effect, and thus assuming independence between measurements and following a similar approach yields a fair approximation as

$$\bar{\iota}(\zeta) = \frac{\bar{M}_{KL}(\zeta) - D_{\text{ref}}^-}{D_{\text{ref}}^+ - D_{\text{ref}}^-},$$

where D_{ref}^- and D_{ref}^+ are computed according to Eqs. (5.20) and (5.21) but with the prior covariance propagated to the reference time according to Eq. (4.6) and posterior covariance updated at the reference time updated according to Eq. (4.16).

Fortunately, the coalescence, or collision entropy term, can be similarly bounded, except with definitive bounds. In the event that zero interaction is occurring between targets (i.e. the targets are separated by an infinite distance in the measurement domain), the interactions between different targets in Eq. (5.17) result in no contribution to the spatial entropy. Recall that larger collision entropies indicate a lower concentration of localized uncertainty. As a result, the upper bound on the collision entropy is

$$\eta_k^+ = 2 \sum_{\ell=1}^L w_k^{(\ell)} - \sum_{i=1}^L (w_k^{(i)})^2 |4\pi \mathbf{W}_k^{(i)}|^{-1/2}.$$

Conversely, if the targets are all collocated, the minimum collision entropy is computed as

$$\eta_k^- = 2 \sum_{\ell=1}^L w_k^{(\ell)} - \sum_{i=1}^L \sum_{j=1}^L \left[w_k^{(i)} w_k^{(j)} |4\pi (\mathbf{W}_k^{(i)} + \mathbf{W}_k^{(j)})|^{-1/2} \right].$$

These bounds enable a coalescence term bounded by zero and one to be defined as

$$\begin{aligned} \eta(\zeta) &= \sum_{k \in \zeta} \eta_k(\mathbf{z}_k) \\ &= \sum_{k \in \zeta} \frac{H_k^{(2)}[f] - \eta_k^-}{\eta_k^+ - \eta_k^-}. \end{aligned} \tag{5.23}$$

Substitution of Eqs. (5.22) and (5.23) into Eq. (5.19) yields an objective that attempts to evenly weight potential risk of misassociating data against the potential information gained for a particular measurement set. Simultaneously considering both of these terms theoretically enables the optimal sensor schedule with regard to this objective to be generated. In selecting an optimization scheme, a couple of factors should be taken into consideration. As the collection of targets increases, the number of available sensor resources increases, and/or the number of candidate measurement times considered grows, this objective function becomes increasingly burdensome from a computational standpoint. Additionally, the weighting of the two terms becomes more influential, as local extrema exist at points in which one of the terms is essentially “neglected” if the other is dominant in the resulting objective evaluation. If, however, the two terms are considered separately and a two-phase optimization solution is applied, the computational effort required is substantially reduced.

This computational requirement is reduced further when considering some of the underlying properties of the objective functions at hand. The expected information gain term exhibits diminishing returns and thus falls into the classification of a submodular set function, as was illustrated in Section 5.2. Leveraging this fact enables submodular optimization techniques to be applied to the term in Eq. (5.22) to generate solutions in a timely manner with some level of freedom with regard to emphasizing optimization run time or the quality of the resulting schedule. Due to the relatively inexpensive optimization of Eq. (5.23), performing two sequential optimizations results in a drastically lower runtime when compared to the simultaneous optimization of the two measures. Both the simultaneous optimization as well as the submodular optimization schemes are detailed further in the following section.

5.6. OPTIMIZATION

The simultaneous optimization of both the expected information gain and the collision entropy given in Eq. (5.19) can be constrained in an attempt to make the optimization more feasible; however, its formulation warrants discussion to illustrate its implementation. Consider the scenario in which a schedule is to be generated for m_o observers, with each schedule consisting of $m_{z,s}$ observations ($s = 1, 2, \dots, m_o$). In this scenario, L targets are considered during the sensor tasking generation. In order to simultaneously consider the expected information gain acquired by the resulting schedules as well as the measurement spatial coalescence “accumulated” by executing these schedules, a single time vector consisting of the schedules for the m_o observers be considered as the optimization parameters; to be explicit, the time vector can be defined as $\mathbf{t} = [\mathbf{t}_1^T, \mathbf{t}_2^T, \dots, \mathbf{t}_{m_o}^T]^T$, where each \mathbf{t}_s is the schedule of $m_{z,s}$ measurement times for sensor s . For simplicity, each sensor is assumed to take a fixed number of measurements of each target, $m_{z,t}$, such that in total it is taking $L \times m_{z,t}$ measurements. The complete schedule is then defined to be $\mathbf{t} = [\mathbf{t}_{1,1}^T, \mathbf{t}_{1,2}^T, \dots, \mathbf{t}_{s,L}^T]^T$, where each $\mathbf{t}_{s,i}$ is a schedule of $m_{z,t}$ measurements of target i for sensor s . The result is a $d = m_o \times L \times m_{z,t}$ -dimensional optimization problem, defined as

$$\mathbf{t}^* = \arg \max_{\mathbf{t}} \{ \iota(\mathbf{t}) + \eta(\mathbf{t}) \} \quad (5.24a)$$

subject to

$$\mathbf{C}_+ \mathbf{t} \geq \mathbf{0}_{d \times 1} \quad (5.24b)$$

$$\mathbf{t}_0 \leq \mathbf{t} \leq \mathbf{t}_f, \quad (5.24c)$$

where \mathbf{C}_+ is a matrix enforcing monotonically increasing time steps comprised of $L \times m_o$ forward difference matrices of dimension $m_{z,t}$ on the diagonal, and \mathbf{t}_0 and \mathbf{t}_f set lower and upper bounds on the search region, respectively.

Equations (5.24a)-(5.24c) provide the framework for an optimization problem that can be executed for a small number of observers/targets/measurements-per-target, but which quickly becomes intractable as the dimension of the problem grows. For many optimization problems, the desire for a near-globally optimal solution does not outweigh the computational burden required to achieve such a solution. This is no different in the problem of sensor tasking; due to the limited number of sensor resources, it is of interest to get the most out of a sensor schedule; unfortunately, considering the vast number of objects to track and the indefinite time interval that must be scheduled, finding the globally optimal schedule is not always feasible. Recall that, as per Section 5.2, both the myopic and forecasted formulation of the expected KL divergence fall into the class of submodular functions. If the optimization is broken into two phases, the KL divergence-based portion of the objective can be optimized utilizing submodular optimization strategies. Furthermore, if this phase of the optimization is executed first and the expected KL divergence maximized, the second stage can be cast as a series of d independent scalar optimizations in which the individual measurement times are altered slightly in a way that minimizes the coalescence term and avoids risky association events. As a result, a two-phase optimization scheme exhibits a second phase that is significantly less complex than the first, generating an optimization problem that requires far fewer computational resources and is less sensitive to the employed optimizer. That is, selection of any appropriate optimizer (i.e. one for continuous, constrained, and non-convex problems) will be sufficient in handling the second phase. The following subsections proceed under this mentality and present differing submodular optimization strategies that can be employed in phase one of the optimization.

5.6.1. Convex Closure. One approach for submodular optimization is to relax the need for a robust global or local maximum and instead expedite the optimization by discretizing the search region. This route still enables the schedules for multiple targets and observers to be optimized simultaneously, meaning each measurement is scheduled with knowledge of past and future observations and multiple observers are able to collaborate

their efforts in maintaining estimates of a multitarget system. Much like the simultaneous optimization presented in Eqs. (5.24a)–(5.24c), this high-dimensional problem becomes prohibitively expensive for larger schedules, but still maintains the ability to consider all the facets of the sensor tasking problem simultaneously from an expected information gain perspective; further, the computational complexity, as well as the quality of the resulting sensor schedule, can be controlled through the discretization of the search region and the number of iterations performed in the optimizer. In the case in which the $m_{z,i}$ highest-weighted candidate observation times are taken for each target, a relatively low convergence tolerance usually suffices in generating an optimal schedule, as the more information-rich candidate times are quickly discovered when compared to other, lower-quality measurement opportunities.

This search space discretization is achieved by generating candidate observation times and assigning an indicator variable to each of them. For the sake of development, recall the optimization variable \mathbf{t} in Eqs. (5.24a)–(5.24c); however, consider the individual schedules for each sensor/target pair to be fixed, e.g. $\mathbf{t}_{s,i}$ consists of candidate observation times from $t_0 = 0$ seconds to $t_f = 3600$ seconds in increments of 30 seconds for each sensor/target pair. Assign to each candidate observation time an indicator variable, defined as

$$\gamma_{s,i,k} = \begin{cases} 1 & \text{sensor } s \text{ observes target } i \text{ at time } t_k, \\ 0 & \text{otherwise.} \end{cases} \quad (5.25)$$

Let $\boldsymbol{\gamma}_{s,i}$ be the set of all indicator variables for the sensor s /target i pair, $\boldsymbol{\gamma}_s$ be the set of all of the vectors $\boldsymbol{\gamma}_{s,i}$ for sensor s , and lastly $\boldsymbol{\gamma}$ be the set of every indicator vector. This enables a new objective function to be defined as the convex closure of the previously proposed objectives, as $f(\boldsymbol{\gamma}) : \{0, 1\}^{d_2} \rightarrow \mathbb{R}$ (where d_2 is the number of candidate observation times) [22]. To be explicit, this new convex function is defined similarly to the previous objectives, except that the measurement model is redefined as the product of the measurement model

and the indicator variable $\gamma_{s,i,k}$, or [82]

$$\mathbf{c}_s(\mathbf{x}_k^{(i)}) \triangleq \gamma_{s,i,k} \mathbf{h}(\mathbf{x}_k^{(i)}) + \mathbf{v}_k. \quad (5.26)$$

When Eq. (5.26) and its Jacobian are used in the Kalman update in Eqs. (4.8a)–(4.9), it is clear that the measurement model Jacobian as well as the measurement residual is zero when $\gamma_{s,i,k}$ is equal to zero, leading to no update, as if the measurement is not processed.

While the dimension of this optimization is larger (as more candidate observations are required than the number of measurements to be scheduled in order to provide a degree of “choice” as to when to schedule measurements), this new function is convex over the boolean hypercube [22]. If the indicator variables are treated as continuous over the interval $[0, 1]$, the convexity can be deduced if one considers the fact that any additional measurements will increase the amount of information gained; more specifically, moving in a particular direction over the boolean hyper cube will result in either a net increase or a net decrease in information, and the function will exhibit either a non-decreasing or non-increasing behavior along that direction, respectively. As a result, the resulting function is convex in $\boldsymbol{\gamma}$, with the global max located at the point in which each indicator variable is unity. If the indicator variables are treated as continuous in the optimizer, the convexity can be leveraged and continuous optimizers employed to generate schedules that simultaneously consider every sensor task to be scheduled with relatively reasonable computational demands when compared to the direct continuous optimization of the measurement times. The resulting optimization becomes

$$\boldsymbol{\gamma}^* = \arg \max_{\boldsymbol{\gamma}} \iota(\boldsymbol{\gamma}), \quad (5.27a)$$

subject to

$$\gamma_{s,i,k} \in [0, 1] \quad (5.27b)$$

$$\mathbf{C}_{m_{z,t}} \boldsymbol{\gamma} = m_{z,t} \mathbf{1}_{(Lm_o)} \quad (5.27c)$$

$$\mathbf{C}_t \boldsymbol{\gamma} = \mathbf{1}_{(d_2/Lm_o)}, \quad (5.27d)$$

where $\mathbf{C}_{m_{z,t}}$ enforces a maximum number of measurements each sensor takes of a particular target, \mathbf{C}_t limits each sensor to taking at most one observation per time step, and $\mathbf{1}_a$ denotes an a -dimensional column vector of ones. Equations 5.27 provides a robust means of generating a sensor schedule that intelligently tasks all measurements with knowledge of one another. It is important to note, however, that the result almost definitely contains many values between zero and one, representing “partial” measurements of the target. These can be addressed as deemed appropriate for the application. For instance, it may be of use to define a threshold γ_{thresh} to determine which values should be rounded up to one and which should be set to zero, generating a new one-dimensional optimization problem, if so desired [82]. When setting a hard limit on the number of tasks to be generated, the top weighted values are taken, e.g. the top $m_{z,t}$ -weighted indicator variables are taken for each sensor/target pair such that the constraints in Eqs. (5.27c) and (5.27d) are met. The result is a polynomial-time solution for the multisensor/multitarget sensor tasking problem that simultaneously considers the schedules for all observers present when generating an optimal schedule.

5.6.2. Greedy Approaches. The convex closure approach to submodular optimization enables the sensor schedule to consider all future and past measurements when generating each individual sensor task, providing a schedule that is coordinated between different observers. The approach is highly dependent on how the optimization search space (i.e. the time interval over which a schedule is to be generated) is discretized; the finer

the time step, the more computationally complex the optimization is but the more robust the solution. In many applications, however, speed is of greater concern than getting every possible bit of information out of a sensor schedule.

In scenarios in which optimization time is of highest concern, greedy optimizers are appealing, as they approach a d -dimensional optimization in which each of the d channels may be dependent upon one another as d scalar optimizations. Individually optimizing each of the d decision variables results in a much faster optimization, as the potential dependencies between each of the decision variables are ignored, effectively limiting the optimization search area to find a locally optimal decision before moving forward [14]. Greedy solutions are particularly attractive in submodular optimization, as the topology of the underlying objective function provides guarantees on the level of optimality with regard to the generated solution. Specifically speaking, if \mathbf{Z}_k^* is a set of k measurement times obtained via greedy optimization, it is guaranteed that [47, 48, 67]

$$f(\mathbf{Z}_k^*) \geq (1 - e^{-k/d}) \max_{\mathbf{Z}} f(\mathbf{Z}), \quad (5.28)$$

where d is the cardinality of the globally optimal measurement set, providing not only a lower bound on the performance of the greedily-generated sensor schedule but also suggesting that to approximate the performance of an optimal measurement set of d measurements, additional measurements may be added to the greedy schedule to increase this lower bound. It is important to note, however, that continuing to add additional measurements tends to result in a “wait-and-see” schedule structure if the remaining portion of the time interval becomes shorter with each additional measurement. Due to the nature of the problem, greedily selected optimization solutions can be structured in a variety of ways [1]. This section seeks to briefly describe some of the intuitive approaches to greedy sensor tasking in the multitarget/multiobserver sensor tasking problem.

5.6.2.1. Greedy selection via observer. Recall $\mathbf{t}_s = [\mathbf{t}_{s,1}^T, \mathbf{t}_{s,2}^T, \dots, \mathbf{t}_{s,L}^T]^T$ is the set of times to observe each of the L targets for sensor s . The set of s optimal sensor schedules $\mathbf{t}_{0:s}$ is obtained in a greedy-in-observer fashion through concatenation of the previous $\mathbf{t}_{0:s-1}$ schedules with a newly optimized schedule, as

$$\mathbf{t}_{0:s} = \begin{bmatrix} \arg \max_{\mathbf{t}_s} \{ \iota(\mathbf{t}_s | \mathbf{t}_{0:s-1}) \} \\ \mathbf{t}_{0:s-1} \end{bmatrix}, \quad (5.29)$$

where \mathbf{t}_s is subject to the monotonically increasing constraint and the bounds given in Eqs. (5.24b) and (5.24c), respectively, but the dimension of the optimization problem is reduced to $d = L \times m_{z,i}$ as a result of optimizing each observer schedule independently; note that $\iota(\cdot | \mathbf{t}_{0:s-1})$ indicates the KL-based objective function evaluated given the previously optimized $s-1$ observer schedules, and the greedy optimization is initialized on the empty set (i.e. $\mathbf{t}_0 \triangleq \emptyset$). This formulation provides schedules that are optimal for each sensor (given previously optimized schedules), breaking down the optimization into more manageable problems while managing to maintain the consideration of observation placement with respect to other measurements that will be taken with that same model from the exact same position with respect to that observer. Additionally, this formulation enables particular observers to be emphasized through the order in which the schedules are optimized; if, for example, two observers are in need of a schedule, in selecting one of the schedules to optimize first, this schedule is generated without any influence from the other schedule. In this way, observations for the first observer can be prioritized by being scheduled first.

5.6.2.2. Greedy selection via target. Let $\mathbf{t}_{1:m_o,i}$ be a set of monotonically increasing measurement times of target i across all observers and time steps. Just as with the previous approach, optimization by greedy selection per target is achieved by appending the previously optimized schedules $\mathbf{t}_{1:m_o,0:i-1}$ (for targets 1 to $i-1$, with the zeroth schedule representing the initialization of the optimization scheme on the empty set, i.e. $\mathbf{t}_{1:m_o,0} \triangleq \emptyset$)

to the optimal schedule $\mathbf{t}_{1:m_o,i}$ for target i given the previous schedules and subject to the monotonically increasing constraints and the bounds in Eqs. (5.24b) and (5.24c), respectively. To be explicit, this approach is mathematically expressed as the concatenation of the $\mathbf{t}_{1:m_o,0:i-1}$ previously optimized schedules and a newly generated schedule, as

$$\mathbf{t}_{1:m_o,1:i}^* = \begin{bmatrix} \arg \max_{\mathbf{t}_{1:m_o,i}} \{t(\mathbf{t}_{1:m_o,i} | \mathbf{t}_{1:m_o,0:i-1})\} \\ \mathbf{t}_{1:m_o,0:i-1} \end{bmatrix}. \quad (5.30)$$

subject to the monotonically increasing constraint and the bounds given in Eqs. (5.24b) and (5.24c), respectively. Note that the dimension of the optimization is reduced $d = m_o \times m_{z,i}$ as a result of considering each target schedule independently, and should be reflected accordingly in the constraints.

Similar to before, this approach benefits from the ability to prioritize other targets by optimizing their schedules first, ensuring that when their generation occurs they are not held to any constraint infractions caused by previously generated schedules. Further, in the case in which the targets being tracked are independent of one another, this style of greedy approach is a good approximation of the optimal schedule, as the schedule of observations of target i does not impact the information-rich portions of the trajectory of target j ($j \neq i$), and consequently, the schedule for target i only influences the schedule for target j by dictating when sensor resources are unavailable. The fidelity of this approximation is apparent in the optimization runtime, as it still requires the coordination between multiple observers in generating a schedule; regardless, reducing the dimension of the optimization by considering each target independently results in a substantial decrease in computational expense.

5.6.2.3. Greedy selection via time. An interesting facet to the particular problem at hand arises when considering the myopic and forecasted formulations of the normalized expected KL divergence with respect to a greedy-in-time optimizer solution. When a single measurement is considered, the myopic and forecasted expected KL divergence objective

measures are equivalent in their evaluations, and thus result in equivalent optimal scheduled measurements in the presence of a single observation (or in a collection of observations of differing targets). Additional constraints should be placed on this particular solution; otherwise, the most direct approach is to consider the curve of the objective function for each sensor/target pair and take the global sensor/target maximum, set that particular time as the lower bound of the optimization, and repeat. Previous studies have shown that myopic implementations of the KL divergence tend to exhibit a “wait-and-see” structure in which sensors wait longer to take an observation [33], which, in this instance, leaves a smaller time interval over which to optimize the remaining measurement times. In this instance, it is useful to again discretize the time interval and optimize each individual time step.

Recall the convex closure formulation, and let β_k be the vector of the indicator variables associated with each candidate sensor action at time t_k . A greedy-in-time step approach sequentially progresses through each time step, selecting the observations that yield the most information as

$$\beta_{0:k}^* = \begin{bmatrix} \arg \max_{\beta_k} \{ \iota(\beta_k | \beta_{0:k-1}) \} \\ \beta_{0:k-1} \end{bmatrix}, \quad (5.31)$$

subject to the same constraints as the convex closure given in Eqs. (5.27b)–(5.27d) and initialized on the empty set ($\beta_0 = \emptyset$). By forcing decisions to be made at each of these discrete time steps, the wait-and-see structure is avoided, and the overall sensor schedule is developed in an intuitive and pragmatic manner. This approach typically exhibits the lightest computational load, but does not permit measurements at future times to be considered when making tasking decisions and simply coordinates efforts across multiple sensors and the multiple targets at each single instant.

5.6.2.4. Further consideration. As another consideration, it is important to note that the selection of one of the aforementioned greedy optimization solutions does not preclude the use of other solutions in tandem with the selected approach. For instance,

when a tasking policy is determined to be executed observer-by-observer, in generating the schedule for each observer, the objective function can be optimized normally without employing a submodular optimization technique. Alternatively, the optimization can be approximated via convex closure or generated greedily, i.e. target-by-target or time step-by-time step. Optimization of the original objective function with no approximations clearly gives the most robust schedules; however, the differing levels of approximation enable a sensor policy to be established with the desired level of robustness while considering the execution runtime when selecting an optimizer.

6. RESULTS

To fully evaluate the performance of the proposed sensor tasking policies derived in this work, two scenarios in which RSOs are tracked and maintained are considered. The first provides an analysis of the behavior of the forecasted approach as well as the divergence-based objectives computed in different spaces, applying brute-force optimization. This same optimization is applied to two conventional information-based tasking objectives to provide a benchmark against which to compare the proposed objectives. The second scenario emphasizes the developments in multitarget considerations and leverages the two-phase optimization strategy defined in Section 5.6. To test the different optimization schemes, a twelve target scenario is considered first with a space-based sensor being tasked. A relative motion model is utilized to leverage the linear dynamic model in an initial evaluation of the optimization strategies. Finally, the need of the submodular optimization techniques is illustrated by moving into nonlinear dynamics in the presence of multiple observers, where one of the observers is ground-based.

6.1. SINGLE-TARGET SCENARIO

The single-target tasking scenario is explored in an effort to illustrate the behavior of the proposed forecasting approach in the absence of other targets. This is investigated in this section in two orbital scenarios, namely an eccentric orbit to provide geometric diversity in the potential measurement sets and a circular orbit to remove this diversity.

6.1.1. Eccentric Orbit. In attempting to determine what facets a forecasted approach might emphasize in a satellite trajectory, a single target in an orbit with an eccentricity of roughly 0.68 is selected to provide some stark differences in the measurement profile geometry between different observation windows. The initial conditions for the elliptical orbit are generated from Keplerian orbital elements obtained from the North American

Aerospace Defense Command (NORAD) SO catalog¹ that are converted into the J2000 coordinate system [18], with the approximate position and velocity given as

$$\mathbf{x}_0 = \begin{bmatrix} \mathbf{r}_0 \\ \mathbf{v}_0 \end{bmatrix} \approx \begin{bmatrix} -232 \text{ km} \\ 10106 \text{ km} \\ 14 \text{ km} \\ -3.2 \text{ km/s} \\ 3.5 \text{ km/s} \\ 6.2 \text{ km/s} \end{bmatrix}. \quad (6.1)$$

The temporal evolution of the target states is governed by classical two-body mechanics, with the equations of motion given in Eq. (3.4). The state in Eq. (6.1) is used to generate the nominal trajectory for the sensor tasking routine as well as to initialize the mean of the target state estimate maintained in the extended Kalman filter, with associated uncertainties of 50 m (3σ) in each position channel and 0.2 m/s (3σ) in each velocity channel.

The satellite is propagated until two distinct observation windows (i.e. windows of time in which the satellite is in the field of regard of one of the observers) are generated. In the interest of providing a clear and concise analysis of the behavior of the forecasted measure, in the case of observation window overlap, the difference is split such that a clean hand-off occurs (i.e. only one observer can detect the satellite at a given instant); however, field of view overlap can be handled in any way deemed appropriate by the user, as the optimization for observation window overlap is not limited to the consideration of a single observer. The two observers considered here are located at the Canberra Deep Space Communication Complex (CDSCC; $35^\circ 24' 05''$ S, $148^\circ 58' 54''$ E) in Australia (denoted as Site *A*) and the Goldstone Deep Space Communications Complex DSN site (GDSCC; $35^\circ 25' 36''$ N, $116^\circ 53' 24''$ W) in California (denoted as Site *B*); the necessary transformation matrices required to rotate the fixed locations of these sites into the J2000 frame are obtained

¹Source: <https://celestrak.com/NORAD/>

via the NASA JPL NAIF SPICE toolkit² [6]. It is assumed that both of these stations are able to detect the satellite when it is at least 10° above the horizon, acquiring range and range-rate measurements via the incoming radio telemetry data according to Eq. (3.6). These measurements are corrupted with zero-mean uncorrelated Gaussian white noise with standard deviations of $\sigma_\rho = 15$ m for range and $\sigma_{\dot{\rho}} = 4$ mm/s for range-rate.

For the initial analysis, m_z measurement times are randomly sampled, drawn from a uniform distribution that spans the two observation windows; this is repeated 1000 times to generate sets of m_z measurements. The reference time is selected to be the final time the target is observable (i.e. $t_{\text{ref}} \approx 9.2$ hours), and each set is processed to generate an associated ${}_x\bar{M}_{KL}$ score (where the x pre-subscript denotes its computation in the state space). The set with the largest measure is then used to initialize the trust-region constrained optimization solver provided by the SciPy Python library. This is performed for $m_z = 5$ to $m_z = 40$ measurements in increments of 5 measurements with the constraint that all measurements must be within one of the observation window time intervals in an attempt to glean insight into the topology generated by the forecasted KL divergence computed in the state space and how that topology morphs when different observation geometries are considered (or more specifically, when the dimension of the optimization problem changes). The resulting sensor schedules are shown in Figure 6.1 alongside the measurement profile.

The most apparent trend in the different sensor schedules is the obvious emphasis near the end of the second observation window, aligning with a rapid change in both the range and range-rate profiles. The rest of the trajectory in this observation window seems to be relatively ignored until enough observation times are available, with the optimization scheme choosing to spread a majority of the rest of the measurement acquisition times over the first observation window. During this portion of the trajectory, the range-rate profile is nearly linear, while the range profile reflects this and decreases uniformly over the duration of the interval. As such, the geometric diversity in the measurement profile over

²Information for the SPICE toolkit can be found at <https://naif.jpl.nasa.gov>

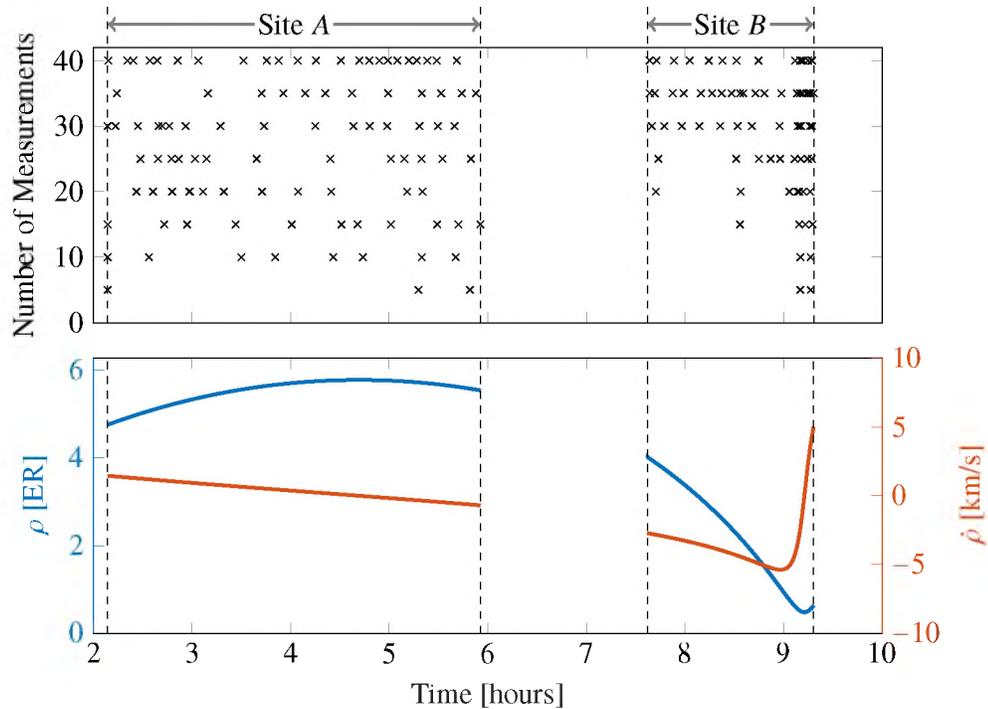


Figure 6.1. The resulting optimal schedules of varying sizes utilizing the forecasted objective (top) and the range and range-rate measurement profiles (bottom).

the observation window is scarce, producing an emphasis on acquiring observations before and after the maximum range until more observation times are available, at which point the observations are spread relatively uniformly over the interval. The collection of schedules suggests that the approach prioritizes rapid changes in the measurement profile, followed by regularly received observations, providing some insight into the underlying hypersurface the optimization is navigating. However, to observe the performance of this optimization criteria, it is best to compare estimation results against other, more conventional sensor tasking approaches.

A second simulation is performed, utilizing the same observation windows considered in the first simulation. However, as the point of interest in this instance is a performance comparison between differing optimization objective functions, the optimization initialization is instead executed as described in Section 4.2.3. That is, the state undergoes a Sundman

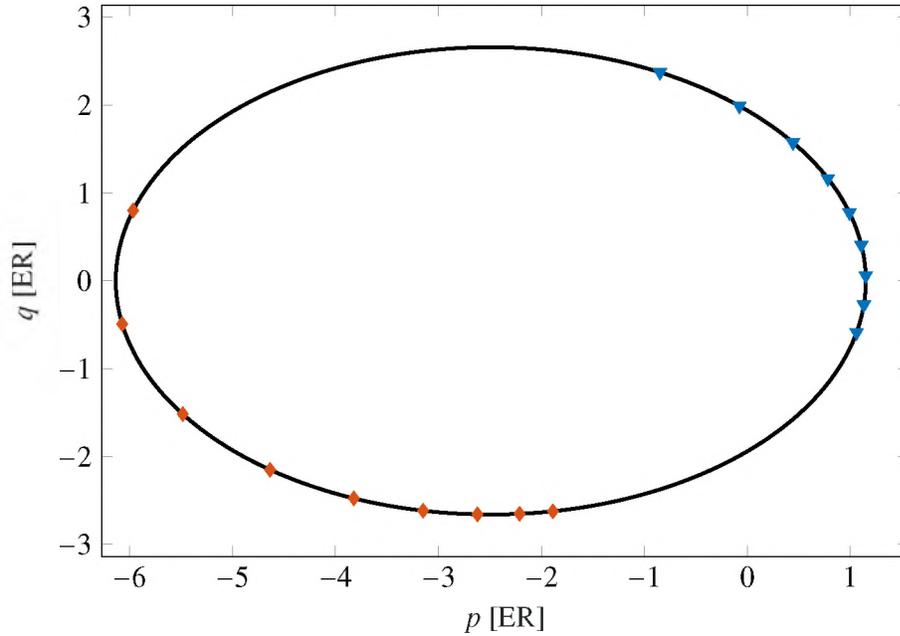


Figure 6.2. The sample states used in generating the initial guess for the optimizer.

transformation in order to generate sample states evenly dispersed over the trajectory of the orbit; the resulting trajectory and the sampled states are shown in the perifocal frame in Figure 6.2, with the different symbols indicating differing fields of view. Note that the roughly 63° inclination of this orbit prevents the Sundman transformation from generating evenly spaced states in this frame, but the sample states are much more uniformly distributed than they would be with a constant Δt .

Using the samples in Figure 6.2, candidate measurement sets consisting of three measurements are constructed for each sensor and evaluated against one another, with the highest scoring sets used to initialize the optimizer. These scores are computed for five different objective functions, namely the forecasted expected KL divergence computed in both the measurement and state space (${}_z\bar{M}_{KL}$ and ${}_x\bar{M}_{KL}$, respectively), the accumulation of the myopic expected KL divergence computed in both the measurement and state space (${}_zM_{KL}$ and ${}_xM_{KL}$, respectively), and the accumulated Fisher information gain (FIG). The

FIG objective is defined on the vector of measurement time indices ξ as

$$\mathbf{FIG}(\xi) = \sum_{\delta \in \xi} \mathbf{H}_\delta^T \mathbf{R}_\delta^{-1} \mathbf{H}_\delta.$$

The trust-region constrained optimization solver is again leveraged, optimizing the objective functions subject to bounds based on the “coarse” optimization solution (i.e. the solution’s neighboring candidate observation times) and the constraint that observation times must be within the observation window time intervals and in ascending order, i.e. for two observations z_i and z_j , it must hold true for the associated measurement times that $t_i > t_j$ if $i > j$.

The resulting schedules are depicted in Figure 6.3. Immediately apparent are the differences between the myopic and forecasted approaches; while perhaps intuitive, it might be easy to assume due to Eq. (4.42) that since the myopic and forecasted approaches should generate identical results for single measurement considerations, that this would extend to multi-observation sets. Clearly that is not the case here. However, note that there is a general consensus on taking observations during the end of the latter observation window when the measurement profile is changing rapidly. However, with regard to the first observation window, the myopic approaches agree that the earlier portion of the trajectory warrants attention, while the forecasted measures split their observations during the beginning and during the end of the interval. This spreading of the optimal data acquisition times illustrated in Figure 6.3 highlights the positive results that rigorously formulated forecasted information theoretic measures are capable of autonomously spreading observations across a long interval rather than succumbing to the “wait and see” mentality and bunching the measurements in a smaller portion of the trajectory as is exhibited by the myopic expected KL and FIG objectives.

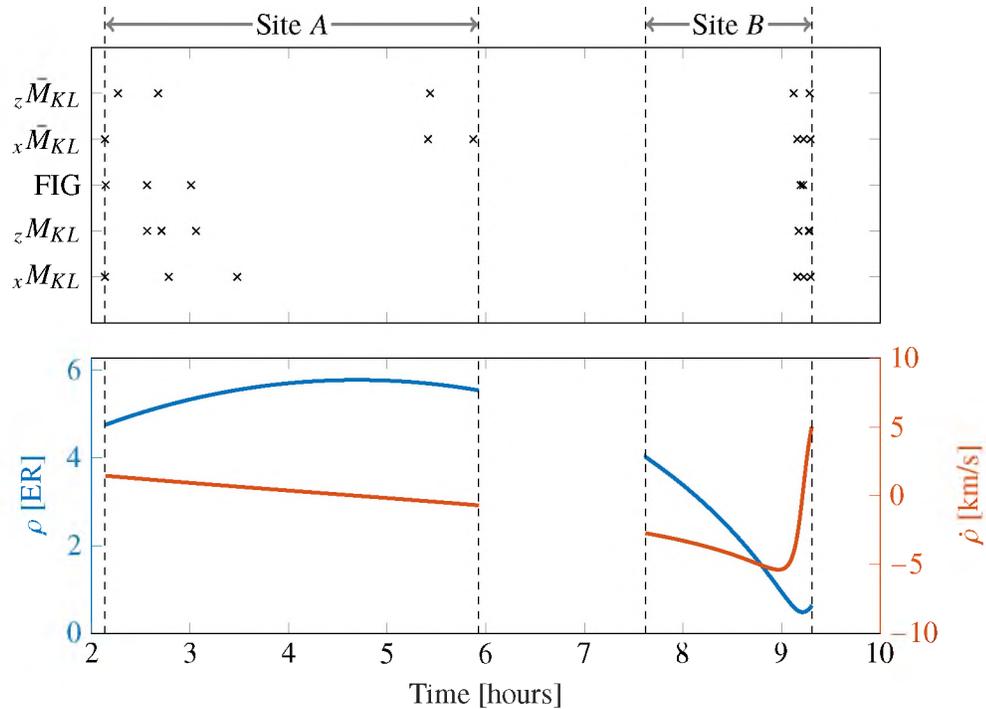


Figure 6.3. The differing sensor tasking schedules (top) plotted alongside the range and range-rate measurement profiles (bottom).

A simulation is then run for each objective function, executing the corresponding schedule and processing the resulting measurement sets in an extended Kalman filter, and the root-sum-square (RSS) values for the state estimate in both position and velocity are computed and tabulated. The RSS values corresponding to the moment before the object leaves the field of regard of a sensor, specifically the final observable time for the first observation window ($t_{1,f}$) and the reference time (i.e. the final observable time for the second observation window, $t_{2,f}$) are tabulated in Tables 6.1 and 6.2. Recall the pre-subscript notation indicates the space in which the measure is computed (x corresponding to the state space, z corresponding to the measurement space), the overbar notation indicates that the forecasted objective is employed, and the lack of an overbar indicates the myopic objective.

Table 6.1. The position RSS values (in meters) at the end of the two observation windows for the eccentric orbit scenario.

	${}_z\bar{M}_{KL}$	${}_x\bar{M}_{KL}$	FIG	${}_zM_{KL}$	${}_xM_{KL}$
$t_{1,f}$	263.07	251.84	323.56	332.16	335.02
$t_{2,f}$	19.42	18.13	21.73	24.37	22.45

Table 6.2. The velocity RSS values (in millimeters per second) at the end of the two observation windows for the eccentric orbit scenario.

	${}_z\bar{M}_{KL}$	${}_x\bar{M}_{KL}$	FIG	${}_zM_{KL}$	${}_xM_{KL}$
$t_{1,f}$	15.49	14.56	20.39	21.43	21.36
$t_{2,f}$	15.79	15.61	16.34	16.73	16.50

Immediately apparent is the obvious margin by which the forecasted methods outperform the FIG and conventional myopic KL in both position and velocity; this performance gain at the selected reference time $t_{2,f}$ is unsurprising, as that is how the approach is designed to operate. However, a noteworthy result is that the forecasted methods also yield an improvement over the myopic approaches at the end of the first observation window, suggesting that the reference time mapping is resulting in more efficient selection of observations over the first field of regard as well. However, as these approaches are designed to yield the best possible state estimate at the selected reference time from an information theoretic standpoint, it can be argued that measures computed at this reference time lead to a biased representation of the results. Considering the speculation on how the forecasted methods are performing over the first observation window with a reference time selected in the second, an interesting measure is the average entropy exhibited by the state estimate over the duration of the simulation. These results are presented in Figure 6.4.

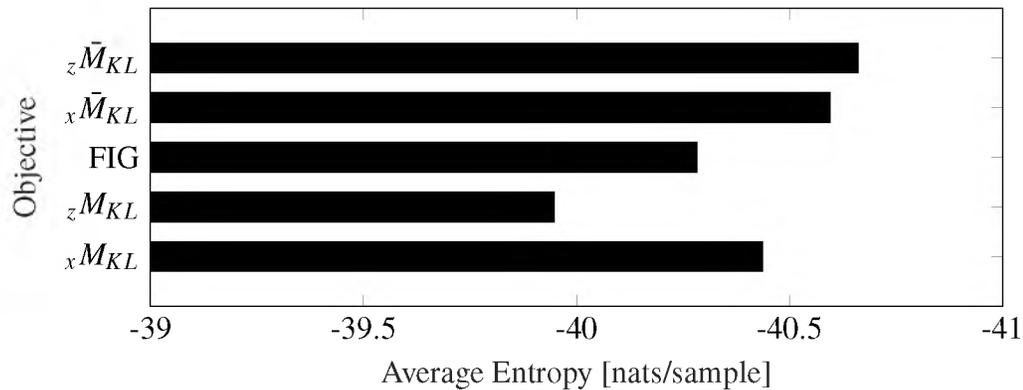


Figure 6.4. Average entropy over the simulation for each sensor tasking objective for the eccentric orbit scenario.

Recall that differential entropy can be negative and that the differential entropy for a continuous deterministic event approaches negative infinity. Thus, the more negative the differential entropy, the less uncertainty there is associated with the event. As such, the average differential entropy in Figure 6.4 is plotted in decreasing entropy so that values further to the right are associated with a more certain event, or there is, on average, more confidence in the target state estimate over the duration of the periods in which the target is detectable. Interestingly, not only do the forecasted methods perform well with respect to the reference time, but also seem to gain some benefit in their sensor schedule generation that enables them to exhibit overall more confidence in the target state estimate over the entire duration of the simulation. The FIG and conventional myopic KL approaches come close to performing similarly to the forecasted objectives; however, it is clear that in this instance that mapping the estimate densities to a single reference time and into a common information space provides more context to the sensor scheduling process, yielding noticeable benefits in the overall performance of the tracking solution.

The results here are encouraging, even when stepping away from the original design intent of this methodology and looking at the performance outside of the reference time for the forecasted methods. However, the orbit under analysis currently is highly eccentric

and highly inclined, providing some information-rich portions of the measurement profile that perhaps allow these methods to generate their optimal sensor schedules at low cost. It could be argued that the rapid changes in the measurement profile create pitfalls for myopic approaches with the intense geometric diversity present in the orbital trajectory. Performing an analysis in the absence of these geometric diversities may provide more insight into the performance of the forecasted approaches.

6.1.2. Nearly Circular Orbit. The previous example illustrates the different priorities exhibited by several sensor tasking objectives, as well as the differences between the myopic and forecasted approaches to KL-based sensor tasking. However, how do these different objectives compare when the measurement profiles between observation windows are nearly identical? More specifically, how are the measurements allocated if the geometric diversity in the two observation windows is relatively sparse? To explore this, a nearly circular orbit is investigated to provide similar measurement profiles from the two observers in the previous simulation. The initial conditions are again generated from Keplerian orbital elements obtained from the SO catalog³ and converted into the J2000 coordinate system [18], with the inertial position and velocity given approximately as

$$\mathbf{x}_0 = \begin{bmatrix} \mathbf{r}_0 \\ \mathbf{v}_0 \end{bmatrix} \cong \begin{bmatrix} 14525 \text{ km} \\ -9479 \text{ km} \\ 18713 \text{ km} \\ 3.25 \text{ km/s} \\ 0.90 \text{ km/s} \\ -2.06 \text{ km/s} \end{bmatrix}. \quad (6.2)$$

Once again the object is propagated until two observation windows are generated, one from the CDSCC DSN site (Site *A*) and another from the GDSCC DSN site (Site *B*). For this simulation, these sites are assumed to be able to generate right ascension

³Source: <https://celestrak.com/NORAD/>

and declination measurements of the satellite whenever it is at least 10° above the horizon (ignoring the effects of the sun to provide similar observation window durations, providing more similar observation arcs) according to Eq. (3.5), with the right ascension and declination measurement model defined in Eq. (3.8). These measurements are corrupted with zero-mean Gaussian white noise with standard deviations of $\sigma_\alpha = \sigma_\delta = 3''$.

Just as before, an initial analysis is performed using the forecasted KL divergence in the state space (${}_x\bar{M}_{KL}$), defining the final time prior to the object leaving the second observation window as the reference time. The first study utilizes a random initialization approach for the optimizer in an attempt to determine the topology generated by the objective function and to observe how the topology evolves as the dimension of the surface grows (i.e. as more observation times are simultaneously considered). The optimization is performed on measurement sets consisting of 5 to 40 measurements in increments of 5; the resulting observation sets are depicted in Figure 6.5 alongside the measurement profile. First note that the measurement profiles are very similar, implying that one arc should not provide more context over the other. Thus, these results illustrate how the forecasted approach performs in the absence of extreme geometric diversity present in the two observation windows.

Note that in both observation windows, each measurement set is taking observations during the portions in which the right ascension is changing most rapidly and when the magnitude of the declination is at its peak for the observation window. This is likely due to the information gained from the out-of-plane motion exhibited by the target, providing much more information than at lower declination points in the trajectory. Outside of these maximal (or minimal) declination points, the observation times for a given observation arc appear to be more or less uniformly distributed over the remainder of the trajectory. However, the latter observation window attracts many more observations than the former. There are a few different factors that could be at play here: the point of maximum declination

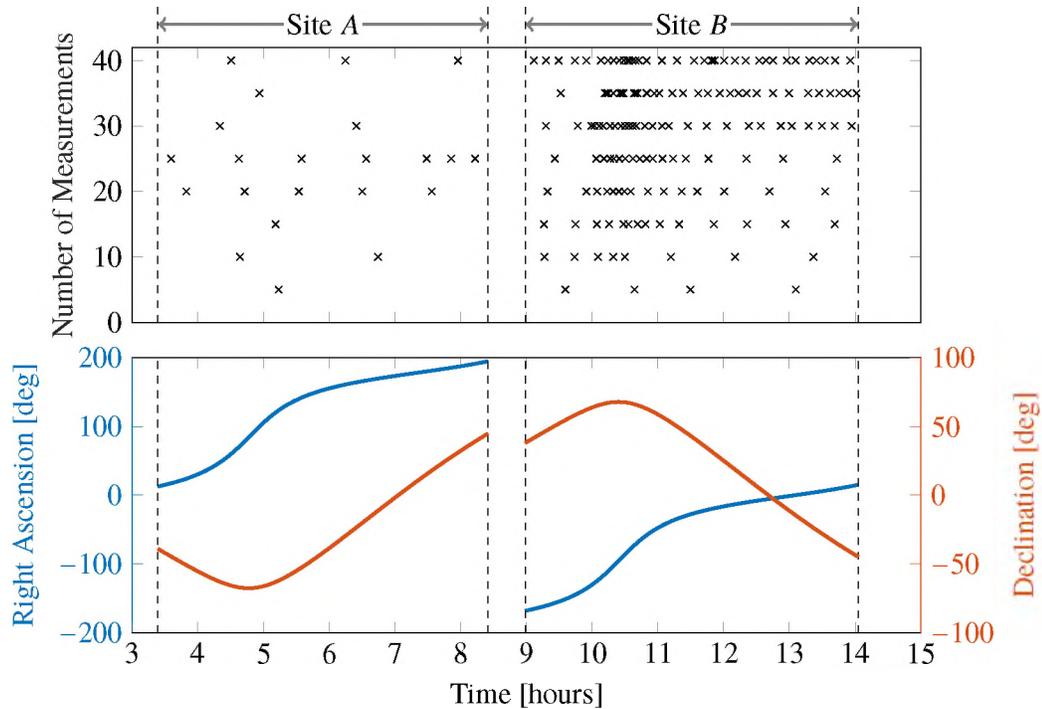


Figure 6.5. The resulting optimal schedules of varying sizes utilizing the forecasted objective (top) and the right ascension and declination measurement profiles (bottom).

magnitude here is marginally higher than in the first observation window, which could play a factor, or this could be a manifestation of wait-and-see behavior exhibited by myopic approaches. To investigate the latter claim, a comparison to myopic approaches is useful.

A second simulation is considered, utilizing the Sundman transformation for sampling the trajectory and generating a “coarse” initial schedule of three measurement times for each observer just as was done in the eccentric orbit simulation; the states used to generate the “coarse” initialization are shown in the perifocal frame in Figure 6.6. These sampled states are used to initialize the optimization for the same set of sensor tasking objectives in the eccentric orbit simulation; the optimized schedules are shown in Figure 6.7. The information divergence measures all seem to agree to some extent on observation times in the first arc with a few subtle differences, all generating substantially different schedules from the FIG-based approach. However, outside of the forecasted KL divergence computed

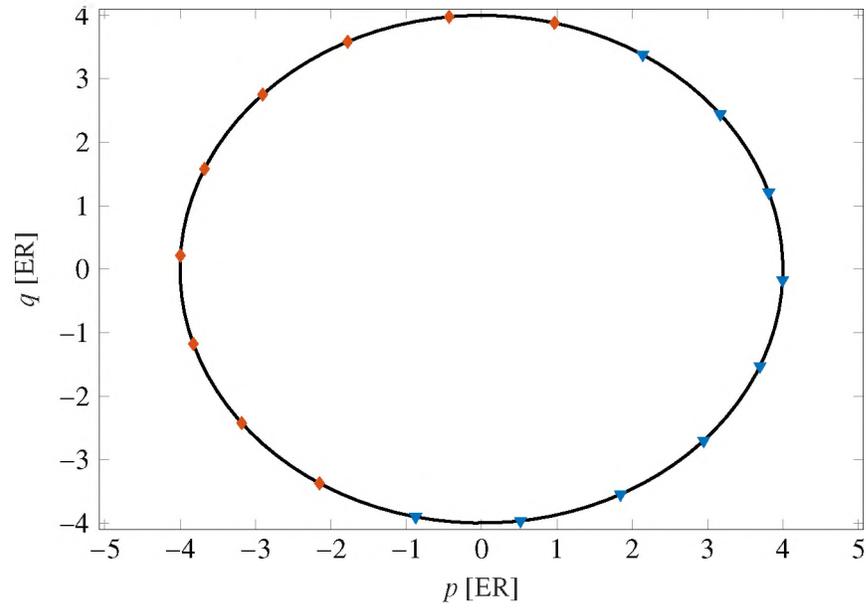


Figure 6.6. The sample states used in generating the initial guess for the interior-point optimization.

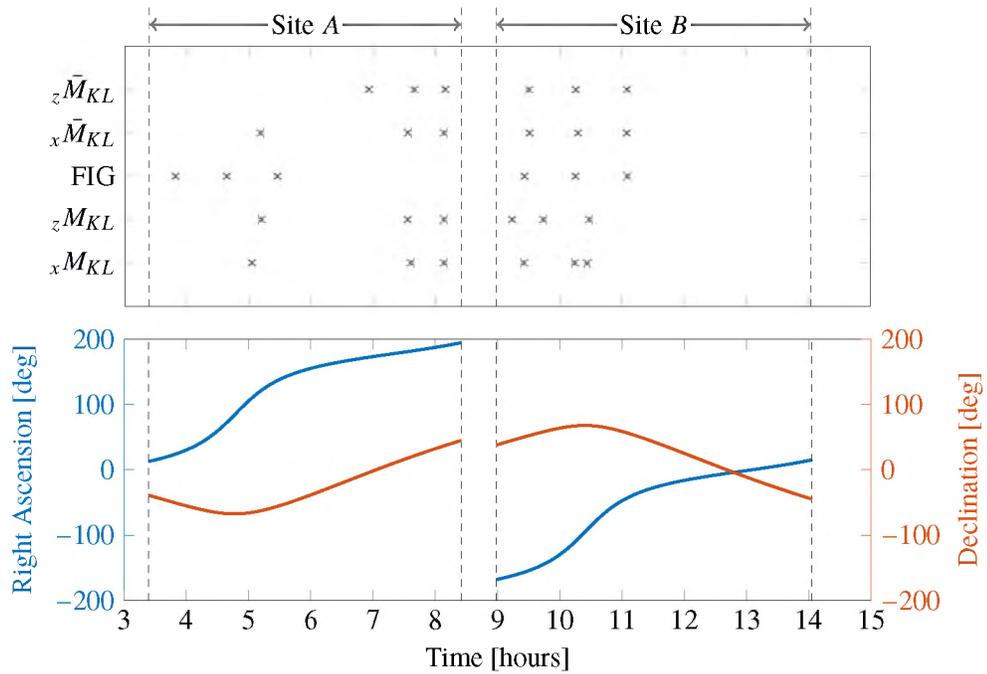


Figure 6.7. The differing sensor tasking schedules (top) and the right ascension and declination measurement profiles (bottom).

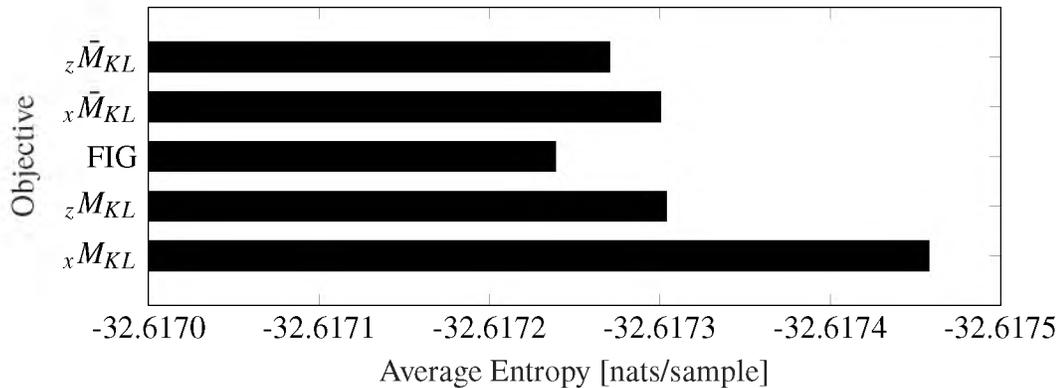


Figure 6.8. The average entropy over the simulation for each sensor tasking objective for the nearly circular orbit scenario.

in the measurement space, each objective ensures to take an observation near the lowest point in the trajectory, again favoring the out of plane motion, a trend that is exhibited by each of the schedules for the second observation window as well.

The resulting RSS values for each sensor schedule are given in Tables 6.3 and 6.4 at the final time of the first observation window as well as at the reference time (i.e. the final time of the latter window). The results across the different sensor tasking objectives are very similar, exhibiting almost negligible differences. This result is unsurprising, for while the peak declination points may provide slightly more information, the slow gradations of the measurement profile suggest that differing measurement times provide similar quantities of information. However, it is encouraging to see the slight gains in performance when considering the forecasted KL divergence in the state space while generating different schedules for each arc, moving away from the FIG strategy, which results in nearly identical measurement times for both arcs. Furthermore, as illustrated by Figure 6.8, it is evident that the average entropies in all of the information theoretic approaches outperform that of the FIG objective. The myopic divergences provide the best average entropy, while the forecasted approaches yield a compromise between average entropy and final RSS values.

Table 6.3. The position RSS values (in meters) at the end of the two observation windows for the nearly circular orbit scenario.

	$z\bar{M}_{KL}$	$x\bar{M}_{KL}$	FIG	zM_{KL}	xM_{KL}
$t_{1,f}$	7.8984	7.8984	7.8993	7.8984	7.8984
$t_{2,f}$	8.4852	8.4798	8.4834	8.4848	8.4813

Table 6.4. The velocity RSS values (in millimeters per second) at the end of the two observation windows for the nearly circular orbit scenario.

	$z\bar{M}_{KL}$	$x\bar{M}_{KL}$	FIG	zM_{KL}	xM_{KL}
$t_{1,f}$	1.0848	1.0848	1.0849	1.0848	1.0848
$t_{2,f}$	1.4406	1.4396	1.4402	1.4405	1.4399

6.2. MULTITARGET SCENARIO

The previous study emphasizes the behavior of the forecasted objective function, illustrating the emphasis on the uncertainty in a target state estimate at the reference time when compared to other, more conventional approaches. Moving into the multitarget tasking domain presents new challenges that must be considered. In the single-target study, brute force optimization forewarned of potential problems with optimization runtimes as the number of targets increases. Additionally, it is of interest to consider potential data association conflicts when generating these schedules.

To evaluate the proposed tasking policy presented in Section 5, a single space-based observer is taken to inhabit a circular orbit at an altitude of 400 km with an inclination of 75°. Twelve targets are represented in the Hill frame and propagated for one hour; during this time span, a sensor schedule consisting of ten measurements for each target is generated according to the policies outlined in Section 5.5. The initial conditions for the targets are

Table 6.5. The initial states for the 12 targets, given in the Hill frame.

Object	x [m]	y [m]	z [m]	\dot{x} [m/s]	\dot{y} [m/s]	\dot{z} [m/s]
$\mathbf{x}_0^{(1)}$	-16.150	18.15	0.0	-1.1110×10^{-3}	32.290×10^{-3}	0.0
$\mathbf{x}_0^{(2)}$	4.960	38.32	0.0	4.4160×10^{-3}	-11.240×10^{-3}	0.0
$\mathbf{x}_0^{(3)}$	5.396	18.64	0.0	-7.1510×10^{-3}	-7.732×10^{-3}	0.0
$\mathbf{x}_0^{(4)}$	-11.340	36.88	0.0	5.3270×10^{-3}	22.310×10^{-3}	0.0
$\mathbf{x}_0^{(5)}$	-8.869	23.99	0.0	-4.3620×10^{-3}	18.920×10^{-3}	0.0
$\mathbf{x}_0^{(6)}$	-14.300	27.12	0.0	-0.7016×10^{-3}	28.940×10^{-3}	0.0
$\mathbf{x}_0^{(7)}$	-3.075	14.79	0.0	-4.5900×10^{-3}	9.951×10^{-3}	0.0
$\mathbf{x}_0^{(8)}$	-0.1675	18.95	0.0	-5.9170×10^{-3}	3.203×10^{-3}	0.0
$\mathbf{x}_0^{(9)}$	-6.249	39.16	0.0	7.5100×10^{-3}	7.478×10^{-3}	0.0
$\mathbf{x}_0^{(10)}$	12.000	-7.00	19.0	9.3460×10^{-3}	-3.438×10^{-3}	2.550×10^{-3}
$\mathbf{x}_0^{(11)}$	-2.000	2.00	1.0	78.3000×10^{-3}	-2.566×10^{-3}	6.444×10^{-3}
$\mathbf{x}_0^{(12)}$	-22.5000	17.700	-23.3	20.8300×10^{-3}	44.230×10^{-3}	19.070×10^{-3}

listed in Table 6.5. Nine of the targets exhibit planar motion in order to generate several collision events in the measurement space, while the targets exhibiting out-of-plane motion exhibit measurement space collisions with either each other, the planar targets, or both.

The Clohessy-Wiltshire model is used to model the dynamics of the objects, such that, for a state defined as the relative position and velocity of the form $\mathbf{x}_k^{(i)} = [(\mathbf{r}_k^{(i)})^T, (\mathbf{v}_k^{(i)})^T]^T$, the dynamics of the state are given by the linear, discrete-time, noiseless system given in Eq. (3.2) where the motion is governed by the Clohessy-Wiltshire equations in Eq. (3.3). Note that this model is only one of many available relative motion models, and is only selected for its linear nature, enabling simple state propagation and facilitating faster generation of results and thus affording accessible analysis across more optimization solutions. Other options for relative motion propagation are available [28, 41, 53, 79] and would apply to this approach as well.

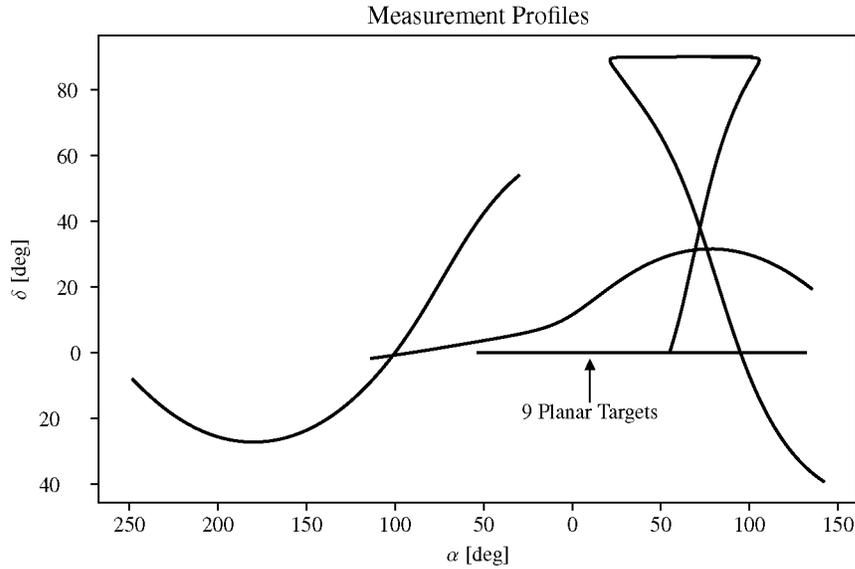


Figure 6.9. The measurement profiles of the twelve targets.

It is assumed that during this hour long time period, right-ascension and declination measurements are available, generated according to Eq. (3.5) and Eq. (3.8). These measurements are corrupted with zero-mean Gaussian white noise with standard deviations of $\sigma_\alpha = \sigma_\delta = 3''$. In order to illustrate the occurrence of collisions in the measurement space, the measurement profiles for each nominal target trajectory are plotted in the measurement space and can be observed in Figure 6.9. Note that the horizontal trajectory contains the nine targets in planar motion.

6.2.1. Continuous Optimizer Configurations. Five different optimizer architectures are applied to the time interval, and the resulting schedules are compared with one another for both the myopic and forecasted KL divergence objective functions as well as the spatial entropy. The two are simultaneously optimized as in Eqs. (5.24a)–(5.24c) with an “upper-bound” on the expected information gains set to 99.99% to determine the efficacy and effectiveness of handling the full optimization and to give a baseline comparison for the two-phase optimizations. For the two-phase solutions, the KL divergences as in Eqs. (5.1) and (5.2) are maximized in the first phase. A continuous optimization is per-

formed to provide a comparison to the submodular solutions, which include convex closure of the objective function and two greedy-in-target solutions (one performing a continuous optimization and another utilizing the convex closure form of the objective). These optimizations are then followed by a second phase in which the collision entropy in Eq. (5.17) is minimized at each measurement time.

For each optimization, a sequential least squares programming (SLSQP) optimizer is employed,⁴ with the convergence criteria based on a function tolerance of 1×10^{-16} for each objective function aside from the convex closure formulation. For the convex objectives, candidate observation times are considered every 30 seconds for each target over the span of the hour, resulting in 120 indicators to consider for each target and a 1440-dimensional optimization problem. The large dimensionality of the objective can make the optimization relatively burdensome. Fortunately, due to the geometry of the objective and the nature of the approach, a conventionally “strict” convergence criteria is not necessary, as the highest weighted measurement times are apparent as the optimizer enters the vicinity of the local extremum and do not change as the function tolerance decreases. As a result, the convergence criteria for the convex closure approaches is set to 1×10^{-6} , and the ten highest weighted measurement times are taken to generate schedules of ten measurements per target.

6.2.2. Optimization Runtimes. The normalized runtimes for each solution, namely the simultaneous optimization of the expected KL divergence and the coalescence term (denoted Simult.), the two-phase solution in which both terms are optimized via the SLSQP optimizer (denoted 2-P Cont.), the two-phase solution leveraging the convex closure of the expected KL divergence (denoted 2-P Conv.), the two-phase greedy-in-target solution in which both terms are optimized via the SLSQP optimizer (denoted 2-P G.I.T. Cont.), and the two-phase greedy-in-target solution leveraging the convex closure of the expected

⁴Initial implementations compared the trust-region constrained optimizer (used in the single-target simulation) against the SLSQP optimizer. The resulting schedules were similar in quality, but the SLSQP solver converged with faster runtimes. As a result, it was selected for use in the analysis

Table 6.6. The myopic objective function runtimes for each of the optimization schemes normalized by the simultaneous optimization.

	Simult.	2-P Cont.	2-P Conv.	2-P G.I.T. Cont.	2-P G.I.T. Conv.
Phase I	1.000	0.113	0.037	9.18×10^{-4}	46.1×10^{-3}
Phase II	-	3.531×10^{-4}	0.012	1.45×10^{-4}	1.38×10^{-4}
Total	1.000	0.113	0.048	1.063×10^{-3}	46.1×10^{-3}

Table 6.7. The forecasted objective function runtimes for each of the optimization schemes normalized by the simultaneous optimization.

	Simult.	2-P Cont.	2-P Conv.	2-P G.I.T. Cont.	2-P G.I.T. Conv.
Phase I	1.000	0.056	0.348	4.309×10^{-3}	0.123
Phase II	-	0.005	0.356	4.689×10^{-3}	0.005
Total	1.000	0.062	0.703	8.997×10^{-3}	0.128

KL divergence (denoted 2-P G.I.T. Conv.), are tabulated in Tables 6.6 and 6.7 for both the myopic and forecasted objective functions employed in the different optimization solutions, respectively. The overall runtime is broken down into the respective phases of the optimization, with each entry normalized against the simultaneous optimization runtime; it is worth noting that the simultaneous optimization consisting of the myopic KL divergence takes over thirty times longer than that of the forecasted approach to generate a schedule, a trend that is consistent across all solutions when comparing the myopic and forecasted measures. The rest of the table details the runtimes of the two-phase optimization. For this particular scenario, the two-phase continuous optimization scheme seems feasible, exhibiting substantial benefits when compared to the simultaneous optimization. However, this is largely due to the linear dynamics of the system and the (relatively) low number of targets to track. In the presence of nonlinear dynamics and/or when the dimension of the optimization argument is sufficiently large (e.g. additional measurements, targets, or observers), this approach proves to be far more computationally demanding.

The final four columns in Table 6.6 and 6.7 summarize the results obtained by applying the same measurement-by-measurement continuous optimization for the second phase, but leverage the submodularity in the first phase. As expected, as more submodular optimization techniques are utilized, the runtimes tend to decrease, with a few exceptions. The benefits of the linear dynamics are once again present in the continuous approach when optimizing target schedules independently, and while it may not hold under nonlinear dynamics or higher dimensional optimizations, this is useful as the only loss of fidelity is the fact that each subsequent target simply cannot be observed at previously tasked times. This is not a particularly detrimental constraint, considering that if a target can produce a lot of information at a previously tasked time, then the measurement can be placed some small amount of time before or after that action. The second exception to the decreasing trend is the first phase of the myopic greedy-in-target convex closure approach; the overall runtime is faster than that of the convex closure when optimizing all of the target schedules simultaneously. However, the speed benefits in each individual iteration of the optimizer do not outweigh the fact that the optimization must be performed 12 times (once for each target).

In short, the two phase optimization is clearly more practical than simultaneously attempting to maximize the information available while also avoiding measurement space collisions. Additionally, while the continuous optimization solutions exploit the linear dynamics and exhibit relatively low runtimes, it is possible to see how the submodular approaches can be selected to reduce the runtime by decomposing the optimization into different, less demanding optimization problems. It is important to remember that while every attempt is made to provide an even comparison across these methods, the convex optimization is extremely dependent upon how the time vector is discretized; more samples across the time interval result in a longer runtime, but better consider all points along the

trajectory. Regardless, the runtimes presented in Tables 6.6 and 6.7 provide a data point in analyzing the different sensor tasking solutions. In order to draw further conclusions, the performance of the resulting schedules must be examined.

6.2.3. Collision Avoidance. Considering the dimension of the optimization, it is difficult to visualize the expected information gain objective function surface and exactly how the schedule is being generated in phase one. Phase two is marginally easier, as the objective can be plotted as a function of time. An attempt to illustrate this is provided in Figures 6.10 and 6.11, where the phase one and two schedules (with measurement times denoted by \times) are plotted along with the collision avoidance objective (that is to be minimized) computed in the absence of any measurement updates (plotted in gray). The collision avoidance objective, η , indicates that in the absence of updates, the overall pdf coalescence generally increases as the targets are propagated, peaking at 1500 seconds and then separating until roughly 2800 seconds before coalescence increases again. Some movements can be detected in response to these collisions when comparing the phase two schedules to their phase one counterpart; both the myopic and forecasted schedules place more measurements at the minimum near the beginning of the simulation, while also moving some measurements away from the increase at the end of the time interval.

This gives a rough idea of how the second phase refines the sensor schedules, as the objective, η , indicates a measure of relative proximity of the targets, but this plot does not indicate the “updated” collision measure as the previous measurements are processed. It is useful evidence that the objective function provides times in which other tasks could be allocated (i.e. the relatively long stretches in which the sensor is not tasked), a trend that is exhibited across all other schedules (aside from the simultaneously optimized myopic divergence); however, aside from this, Figures 6.10 and 6.11 do not communicate much.

Since the optimization is performed on each measurement sequentially, each measurement time necessarily maintains or reduces the collision level it is observing. Consequently, it may be of more use to see, on average, how much this objective is reduced.

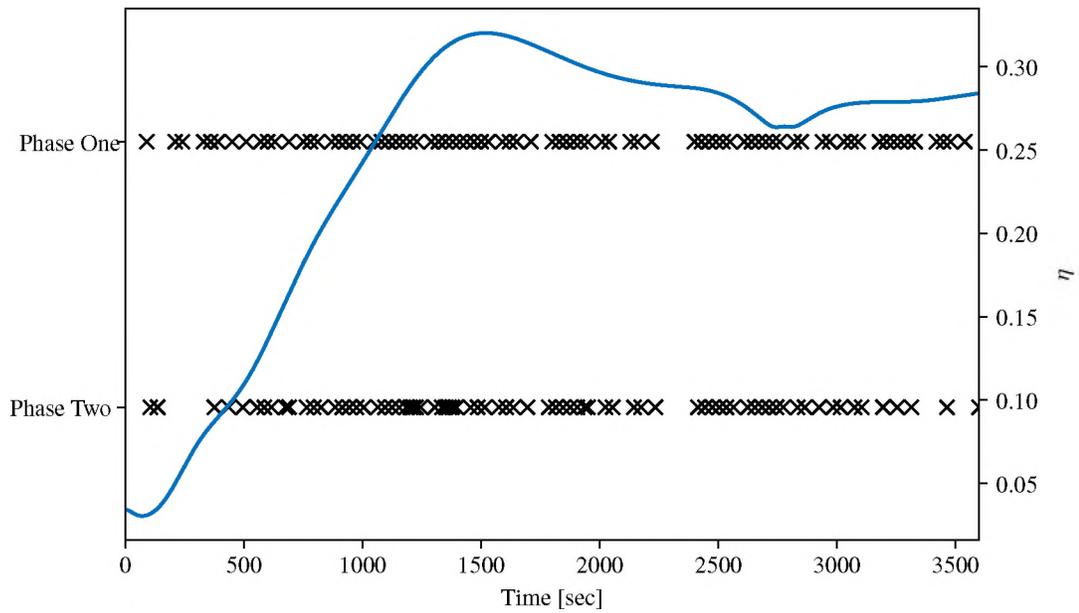


Figure 6.10. The phase one and two schedules generated via the myopic convex closure objective function optimized target-by-target, plotted with the collision avoidance measure.

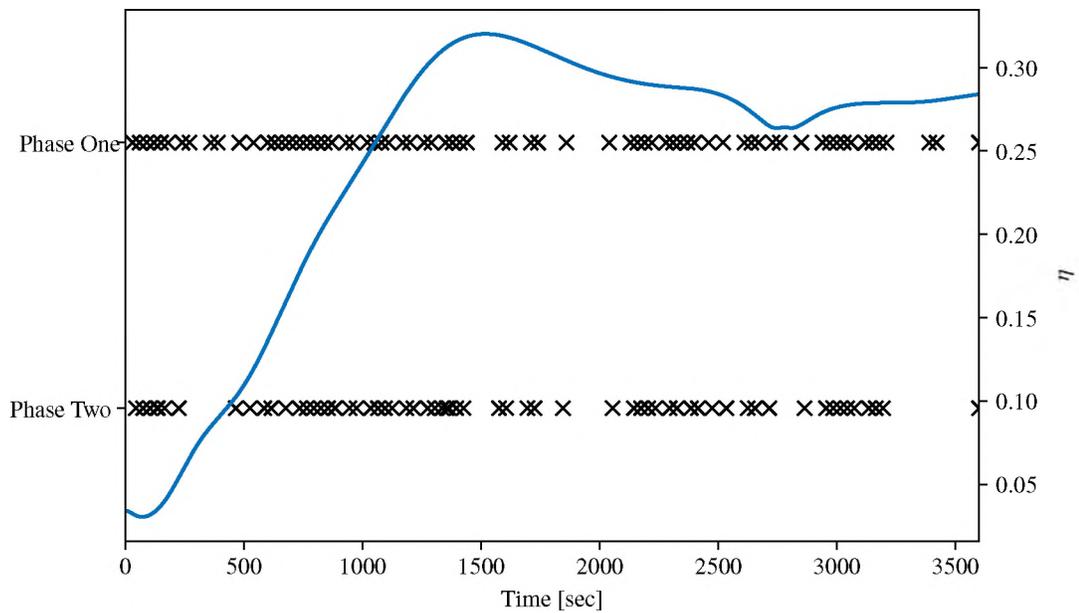


Figure 6.11. The phase one and two schedules generated via the forecasted convex closure objective function optimized target-by-target, plotted with the collision avoidance measure.

The average value over each measurement time is shown for both phase one and phase two in Figures 6.12 and 6.13. In each instance, the two phase optimization results in a lower collision measure after the second phase, with each second phase schedule exhibiting a lower measure when compared against the simultaneous optimization solution aside from the second phase of the continuously optimized forecasted measure.

While it may be tempting to draw conclusions from the magnitudes of these measures across the different objective types (myopic versus forecasted) as well as the different optimization solutions, it is important to note that this measure is independent of the first phase, so a particular trend cannot be guaranteed (i.e. it is coincidental that the continuous optimization of the myopic objective ends up with a relatively low collision measure in phase one). The main results exhibited by Figures 6.12 and 6.13 is that the objective is successful in reducing these potentially difficult data association scenarios as well as the fact that optimizing the two terms separately does not preclude similar performance with respect to this collision avoidance when compared to a solution that optimizes both objectives simultaneously. The results are still encouraging, however, considering that this optimization is performed second and thus is subject to most of the performance degradation due to the constraints placed on the optimization by the phase one schedule.

6.2.4. Schedule Performance. In order to better evaluate the performance of these objective functions, it is necessary to observe the tracking performance as the resulting sensor schedules are executed. Given that the objective measures are developed in information space, and in the interest of concise analysis, the sum of the individual target entropies is used to quantify the uncertainty present in the multitarget state. This is directly a function of the expected information gain objective, and since this is optimized prior to considering measurement space collisions in the two-phase approaches, one would expect that performing a continuous optimization with respect to solely the information gain term should exhibit the same if not better results as the simultaneous optimization. To verify this, each schedule is processed in an EKF separately, and the sum of the individual target entropies is logged

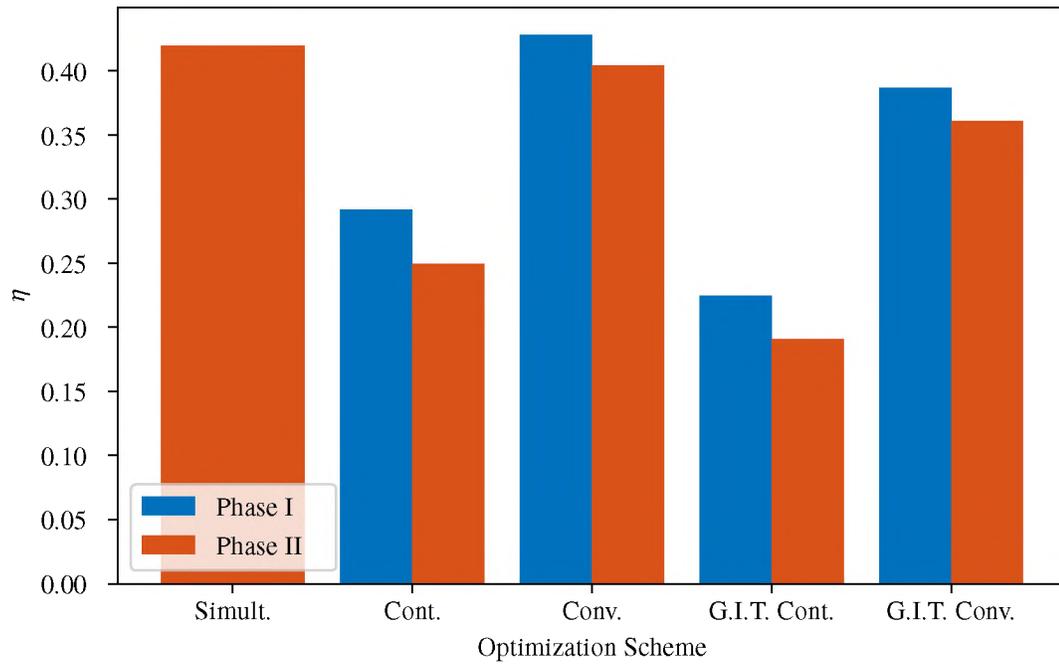


Figure 6.12. The average collision measures at each measurement time for the different myopic schedules.

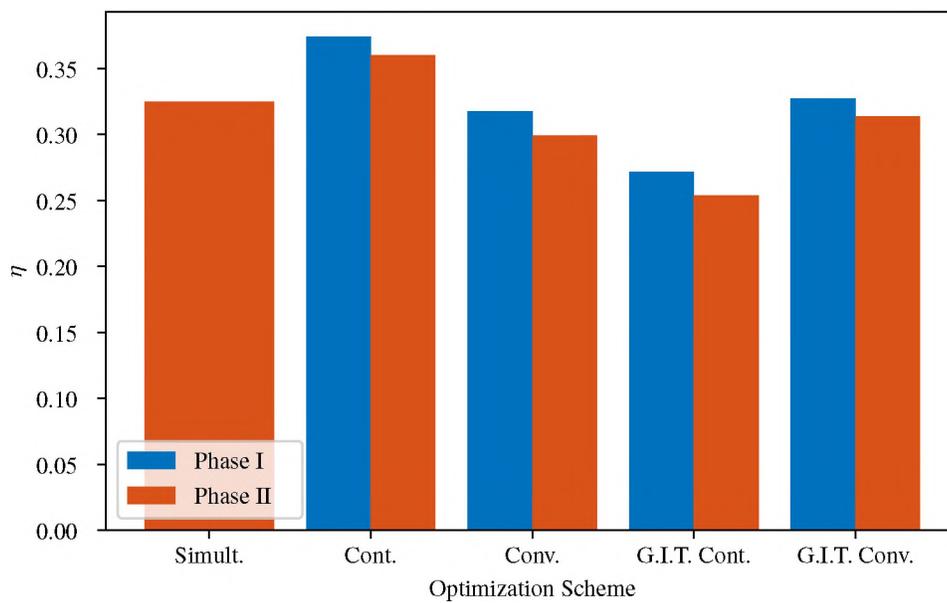


Figure 6.13. The average collision measures at each measurement time for the different forecasted schedules.

over the duration of the simulation. The results for both of the simultaneously optimized schedules are plotted alongside their continuously optimized first-phase counterparts in the two plots in the top left of Figure 6.14, with the simultaneous optimization plotted in blue and the continuous plotted in orange. What these curves specifically show is the sum of the Shannon entropy for the individual targets at each measurement index before and after a measurement is processed, enabling the contribution of each individual measurement to be observed more clearly than if it were plotted against time. Both of the schedules perform similarly across the time interval, with the continuously optimized myopic schedule just barely outperforming the simultaneous schedule at the end of the simulation. The two forecasted schedules perform similarly, but with the continuous optimization of the objective function maintaining a lower system entropy throughout most of the time interval.

Figure 6.14 goes on to illustrate the differences in performance when different sub-modular optimization strategies are applied. The two plots in the top right show the convex closure forms of the objective functions and their performance in approximating the solution in the continuous optimization. The forecasted approach achieves this approximation closely, whereas the myopic suffers somewhat noticeably over the hour long interval. This performance is absolutely impacted by local minima that exist within the objective function, but recall that the convex closure is also extremely sensitive to the discretization of the time vector; without enough candidate observation times, information-rich portions of a target trajectory are missed.

In an attempt to reduce the computational burden of optimizing the convex closure and marginally improve the initialization of the optimization, target independence can be leveraged and each target schedule can be initialized and optimized independently. This greedy approach is first applied with the expected information gain optimized as a continuous objective function in the bottom left two plots of Figure 6.14. Note that the performances match much closer here, though this could also be a benefit brought on by the continuous optimization which, recall, exhibits fast performance due to the linearity of

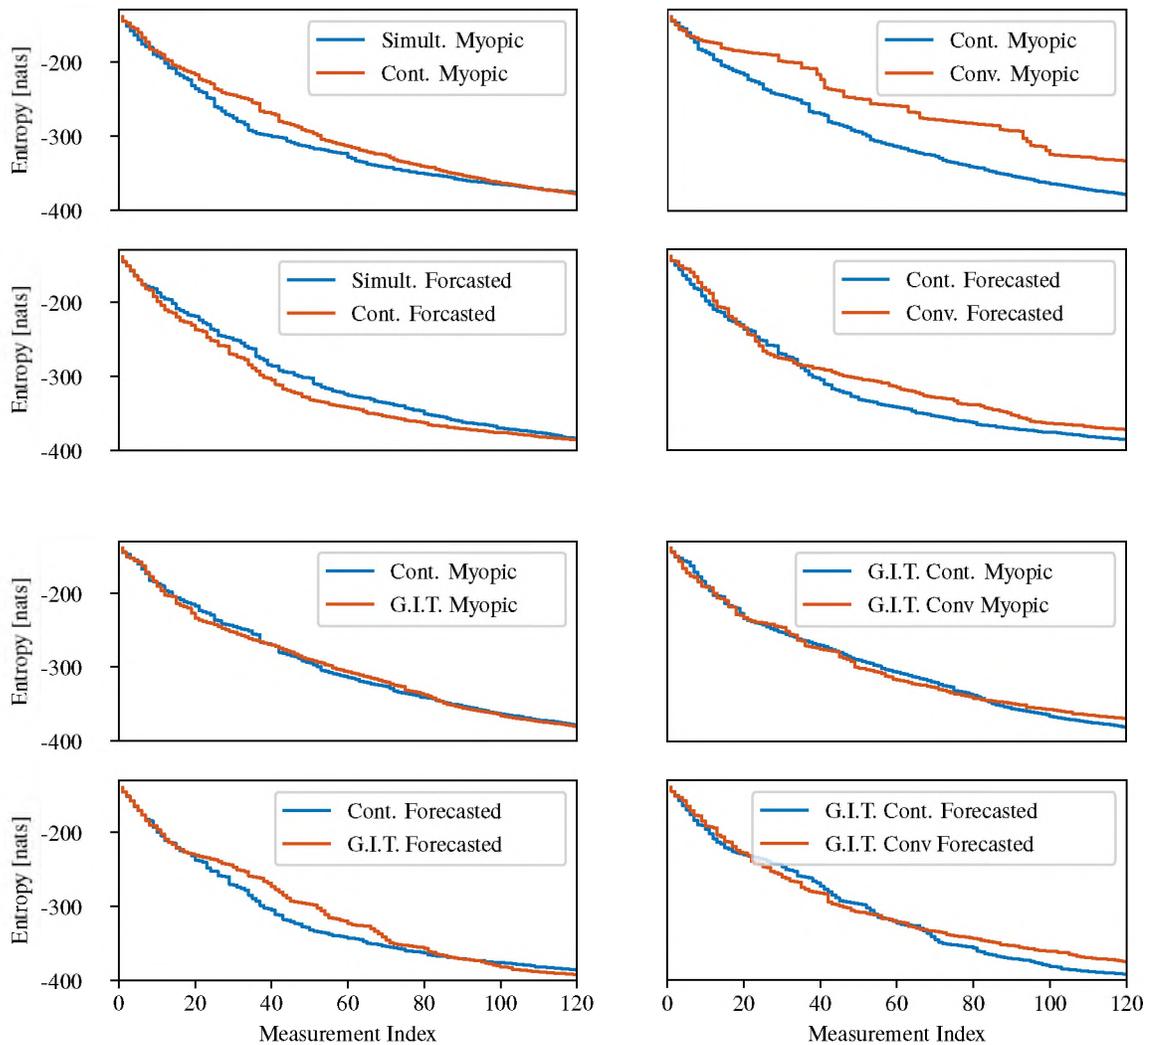


Figure 6.14. The resulting myopic and forecasted phase one schedules for each of the different submodular optimization schemes as compared to the continuous optimization of the first phase as well as the simultaneous optimization.

the dynamics. A more practical approach is to optimize each target schedule independently with the convex closure form of the objective function, as is shown in the bottom right of Figure 6.14. Once again, both optimization strategies perform relatively similarly, with some minor information loss brought on by missed measurement opportunities resulting from the target-by-target optimization. However, the end result is an optimization scheme that affords relatively fast runtimes when compared to other solutions.

Table 6.8. The system entropy (in nats) at the end of the hour long time interval and the different phase two sensor schedules have been processed.

	Simult.	Cont.	Conv.	G.I.T. Cont.	G.I.T. Conv.
Myopic	-376.91	-384.47	-351.95	-385.08	-374.97
Forecasted	-384.13	-387.13	-374.79	-393.40	-382.89

The final question to address regarding these schedules is the impact phase two optimization has on the schedule performance. Since the second phase is only capable of shifting individual measurement times, due to the constraints, these times are not altered drastically, and thus the performance over the duration of the simulation appears very similar. For conciseness and to provide a side-to-side comparison of the schedules, the system entropy at the end of the time interval is tabulated for each schedule in Table 6.8. As expected, the continuous optimization schemes perform the best from an entropic perspective. However, the submodular strategies provide a good approximation to this performance, and illustrate that further reduction of the dimension of the optimization problem (namely, when performing greedy optimization) can actually be beneficial from an entropic perspective in addition to reducing runtime.

6.2.5. Comments on Multiple Observers. Lastly, the effectiveness of the submodular optimization strategies is investigated in a higher-dimension problem. To achieve this, the previous simulation is extended to include a ground-based observer to coordinate with the space-based sensor in maintaining the target state estimates in the Earth-centered inertial frame. This ground-based observer is taken again to be the GDSCC site in California from the single-target study, generating range and range-rate measurements corrupted with zero-mean uncorrelated Gaussian white noise with standard deviations of $\sigma_\rho = 15$ m and $\sigma_{\dot{\rho}} = 4$ m/s for range and range-rate, respectively. Recall that the space-based sensor is capable of generating right-ascension and declination measurements corrupted with zero-mean Gaussian white noise with standard deviation $\sigma_\alpha = \sigma_\delta = 3''$. The initial relative states

for the targets remain the same from the previous example, and the space-based observer is initialized at the beginning of the hour-long time interval with Cartesian state

$$\mathbf{x}_{SS}(t_0) \cong \begin{bmatrix} -6416 \text{ km} \\ -122 \text{ km} \\ 2182 \text{ km} \\ 1.13 \text{ km/sec} \\ -6.99 \text{ km/sec} \\ 2.93 \text{ km/sec} \end{bmatrix},$$

where the SS superscript denotes the state is that of the space sensor. This results in a period of roughly five minutes for each target in which the objects are visible from the ground-based sensor, occurring right at the beginning of the simulation.

Initial simulations made it immediately apparent that the nonlinear dynamics and the overlapping observation windows for both of the sensors rendered the simultaneous optimization, as well as the continuous optimization of the expected information gain, impractical, taking an exorbitant amount of time and failing to converge when relaxing the convergence criteria. In order to generate schedules to determine some idea of the multi-observer performance, the two-phase optimization is necessary. The convex closure form of the expected information gain is utilized, both in optimizing each observer schedule independently (in this instance, optimizing the space-based schedule first, followed by the ground-based schedule) in a greedy-in-observer (denoted G.I.O.) approach as well as performing this optimization target-by-target (denoted G.I.T.).

Through leveraging the submodular optimization strategies, the SLSQP solver is able to converge to solutions for both the myopic and forecasted divergences, yielding schedules in reasonable runtimes. Independently optimizing the convex closure of the expected information gain provides a tractable solution, as compared to the simultaneous optimization. Further, an additional reduction in runtime is achieved when optimizing each

Table 6.9. The phase one optimization runtimes for the greedy-in-observer (G.I.O.) and the greedy-in-observer and greedy-in-target (G.I.O. G.I.T.) convex closure objective functions, normalized against the G.I.O. convex closure runtime.

	Myopic		Forecasted	
	G.I.O. Conv.	G.I.O. G.I.T. Conv.	G.I.O. Conv.	G.I.O. G.I.T. Conv.
SS	0.34	0.21	0.19	0.16
GDSCC	0.66	0.03	0.18	0.17
Total	2703	0.24	0.37	0.33

Table 6.10. The system entropy (in nats) at the end of the hour long time interval for each convex closure optimization, optimizing the space-based sensor schedule first followed by the GDSCC.

	G.I.O. Conv.	G.I.O. G.I.T. Conv.
Myopic	-568.7	-523.7
Forecasted	-574.7	-568.4

target schedule independently, with nearly a 76% reduction for the myopic measure and a roughly 12% reduction for the forecasted; the specific optimization runtimes are tabulated in Table 6.9. While some scheduling decisions are made differently when each target schedule is generated, the end result is still relatively similar as illustrated by Table 6.10.

The introduction of the second observer, aside from introducing new measurements and new information, largely does not impact the sensor tasking procedure. The schedules are coordinated to some extent, as a previously optimized observer schedule must be considered when generating additional observer schedules, but the construction of the objective and the execution of the optimization remains constant otherwise. However, much like the greedy-in-target approach, a sort of prioritization occurs when selecting the order of the sensor schedule optimizations in a greedy-in-observer approach. Though this prioritization is much less influential when compared to the greedy-in-target approach, as

Table 6.11. The system entropy (in nats) at the end of the hour long time interval for each convex closure optimization, optimizing the GDSCC schedule first followed by the space-based sensor.

	G.I.O. Conv.	G.I.O G.I.T. Conv.
Myopic	-568.7	-547.2
Forecasted	-578.3	-574.5

one observer's schedule does not prevent another from making tasking decisions. This idea is illustrated in Table 6.11, in which the order of the observer schedule optimizations is switched and the GDSCC schedule is generated prior to the space-based observer.

While some differences exist between Tables 6.10 and 6.11, they are minor. This is not to suggest that order does not matter in every circumstance. In this scenario, the two observers are acquiring different types of data, and thus find different facets of the target trajectories more informative. Due to these differences, one schedule should not significantly impact the other. When performing the optimizations on observers with similar measurement types, the results may indicate otherwise, though this question is left to future investigations.

The results of both the single- and multiobserver multitarget studies expand upon the single-target investigation. The forecasted objective maintains the same benefits from a target tracking perspective, specifically with regard to emphasizing the uncertainty at the reference time, while also exhibiting faster runtimes when compared to the myopic objective. Additionally, the coalescence term results in tasking decisions that are less likely to observe coalescence in the measurement space, reducing the risk when performing data association. While optimization of the proposed objective in Eq. (5.19) is costly from a computational standpoint, the cost is mitigated when leveraging the two phase optimization. This two phase optimization allows for submodular optimization strategies to be employed, providing more feasible runtimes for practical application. The different strategies allow

for an optimization solution to be dictated by the constraints of the available computational resources, however the results indicate the significant speed benefits exhibited by the greedy strategies with a relatively low loss of information.

7. CONCLUSIONS

7.1. RESEARCH SUMMARY

The work presented in this dissertation has focused on the problem of sensor tasking under the context of resident space object catalog maintenance. Due to the stochastic nature of the problem of target tracking, Section 2 laid the foundation for a probabilistic treatment of the target states and measurement model while also providing a background in information theory for context in the later developments. In the interest of completeness, Section 3 detailed the specific considerations taken when modeling the dynamics of the problem and the measurement models employed in the work.

Section 4 provided a thorough analysis of a conventional sensor tasking approach in a single-target tracking scenario, deriving the full distribution for the Kullback-Leibler (KL) divergence when used to describe the directed distance between *a priori* and *a posteriori* densities as they are maintained in an extended Kalman filter (EKF). This distribution allowed for a better understanding of the approximations made when the first moment of the KL divergence is employed as a sensor tasking objective. The section continued to then generalize the consideration of the first moment through a forecasted approach, providing a means of emphasizing uncertainty precision at an arbitrary reference time. Following the development of the forecasted divergence, the differences between the forecasted KL measure and its myopic counterpart were examined in a brief tasking scenario, illustrating their connection when considering a single measurement time and the advantages the forecasted measure provides when considering sets of multiple measurements. In the interest of addressing the initialization for a single target in optimizing the tasking objective, the Sundman transformation was presented and applied to provide a coarse sensor schedule

to be refined by an optimization solver. The section concluded with some discussions of the versatility of the measure, specifically its computation in different spaces and in the presence of other estimated parameters.

The sensor tasking problem was then lifted into the multitarget domain in Section 5. The assumption of independent targets was taken, enabling the myopic and forecasted KL objectives in Section 4 to be detailed in a straightforward manner. This was followed by a proof illustrating that the objectives are both submodular set functions, enabling submodular optimization strategies to be applied when generating schedules.

Section 5 went on to discuss the issue of data association to provide more context to the sensor tasking problem in the presence of multiple targets, illustrating that filters may be vulnerable if optimized schedules dictate that observations be taken during “collisions” between targets in the measurement space. One common solution to data association was presented, namely gating the squared Mahalanobis distance for acquired measurements. The section proposed an additional tasking objective that avoids these scenarios through the use of information entropies, specifically the Rényi entropy. A brief scenario was given to illustrate the use of information entropies in characterizing events in which coalescence between target state estimates in the measurement space is high. In order to avoid observing these events and thus providing a filter data in which it may be difficult to discern the origin of the measurements, an objective was constructed that accounted for the expected information gain provided by a measurement set as well as the potential for observing a collision in the measurement space. A detailed investigation into the optimization of this joint objective was given, including a two-phase approach in which the divergence-based objective can be optimized prior to the entropic objective, enabling submodular optimization strategies to be employed. A selection of submodular optimization schemes was presented to conclude the section.

Lastly, the theoretical discussions and developments were investigated through numerical simulations in Section 6. The first set of simulations sought to investigate the behavior and performance of the proposed forecasted measures in the presence of a single target. This was achieved by generating sensor schedules for two ground-based observers, applying brute-force optimization over an observation window consisting of a single pass of a resident space object (RSO) pulled from a catalog. The forecasted KL-based objective was compared to the more traditional myopic approach as well as the Fisher information gain (FIG), another commonly used tasking objective. Two different orbits were evaluated for the RSO, the first being an eccentric orbit that provided a geometrically interesting measurement profile to observe the objectives' behaviors when a clear information-rich portion of the trajectory is available, followed by a circular orbit which lacked this geometric diversity in the measurement profiles. The results indicated that the forecasted measures not only operated as designed, minimizing the volume of uncertainty at the reference time, but also exhibited several other benefits in their performance. The resulting schedules took measurements at information rich points in time while also autonomously spreading the measurements over the rest of the time interval. Additionally, when data taken at the scheduled times were processed in an EKF, the average entropy of the target state estimate over the observation window was lower than the more traditional methods, indicating that guaranteeing a minimal uncertainty volume at a reference time does not necessarily sacrifice performance with respect to other measurement times.

Section 6 concluded with a multitarget simulation in which the focus of the investigation was placed on the two phase optimization. A linear example was given with a space-based observer tracking a dozen targets in the Hill frame. This scenario facilitated the analysis of several different optimization strategies and provided a comparison of the runtimes and the performances of the submodular optimization strategies to their continuous counterparts. The results illustrated the reduction in the time required to achieve optimization convergence as well as an idea of the loss of schedule quality that can be

expected, which was shown to be minimal. The simulation also illustrated the contribution of the entropy collision term in the optimization process, reducing the average coalescence present in the observations acquired when executing the resulting sensor schedules. A final case demonstrated the advantages of the two phase optimization in the presence of multiple observers by introducing a ground-based observer and maintaining the target state estimates in the Earth-centered inertial frame. As expected, coordination of the observers when performing a simultaneous optimization of the two objective measures, as well as the continuous optimization in the two phase approach, proved to be infeasible from a computational standpoint, while the submodular strategies afforded an achievable solution with small losses in performance when independently optimizing each target/observer pair schedule.

Ultimately, the end result was a sensor tasking policy that provides a minimal uncertainty volume solution at an arbitrary time along with an optimization methodology that enables schedule robustness and optimization runtime to be tuned based on the constraints of the application. While the most appropriate selection of an optimization solution is problem specific, the prioritization of information content exhibited by the two phase optimization provides advantages over the simultaneous optimization of the expected KL divergence and target coalescence terms. When considering independent targets, the submodular strategies, namely the greedy-in-observer and greedy-in-target optimization schemes, enable an optimizer to better navigate the objective function by reducing the dimension of the optimization problem, providing better performance from both a runtime and schedule quality perspective.

7.2. FUTURE RESEARCH DIRECTIONS

The work presented in this dissertation explored the solutions to the questions posed in the introduction, though inevitably more questions were raised in the pursuit of the answers. This research investigated the performance of forecasting the expected KL

divergence, but does not limit its application to this sole information divergence measure. Other divergences are of interest (e.g. the Cauchy-Schwarz divergence in the multitarget tasking problem) as other measures may be better-suited to the multitarget problem and may benefit from the approaches from either a tracking performance or runtime perspective. Additionally, the work presented here assumed sufficient *a priori* knowledge of the targets to be maintained, providing relatively small uncertainties and enabling the EKF to be utilized in the presence of nonlinear dynamics. Future work could consider handling nonlinearities in other ways to make the forecasting approach accessible in applications with larger uncertainties. Considering different divergences or different filters may result in computationally expensive optimizations, but the presented optimization strategies may alleviate some of the computational demand if the objective remains submodular.

Further, the two phase objective function and the submodular optimization strategies make other investigations more accessible and enable other questions to be explored. For instance, in the multiobserver scenario presented in Section 6, the order in which the observer schedules were generated had little impact on the overall target tracking performance. Are there scenarios in which the problem geometry or the differences in the available sensor resources increase the effect of the order of a greedy-in-observer optimization? The faster optimization runtimes afforded by the presented schemes provide a more feasible approach to this and similar explorations.

APPENDIX A.

THE SHANNON ENTROPY FOR A POISSON RFS

If the multitarget RFS is assumed to be an i.i.d. cluster process, Eq. (2.7) is substituted into Eq. (2.14) and the set integral definition in Eq. (2.8) applied to yield the Shannon entropy for an i.i.d. cluster RFS, which is given by

$$H[f] = - \sum_{n=0}^{\infty} \frac{1}{n!} \int_{\mathcal{X}^n} \left[n! \rho(n) \prod_{k=1}^n s(\mathbf{x}^{(k)}) \right] \log \left(n! \rho(n) \prod_{i=1}^n s(\mathbf{x}^{(i)}) \right) d\mathbf{x}^{(1)} \dots d\mathbf{x}^{(n)},$$

where \mathcal{X}^n is the Cartesian product of n copies of the state space. Fortunately, the product rule of logarithms enables the decomposition of the logarithm into terms dependent upon the single-target spatial densities and terms dependent on the cardinality n . Thus, the entropy may be expressed as

$$H[f] = - \sum_{n=0}^{\infty} \left\{ \rho(n) \int_{\mathcal{X}^n} \left[\prod_{k=1}^n s(\mathbf{x}^{(k)}) \right] \log \{n! \rho(n)\} d\mathbf{x}^{(1)} \dots d\mathbf{x}^{(n)} \right\} \\ - \sum_{n=0}^{\infty} \left\{ \rho(n) \int_{\mathcal{X}^n} \left[\prod_{k=1}^n s(\mathbf{x}^{(k)}) \right] \sum_{i=1}^n \log \{s(\mathbf{x}^{(i)})\} d\mathbf{x}^{(1)} \dots d\mathbf{x}^{(n)} \right\}.$$

The product of the single-target spatial densities applies a sifting-like effect; since each logarithm term is only dependent upon $\mathbf{x}^{(i)}$, the remaining $n - 1$ integrals can be evaluated over each spatial density. The result of each of the $n - 1$ evaluations is unity, as the single-target spatial densities are taken to be valid pdfs. Therefore, only a sum of integrals over the target state space, \mathcal{X} , remains, and the result is that the entropy is given by

$$H[f] = - \sum_{n=0}^{\infty} \left[\rho(n) \log \{n! \rho(n)\} \right] - \sum_{n=0}^{\infty} \left[\rho(n) \sum_{i=1}^n \int_{\mathcal{X}} s(\mathbf{x}^{(i)}) \log \{s(\mathbf{x}^{(i)})\} d\mathbf{x}^{(i)} \right].$$

While the variable $\mathbf{x}^{(i)}$ has been maintained up to this point to distinguish between the multiple integration dimensions, it is recognized that the i.i.d. assumption can now be used to simplify the expression by replacing $\mathbf{x}^{(i)}$ with a non-indexed \mathbf{x} since each n -tuple integral has been reduced to a single integral over the state space \mathcal{X} . The result of dropping

the index is a sum of n identical integrals, or

$$H[f] = - \sum_{n=0}^{\infty} \left[\rho(n) \log \{n! \rho(n)\} \right] - \sum_{n=0}^{\infty} [n \rho(n)] \int_{\mathcal{X}} s(\mathbf{x}) \log \{s(\mathbf{x})\} d\mathbf{x}.$$

The summation in the second term is simply the definition of the mean of the cardinality distribution. Denoting this mean by μ , it follows that

$$H[f] = - \sum_{n=0}^{\infty} \left[\rho(n) \log \{n! \rho(n)\} \right] - \mu \int_{\mathcal{X}} s(\mathbf{x}) \log \{s(\mathbf{x})\} d\mathbf{x}. \quad (\text{A.1})$$

Equation (A.1) is the Shannon entropy for an RFS under the assumption that it is distributed according to an i.i.d. cluster process.

The expression for the Shannon entropy of an i.i.d. cluster process given in Eq. (A.1) can be specialized by making an assumption on the cardinality distribution. Substituting the Poisson cardinality distribution given by Eq. (2.5) into the entropy relationship of Eq. (A.1), while noting that $\mu = \lambda$, a closed-form solution for the infinite summation can be obtained, as

$$\begin{aligned} - \sum_{n=0}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \log \{\lambda^n e^{-\lambda}\} &= -e^{-\lambda} \sum_{n=0}^{\infty} \frac{\lambda^n}{(n-1)!} (\log \lambda - \lambda) \\ &= -\lambda e^{-\lambda} (\log \lambda - \lambda) \sum_{n=0}^{\infty} \frac{n \lambda^n}{n!}. \end{aligned} \quad (\text{A.2})$$

Note that the summation can be reindexed to begin at $n = 1$ (as the term for $n = 0$ has no contribution), and the infinite summation converges to e^λ . Leveraging this in Eq. (A.2) and substituting into Eq. (A.1) yields

$$H[f] = \lambda - \lambda \log \lambda - \lambda \int_{\mathcal{X}} s(\mathbf{x}) \log \{s(\mathbf{x})\} d\mathbf{x}. \quad (\text{A.3})$$

Solving for the spatial density in terms of the intensity function from Eq. (2.10) and substituting the result into Eq. (A.3), it follows that the Shannon entropy may be expressed solely in terms of the rate parameter, λ , and the intensity function, $v(\mathbf{x})$, as

$$H[f] = \lambda - \lambda \log \lambda - \int_{\mathcal{X}} v(\mathbf{x}) \log \{v(\mathbf{x})\} d\mathbf{x} + \int_{\mathcal{X}} v(\mathbf{x}) d\mathbf{x} \log \lambda .$$

The integral in the final term is replaced by recalling from Eq. (2.9) that it is simply the rate parameter, λ ; therefore,

$$H[f] = \lambda - \int_{\mathcal{X}} v(\mathbf{x}) \log \{v(\mathbf{x})\} d\mathbf{x} .$$

APPENDIX B.

THE RÉNYI ENTROPY FOR A POISSON RFS

The multitarget Rényi divergence is defined to be

$$H^{(\alpha)}[p] = \frac{1}{1-\alpha} \log \left\{ \int_{\mathcal{X}^n} p^\alpha(\mathbf{X}) \delta \mathbf{X} \right\}.$$

When \mathbf{X} is assumed to be a Poisson RFS, the multitarget pdf in Eq. (2.4) and the set integral definition in Eq. (2.8) are applied to yield

$$H^{(\alpha)}[f] = \frac{1}{1-\alpha} \log \left\{ \sum_{n=0}^{\infty} \frac{1}{n!} \int_{\mathcal{X}^n} [n! \rho(n) s(\mathbf{x}^{(1)}) \cdots s(\mathbf{x}^{(n)})]^\alpha d\mathbf{x}^{(1)} \cdots d\mathbf{x}^{(n)} \right\}.$$

The indices on the $\mathbf{x}^{(i)}$ distinguish between the individual target integrations. However, due to the independence of the individual targets, these integrations can be separated and the indices can be dropped and replaced with a product of n integrals as

$$\begin{aligned} H^{(\alpha)}[f] &= \frac{1}{1-\alpha} \log \left\{ \sum_{n=0}^{\infty} \frac{1}{n!} (n!)^\alpha \rho^\alpha(n) \int_{\mathcal{X}} s^\alpha(\mathbf{x}^{(1)}) d\mathbf{x}^{(1)} \cdots \int_{\mathcal{X}} s^\alpha(\mathbf{x}^{(n)}) d\mathbf{x}^{(n)} \right\} \\ &= \frac{1}{1-\alpha} \log \left\{ \sum_{n=0}^{\infty} \frac{1}{n!} (n!)^\alpha \rho^\alpha(n) \left[\int_{\mathcal{X}} s^\alpha(\mathbf{x}) d\mathbf{x} \right]^n \right\}. \end{aligned} \quad (\text{B.1})$$

Substituting Eq. (2.5) for the cardinality distribution in Eq. (B.1) and reducing yields

$$H^{(\alpha)}[f] = \frac{1}{1-\alpha} \log \left\{ e^{-\alpha\lambda} \sum_{n=0}^{\infty} \frac{1}{n!} \left[\int_{\mathcal{X}} (\lambda s(\mathbf{x}))^\alpha d\mathbf{x} \right]^n \right\}. \quad (\text{B.2})$$

Note that the terms under the summation are in the form of an exponential, namely

$$e^z = \sum_{n=0}^{\infty} \frac{z^n}{n!}.$$

Leveraging this fact, the summation can be replaced with its closed-form solution, yielding

$$H^{(\alpha)}[f] = \frac{1}{1-\alpha} \log \left\{ e^{-\alpha\lambda + \int_{\mathcal{X}} (\lambda s(\mathbf{x}))^\alpha d\mathbf{x}} \right\}.$$

As stated previously, this work utilizes the natural logarithm when computing information theoretic measures. Consequently, the Poisson RFS Rényi entropy becomes

$$H^{(\alpha)}[f] = \frac{1}{1-\alpha} \left[-\alpha\lambda + \int_{\mathcal{X}} (\lambda s(\mathbf{x}))^\alpha d\mathbf{x} \right].$$

From the definition of the intensity function for a Poisson RFS in Eq. (2.10), the Rényi entropy is given in terms of the rate parameter, λ , and the intensity function, $v(\mathbf{x})$, to be

$$H^{(\alpha)}[f] = -\frac{\alpha\lambda}{1-\alpha} + \frac{1}{1-\alpha} \int_{\mathcal{X}} v^\alpha(\mathbf{x}) d\mathbf{x}.$$

REFERENCES

- [1] Adurthi, N., *Conjugate Unscented Transform Based Methods for Uncertainty Quantification, Nonlinear Filtering, Optimal Feedback Control and Dynamic Sensing*, Ph.D. thesis, State University of New York at Buffalo, 2016.
- [2] Adurthi, N., Singla, P., and Majji, M., ‘Conjugate unscented transform based approach for dynamic sensor tasking and space situational awareness,’ in ‘2015 American Control Conference (ACC),’ IEEE, 2015 pp. 5218–5223.
- [3] Adurthi, N., Singla, P., and Majji, M., ‘Mutual information based sensor tasking with applications to space situational awareness,’ *Journal of Guidance, Control, and Dynamics*, 2020, **43**(4), pp. 767–789, doi:10.2514/1.G004399.
- [4] Alspach, D. L. and Sorenson, H. W., ‘Nonlinear Bayesian estimation using Gaussian sum approximations,’ *IEEE Transactions on Automatic Control*, 1972, **17**(4), pp. 439–448.
- [5] Amari, S.-I., ‘ α -divergence is unique, belonging to both f -divergence and Bregman divergence classes,’ *IEEE Transactions on Information Theory*, 2009, **55**(11), pp. 4925–4931.
- [6] Annex, A., ‘SpiceyPy, a Python Wrapper for SPICE,’ in ‘Third Planetary Data Workshop and The Planetary Geologic Mappers Annual Meeting,’ volume 1986, 2017 p. 7081.
- [7] Baldessari, B., ‘The distribution of a quadratic form of normal random variables,’ *The Annals of Mathematical Statistics*, 1967, **38**(6), pp. 1700–1704.
- [8] Bar-Shalom, Y. and Fortmann, T., *Tracking and Data Association*, Mathematics in Science and Engineering Series, Academic Press, 1988, ISBN 9780120797608.
- [9] Battin, R. H., *An Introduction to the Mathematics and Methods of Astrodynamics, Revised Edition*, American Institute of Aeronautics and Astronautics, 1999, ISBN 978-1-56347-342-5, doi:10.2514/4.861543.
- [10] Blackman, S. S., ‘Multiple hypothesis tracking for multiple target tracking,’ *IEEE Aerospace and Electronic Systems Magazine*, 2004, **19**(1), pp. 5–18.
- [11] Bodenham, D. A. and Adams, N. M., ‘A comparison of efficient approximations for a weighted sum of chi-squared random variables,’ *Statistics and Computing*, 2016, **26**(4), pp. 917–928, ISSN 1573-1375, doi:10.1007/s11222-015-9583-4.
- [12] Carpenter, J. R. and D’Souza, C. N., ‘Navigation filter best practices,’ 2018.

- [13] Castaño-Martínez, A. and Fernando, L.-B., ‘Distribution of a sum of weighted central chi-square variables,’ *Comm. Statist. Theory Methods*, 2005, **34**, pp. 515–524, doi: 10.1081/STA-200052148.
- [14] Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C., *Introduction to Algorithms*, MIT press Cambridge, MA, second edition, 2003, ISBN 0-262-03293-7.
- [15] Csiszár, I., *Stochastics: Information Theory*, Springer, 2006.
- [16] Dames, P. M., *Multi-Robot Active Information Gathering using Random Finite Sets*, Ph.D. thesis, University of Pennsylvania, Philadelphia, Pennsylvania, 2015.
- [17] Darling, J. E. and DeMars, K. J., ‘Minimization of the Kullback-Leibler divergence for nonlinear estimation,’ *Journal of Guidance, Control, and Dynamics*, 2017, **40**(7), pp. 1739–1748.
- [18] Davies, M. E., Abalakin, V. K., Brahic, A., Bursa, M., Chovitz, B., Lieske, J., Seidelmann, P., Sinclair, A., and Tiufin, I., ‘Report of the iau/iag/cospar working group on cartographic coordinates and rotational elements of the planets and satellites-1991,’ *Celestial Mechanics and Dynamical Astronomy*, 1992, **53**, pp. 377–397.
- [19] DeMars, K. J. and Gualdoni, M. J., ‘Information theoretic approaches to space object collision,’ in ‘Space Debris, 7th European Conference on,’ ESA, 2017 .
- [20] DeMars, K. J. and Jah, M. K., ‘Evaluation of the information content of observations with application to sensor management for orbit determination,’ *Advances in the Astronautical Sciences*, 2011, **142**, pp. 3187–3206.
- [21] Der, G. J., *An elegant state transition matrix*, 1997, doi:10.2514/6.1996-3660.
- [22] Dughmi, S., ‘Submodular functions: Extensions, distributions, and algorithms. a survey,’ 2009.
- [23] Friedman, A. M. and Frueh, C., ‘Determining characteristics of artificial near-Earth objects using observability analysis,’ *Acta Astronautica*, 2018, **144**, pp. 405–421.
- [24] Frueh, C., Fiedler, H., and Herzog, J., ‘Realistic sensor tasking strategies,’ in ‘Advanced Maui Optical and Space Surveillance Technologies Conference, 2016, Hawaii, USA,’ 2016 .
- [25] Frueh, C., Fielder, H., and Herzog, J., ‘Heuristic and optimized sensor tasking observation strategies with exemplification for geosynchronous objects,’ *Journal of Guidance, Control, and Dynamics*, 2017, **41**(5), pp. 1036–1048.
- [26] Gabler, S. and Wolff, C., ‘A quick and easy approximation to the distribution of a sum of weighted chi-square variables,’ *Statistische Hefte*, 1987, **28**(1), pp. 317–325, ISSN 1613-9798, doi:10.1007/BF02932611.

- [27] Gehly, S., Jones, B., and Axelrad, P., 'Sensor allocation for tracking geosynchronous space objects,' *Journal of Guidance, Control, and Dynamics*, 2018, **41**(1), pp. 149–163, doi:10.2514/1.G000421.
- [28] Gim, D.-W. and Alfriend, K. T., 'State transition matrix of relative motion for the perturbed noncircular reference orbit,' *Journal of Guidance, Control, and Dynamics*, 2003, **26**(6), pp. 956–971.
- [29] Glenn Shepherd, A., 'Space surveillance network,' Shared Space Situational Awareness Conference, Colorado Springs, CO, 2006.
- [30] Goodyear, W. H., 'Completely general closed-form solution for coordinates and partial derivative of the two-body problem,' *The Astronomical Journal*, 1965, **70**, p. 189.
- [31] Graybill, F. A. and Marsaglia, G., 'Idempotent matrices and quadratic forms in the general linear hypothesis,' *The Annals of Mathematical Statistics*, 1957, **28**(3), pp. 678–686, ISSN 00034851.
- [32] Gu, D. and Hu, H., 'Rényi entropy based target tracking in mobile sensor networks,' *IFAC Proceedings Volumes*, 2011, **44**(1), pp. 13558 – 13563, ISSN 1474-6670, doi: <https://doi.org/10.3182/20110828-6-IT-1002.00602>, 18th IFAC World Congress.
- [33] Gualdoni, M. J. and DeMars, K. J., 'An improved representation of measurement information content via the distribution of the Kullback-Leibler divergence,' *Advances in the Astronautical Sciences. Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, 2017, **162**(17-810).
- [34] Gualdoni, M. J. and DeMars, K. J., 'Evaluating the performance of information-based multi-observation sensor tasking,' in 'John L. Junkins Dynamical Systems Symposium,' 2018 .
- [35] Gualdoni, M. J. and DeMars, K. J., 'The information content of data arcs for multi-step sensor tasking,' *Proceedings of 2018 Space Flight Mechanics Meeting*, 2018, (2018-0727), doi:10.2514/6.2018-0727.
- [36] Gualdoni, M. J. and DeMars, K. J., 'An optimization formulation of information theoretic sensor tasking,' *Advances in the Astronautical Sciences. Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, 2018, **163**(18-471).
- [37] Hero, A. O., Kreucher, C. M., and Blatt, D., 'Information theoretic approaches to sensor management,' in 'Foundations and applications of sensor management,' pp. 33–57, Springer, 2008.
- [38] Hill, K., Sydney, P., Cortez, R., Hamada, K., Nishimoto, D., Luu, K., and Schumacher Jr, P. W., 'Dynamic tasking of networked sensors using covariance information,' Technical report, Pacific Defense Solutions, LLC, 2010.
- [39] Hintz, K. J. and McVey, E. S., 'Multi-process constrained estimation,' *IEEE Transactions on Systems, Man, and Cybernetics*, 1991, **21**(1), pp. 237–244.

- [40] Hoffman, K. and Kunze, R., *Linear Algebra*, Prentice-Hall, second edition, 1971.
- [41] Inalhan, G., Tillerson, M., and How, J. P., 'Relative dynamics and control of spacecraft formations in eccentric orbits,' *Journal of Guidance, Control, and Dynamics*, 2002, **25**(1), pp. 48–59.
- [42] Jazwinski, A. H., *Stochastic Processes and Filtering Theory*, Courier Corporation, 2007.
- [43] Kalman, R. E., 'A New Approach to Linear Filtering and Prediction Problems,' *Journal of Basic Engineering*, 1960, **82**(1), pp. 35–45, ISSN 0021-9223, doi: 10.1115/1.3662552.
- [44] Kangsheng, T. and Guangxi, Z., 'Sensor management based on Fisher information gain,' *Journal of Systems Engineering and Electronics*, 2006, **17**(3), pp. 531–534, ISSN 1004-4132, doi:10.1016/S1004-4132(06)60091-1.
- [45] Kelso, T., 'Iridium 33/Cosmos 2251 collision,' *Celestrak.com*, 2012, **5**.
- [46] Kelso, T., Parkhomenko, N., Shargorodsky, V., Vasiliev, V., Yurasov, V., Nazarenko, A., Tanygin, S., and Hiles, R., 'What happened to BLITS? an analysis of the 2013 January 22 event,' in 'Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference,' volume 4, 2013 .
- [47] Krause, A. and Guestrin, C., 'Near-optimal observation selection using submodular functions,' in 'AAAI,' volume 7, 2007 pp. 1650–1654.
- [48] Krause, A. and Guestrin, C., 'Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach,' in 'Proceedings of the 24th international conference on Machine learning,' ACM, 2007 pp. 449–456.
- [49] Kreucher, C., Kastella, K., and Hero, A. O., 'Multi-target sensor management using alpha-divergence measures,' in 'Information Processing in Sensor Networks,' volume 2634, Springer, 2003 pp. 209–222.
- [50] Kreucher, C. M., Kastella, K., and Hero III, A. O., 'Information based sensor management for multitarget tracking,' in 'Proceeds of SPIE,' volume 5204, 2003 pp. 480–489.
- [51] Kullback, S., *Information Theory and Statistics*, Courier Corporation, 1997.
- [52] Kullback, S. and Leibler, R. A., 'On information and sufficiency,' *The Annals of Mathematical Statistics*, 1951, **22**(1), pp. 79–86.
- [53] LeGrand, K. A., *Space-Based Relative Multitarget Tracking*, Master's thesis, Missouri University of Science and Technology, 2015.
- [54] Little, B. D. and Frueh, C. E., 'Space situational awareness sensor tasking: Comparison of machine learning with classical optimization methods,' *Journal of Guidance, Control, and Dynamics*, 2020, **43**(2), pp. 262–273.

- [55] Lovász, L., *Submodular functions and convexity*, pp. 235–257, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN 978-3-642-68874-4, 1983, doi:10.1007/978-3-642-68874-4_10.
- [56] Mahalanobis, P., ‘On the generalised distance in statistics,’ in ‘Proceedings of the National Institute of Sciences of India,’ volume 2, 1936 pp. 49–55.
- [57] Mahler, R. P. S., *Statistical Multisource-Multitarget Information Fusion*, Artech House, Inc., 2007.
- [58] Mahler, R. P. S., *Advances in Statistical Multisource-Multitarget Information Fusion*, Artech House, Inc., 2014.
- [59] Marcus, M. and Minc, H., *A survey of matrix theory and matrix inequalities*, volume f14, Courier Corporation, 1992.
- [60] McCray, W. P., *Keep Watching the Skies!: the Story of Operation Moonwatch & the Dawn of the Space Age*, Princeton University Press, 2008.
- [61] Melsa, J. L. and Sage, A. P., *An Introduction to Probability and Stochastic Processes*, Courier Corporation, 2013.
- [62] Moran, W., Suvorova, S., and Howard, S., ‘Application of sensor scheduling concepts to radar,’ in ‘Foundations and Applications of Sensor Management,’ pp. 221–256, Springer, 2008.
- [63] Moschopoulos, P. and Canada, W., ‘The distribution function of a linear combination of chi-squares,’ *Computers & mathematics with applications*, 1984, **10**(4-5), pp. 383–386, doi:10.1016/0898-1221(84)90066-X.
- [64] Müller-Lennert, M., Dupuis, F., Szehr, O., Fehr, S., and Tomamichel, M., ‘On quantum Rényi entropies: a new generalization and some properties,’ *Journal of Mathematical Physics*, 2013, **54**(12), p. 122203, doi:10.1063/1.4838856.
- [65] Murty, K. G. and Kabadi, S. N., ‘Some NP-complete problems in quadratic and nonlinear programming,’ *Mathematical Programming*, 1987, **39**(2), pp. 117–129, ISSN 1436-4646, doi:10.1007/BF02592948.
- [66] Nader, R. and Kelso, T., ‘The pegasus incident: The loss of the first ecuadorian satellite and its recovery,’ in ‘65th International Astronautical Congress 2014 SPACE,’ 2014 .
- [67] Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L., ‘An analysis of approximations for maximizing submodular set functions - I,’ *Mathematical programming*, 1978, **14**(1), pp. 265–294.
- [68] Nesterov, Y. and Nemirovskii, A., *Interior-point polynomial algorithms in convex programming*, volume 13, Siam, 1994.

- [69] Rényi, A., ‘On measures of entropy and information,’ in ‘Proceedings of the fourth Berkeley symposium on mathematical statistics and probability,’ volume 1, 1961 pp. 547–561.
- [70] Rényi, A., *Foundations of Probability*, Courier Corporation, 2007.
- [71] Rezaeian, M. and Vo, B.-N., ‘The entropy of random finite sets,’ in ‘Proceedings of the International Symposium on Information Theory,’ 2009 .
- [72] Roa, J., *Regularization in Orbital Mechanics: Theory and Practice*, volume 42, Walter de Gruyter GmbH & Co KG, 2017.
- [73] Robinson, S., Scarritt, S., and Goodman, J. L., ‘Encke-beta predictor for orion burn targeting and guidance,’ Proceedings of the 39th Annual AAS Guidance and Control Conference, 2016, (16-112).
- [74] Schildknecht, T., ‘Optical surveys for space debris,’ *The Astronomy and Astrophysics Review*, 2007, **14**(1), pp. 41–111.
- [75] Shannon, C. E., Weaver, W., and Burks, A. W., *The Mathematical Theory of Communication*, 1951.
- [76] Singull, M. and Koski, T., ‘On the distribution of matrix quadratic forms,’ *Communications in Statistics - Theory and Methods*, 2012, **41**(18), pp. 3403–3415, doi: 10.1080/03610926.2011.563009.
- [77] Solomon, H. and Stephens, M. A., ‘Distribution of a sum of weighted chi-square variables,’ *Journal of the American Statistical Association*, 1977, **72**(360a), pp. 881–885, doi:10.1080/01621459.1977.10479976.
- [78] Sorenson, H. W. and Alspach, D. L., ‘Recursive Bayesian estimation using Gaussian sums,’ *Automatica*, 1971, **7**(4), pp. 465–479.
- [79] Stringer, M. T., Newman, B., Lovell, T. A., and Omran, A., ‘Second order nonlinear initial value solution for relative motion using Volterra theory,’ in ‘Proceedings of the 23rd AAS/AIAA Space Flight Mechanics Conference,’ AAS 13-469, 2013 pp. 2133–2149.
- [80] Sturdevant, R. W., ‘From satellite tracking to space situational awareness: the USAF and space surveillance, 1957-2007,’ *Air Power History*, 2008, **55**(4), pp. 4–23.
- [81] Tapley, B. D., Schutz, B. E., and Born, G. H., *Statistical Orbit Determination*, Elsevier, 2004.
- [82] Tuggle, K. and Akella, M., ‘Information-theoretic sensor scheduling under communication constraints,’ *Advances in the Astronautical Sciences. Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, 2018, **163**(18-371).

- [83] Vincent, M. A., 'How state error covariance matrices evolve in six dimensions,' in 'Proc. Of the Advanced Maui Optical and Space Surveillance Technologies Conference,' 2018 .
- [84] Vo, B.-N. and Ma, W.-K., 'The Gaussian mixture probability hypothesis density filter,' *IEEE Transactions on Signal Processing*, 2006, **54**(11), pp. 4091–4104, doi: 10.1109/TSP.2006.881190.
- [85] Walpole, R. E., Myers, R. H., Myers, S. L., and Ye, K., *Probability and Statistics for Engineers and Scientists*, Pearson Education, 2nd edition, 1978, ISBN 0-02-424110-5.
- [86] Weeden, B., 'Iridium-Cosmos collision fact sheet,' Space World Foundation, November, 2009, **10**, p. 2010.
- [87] Weeden, B., 'Space situational awareness fact sheet,' Secure World Foundation, 2014.
- [88] Wiesel, W. E., *Spaceflight dynamics*, McGraw-Hill Science Engineering, New York, NY, 1997.
- [89] Williams, P. S., Spencer, D. B., and Erwin, R. S., 'Coupling of estimation and sensor tasking applied to satellite tracking,' *Journal of Guidance, Control, and Dynamics*, 2013, **36**(4), pp. 993–1007.
- [90] Zamir, R. and Feder, M., 'A generalization of the entropy power inequality with applications,' *IEEE Transactions on Information Theory*, 1993, **39**(5), pp. 1723–1728, ISSN 1557-9654, doi:10.1109/18.259666.

VITA

Matthew James Gualdoni was born in St. Charles, MO and raised by his parents, Janice and James, alongside his older brother, Michael. Following his graduation from Saint Charles High School in 2010, Matt enrolled in Missouri University of Science and Technology (Missouri S&T) to pursue degrees in both computer and electrical engineering, eventually graduating summa cum laude in December of 2014, with further distinguished honors in electrical engineering.

Following graduation, Matt remained at Missouri S&T, enrolling in the aerospace engineering graduate program under his advisor, Dr. Kyle J. DeMars, and taking on a position as a graduate research assistant in the Algorithms Research for the Estimation of Uncertain Systems (AREUS) laboratory. He received a Graduate Assistance in Areas of National Need (GAANN) fellowship in 2017, enabling him to continue his research efforts and pursue his Ph.D. During his doctoral studies, Matt worked several internships at NASA's Johnson Space Center (JSC), where he was accepted into the NASA Pathways Coop Program in 2018. In August 2020, he received his Ph.D. degree in Aerospace Engineering from Missouri University of Science and Technology. Following the completion of his Ph.D., Matt began his full-time employment at NASA JSC as a member of the GNC Autonomous Flight Systems Branch.