ABSTRACT

Title of dissertation: SATELLITE CLUSTER TRACKING VIA EXTENT ESTIMATION

Christopher R. Binz
Doctor of Philosophy, 2019

Dissertation directed by: Dr. Liam Healy
Department of Aerospace Engineering

Clusters of closely-spaced objects in orbit present unique tracking and prediction challenges. Association of observations to individual objects is often not possible until the objects have drifted sufficiently far apart from one another. This dissertation proposes a new paradigm for initial tracking of these clusters of objects: instead of tracking the objects independently, the cluster is tracked as a single entity, parameterized by its centroid and extent, or shape. The feasibility of this method is explored using a decoupled centroid and extent estimation scheme. The dynamics of the centroid of a cluster of satellites are studied, and a set of modified equinoctial elements is shown to minimize the discrepancy between the motion of the centroid and the observation-space centroid. The extent estimator is formulated as a matrix-variate particle filter. Several matrix similarity measures are tested as the filter weighting function, and the Bhattacharyya distance is shown to outperform the others in test cases. Finally, the combined centroid and extent filter is tested on a set of three on-orbit breakup events, generated using the NASA standard breakup

model and simulated using realistic force models. The filter is shown to perform well across low-Earth, geosynchronous, and highly-elliptical orbits, with centroid error generally below five kilometers and well-fitting extent estimates. These results demonstrate that a decoupled centroid and extent filter can effectively track clusters of closely-spaced satellites. This could improve spaceflight safety by providing quantitative tracking information for the entire cluster much earlier than would otherwise be available through typical means.

SATELLITE CLUSTER TRACKING VIA EXTENT ESTIMATION

by

Christopher R. Binz

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2019

Advisory Committee:
Dr. Liam Healy, Advisor
Professor Raymond Sedwick, Chair
Professor David Akin
Professor Derek Paley
Professor Christine Hartzell
Professor David Lovell

Acknowledgments

This dissertation would not have been written if not for the professional, academic, and personal support I received from a number of people. I gratefully acknowledge the guidance provided by my advisor, Liam Healy, throughout this entire process. Early conversations with Alan Segerman, Paul Schumacher, Chris Sabol, and others helped shape my understanding of the space surveillance field, as well as the direction of this research, and I am grateful to have been able to work with them early in my career.

My participation in the Ph.D. program was generously supported by my employer, the U.S. Naval Research Laboratory, and my management and coworkers provided much-needed encouragement, feedback, and advice throughout. I'd like to specifically thank Alan Segerman, Erin Fowler, Zach Sibert, and Greg Scott for helping to keep me focused and sane.

Finally, I want to thank my family and friends, who created the environment that allowed me to spend the time necessary to finish, while maintaining a balanced, fulfilling life. My parents Valerie and Richard; my sister Lisa; my friends Kyle, Suzy, Mike, and Claire; my wife's family; and many others (including several of the people already mentioned above) deserve thanks for fostering this environment. Last, and certainly not least, my wife Laura deserves more appreciation than I can put into words for remaining confident in me throughout, while doing a wonderful job helping to raise our daughter Cecily. I could not have finished this without her.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

Space surveillance–the tracking of objects in orbit around the Earth, usually in a non-cooperative way–is as old as spaceflight itself. With the launch of *Sputnik-1* in October 1957, Western nations urgently recorded optical measurements through telescopes and measurements of the Doppler shift of the signal from its radio beacon in efforts to track the spacecraft [1][2, Appendix A]. Later, developments in electronics and processing technology made it possible to track satellites using radar [3]. In the decades that followed, networks of sensors such as the U.S. Space Surveillance Network (the largest such system) used radar and optical measurements to update the orbits of all observable, man-made objects in orbit around the Earth. With some exceptions, the process for doing so has remained largely unchanged since its inception. However, several factors are converging to form looming challenges.

First, sensor technology is improving–ground-based sensors are increasing attainable measurement precision, decreasing the minimum detectable size, and expanding the volume of space they can survey at once. The Space Fence [4] and the Space Surveillance Telescope (SST) [5] are two such sensors that will become operational in the near future, representing a significant advance in radar and optical satellite tracking technology, respectively. As a result, existing objects that were

previously below the detectability threshold will become observable, increasing the total number of objects that must be tracked.

The second factor is the growth of the satellite population. Increased access to and use of space by many nations and, more recently, non-governmental organizations, corporations, and even individuals, has already overturned the status quo of space being the domain of only world superpowers. As the satellite population ages, on-orbit breakups are becoming more common as well, adding significantly to the number of objects that must be tracked. The rise of the CubeSat form factor has enabled launches that carry and deploy dozens of separate payloads at a time [6], [7]. Missions that themselves deploy hundreds of tiny "chipsats" have already been developed [8]. Additionally, several commercial companies have presented plans for what have been termed "megaconstellations"—systems consisting of hundreds to thousands of satellites in low-Earth orbit [9]. Overall, the number of objects that must be tracked, as well as the number of missions that must be protected, is growing at a very rapid pace, making space surveillance much harder and more critical than ever. Figure 1.1 shows the dramatic growth in the satellite population since the launch of *Sputnik-1*.

Historically, on-orbit breakups have been rare, and objects were in such disparate orbits that breakups rarely put other operational satellites in immediate danger. This, combined with the inherent accuracy limitations of the standard general perturbations orbit model (SGP4 and its predecessors [10]), meant that breakup processing largely consisted of waiting for the objects to naturally drift apart from one another, then individually cataloging the pieces large enough to be consistently

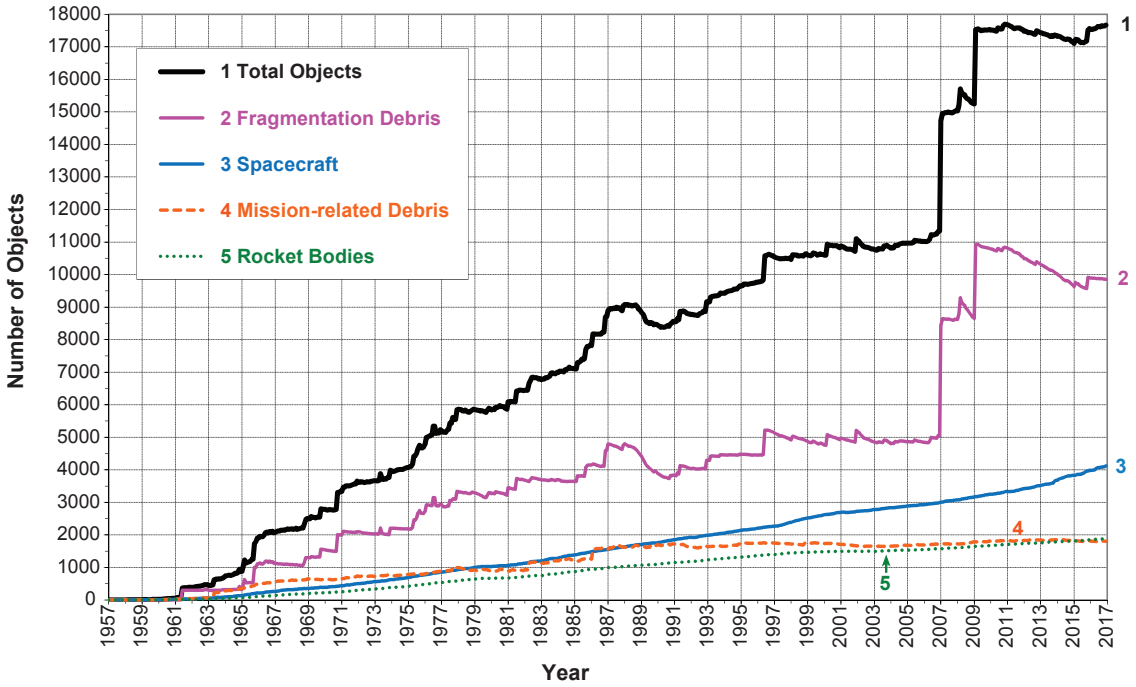

Figure 1.1: Growth in the number of satellites over time. Reproduced from the *Orbital Debris Quarterly News*, volume 21, issue 1, February 2017. `https://www.orbitaldebris.jsc.nasa.gov/quarterly-news/newsletter.html`

tracked once the identity was unambiguous.

In fact, the current nature of the space surveillance enterprise makes processing on-orbit debris events a time-consuming and difficult process–sensor resolution is limited, sensor locations (and thus measurements) are disparate, and sensor and processing resources, assigned the task of maintaining and updating the catalog, are constrained. This is evident when looking at the timelines for publicly-released two-line element sets (TLEs) associated with major debris events.

Figure 1.2 shows the number of objects added to the public catalog in the year following the January 2007 Chinese antisatellite test [11], based on the epoch of the first available TLE for a given object. The first objects were cataloged about seven days after the event. Batches of objects are then added regularly; however, by the

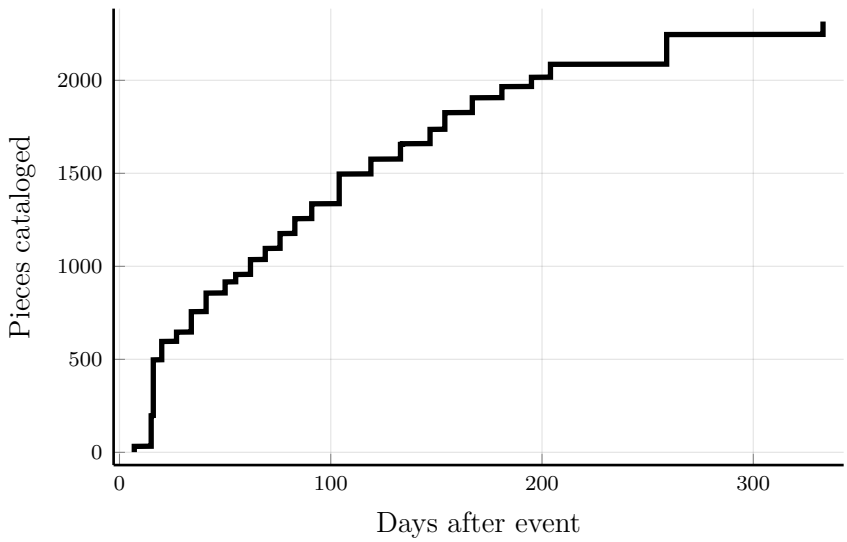

Figure 1.2: Number of TLEs published for the 2007 Chinese antisatellite test.

end of the first year, only about two thirds of the total debris population had been cataloged. The number of cataloged objects for the collision between Iridium 33 and Cosmos 2251 in February 2009 is shown in Figure 1.3, and a similar trend can be seen. In another case, a 2012 explosion of a Russian *Briz-M* upper stage (NORAD catalog #38746) resulted in the SSN "tracking over 500 pieces of debris" [12]. The initial set of 80 two-line element sets (TLEs) was published 8 days after the assessed breakup time, and only 111 debris pieces were ever added to the public catalog [13]. These cataloging delays are not limited to debris events: in recent mass-satellite deployment scenarios, initial tracking and identification of the payloads can take days or even weeks to settle [14, Section: Falcon 9]. From these examples, it is clear that the vast majority of a debris cloud cannot be cataloged in the very early stages after the precipitating event.

The dynamics of a debris cloud in orbit is far from being an understudied topic: there is a significant amount of published literature addressing debris cloud

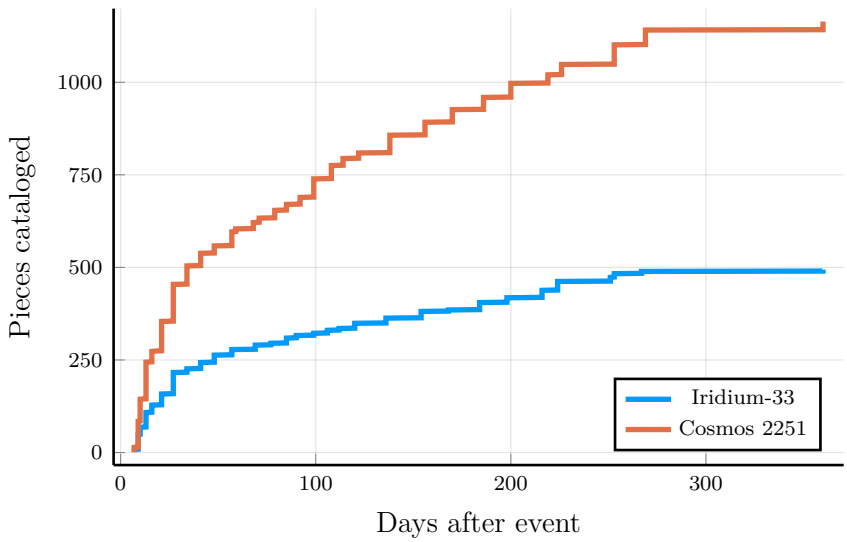

Figure 1.3: Number of TLEs published for the 2009 Iridium-Cosmos collision.

modeling ([15]–[19]; see Chapter 2). While this work is important for understanding the evolution of these clouds of objects, no attempt has been made to couple this with estimation techniques–that is, direct observation and tracking of these clouds as clouds. The technique presented here is an attempt to do just that.

The novel tracking technique presented here is based on group tracking techniques originally applied to the problem of tracking a missile warhead in a cloud of decoys and chaff [20]. Instead of tracking each member of the cluster, or group, individually, the entire cloud is tracked as a single entity, along with the shape, or extent, of the cloud. This alleviates the need for a very dense data set to identify individual objects between frames of measurements, and still allows for an accurate assessment of the hazardous volume associated with such a cluster. By treating each set of observations as a collective, the computational requirements for solving the assignment problem (in which observations are assigned to targets) may be avoided altogether.

This method is intended to be used only when the cloud of objects is relatively dense–that is, when the assignment problem would prove to be excessively difficult. This corresponds to the early phases of a breakup on-orbit, or a mass deployment event–what Jehn refers to as the "pulsating ellipsoid" phase [21]—before natural orbit perturbations cause the objects to drift away from each other, forming a ring (see Figure 2.2 and Section 2.1.3 for more detail on these phases).

Broadly, the group tracking process consists of the following steps. First, the entire set of observations is screened for potential outliers, so that only cluster "members" are included (a process called *gating*). Then, the measured centroid is used to update the centroid of the cluster, and the extent of the cloud of observations is used to update the extent of the cluster in state space. This allows for an accurate description of the hazardous area presented by a debris cloud well before the pieces are separated enough to be cataloged in the traditional way.

This dissertation proposes and evaluates a solution to this problem. The key contributions can be summarized as follows. A characterization method and an estimation method for clusters of closely-spaced objects in orbit are developed. A matrix-variate particle filter is implemented in order to demonstrate the efficacy of this estimation method. Finally, simulations involving realistic fragmentation models and dynamics are performed across a variety of orbit regimes in order to demonstrate the performance of this method.

This dissertation has the following structure. Chapter 2 presents a review of the existing literature in space surveillance, multiple target tracking, and related fields. Chapter 3 introduces the dynamic models used throughout and defines extent

in this context. In Chapter 4, the unique dynamics of the centroid of a tracked cluster are explored, along with the effects of different coordinate systems. Chapter 5 presents the filter formulation, which requires careful selection of the likelihood function. Results for three different test cases are presented in Chapter 6, and conclusions and recommendations for future study are given in Chapter 7.

Chapter 2: Literature Review

This research draws on a body of literature from not only the astrodynamics community, but also the broader target tracking field, particularly the topic of group tracking. This chapter presents a review of the relevant literature from both of these fields, beginning with space surveillance.

2.1 Space Surveillance

Until fairly recently, most space surveillance work used analytical methods (or "general perturbations"[22], [23]) such as SGP4 [10] and PPT3 [24] for orbit prediction. This was largely due to constraints on computational resources–high-precision numerical integration of several thousand objects was simply not feasible. Eventually, computing capabilities advanced to the point where numerical integration (also known as special perturbations (SP) propagation) was feasible for the entire catalog [25], [26].

The orbit determination aspect experienced a shift along with the change in propagation methods. While analytical models were still in use, batch least squares differential correction was used to fit observations to the mean elements used in the model [27]. The shift to SP propagation meant a change to fitting

osculating elements to numerically-integrated models (also using the batch least squares process), resulting in more accurate orbits [26].

Additionally, a significant amount of work has been presented in adapting other estimation methods to the space surveillance problem. In some applications, batch processing of many observations results in an unacceptable time delay, and so online estimation methods are more desirable in such situations. The extended Kalman filter (EKF) linearizes the problem around the current time step, and as such can be applied to any nonlinear estimation problem, often producing similar results to a batch least squares process [28, Section 8.4][29]. The EKF has been shown to be inadequate in some situations, however [30], and the introduction of the unscented Kalman filter [31] was quickly adopted for the space object tracking problem [32]. The UKF eliminates the need for local linearization of the problem, and the associated computation of partial derivative matrices, by simply propagating a set of deterministically-chosen points to represent the target uncertainty distribution. This simplifies the implementation, and provides robustness in situations where an EKF might diverge. The UKF is a central element in many modern space object tracking methods [33]–[35].

Applications beyond single point target tracking have been largely limited to solving the uncorrelated track (UCT) problem (excepting several notable multi-target tracking studies discussed below). When a radar or telescope records a track[1] of observations that does not match the predicted track of a known object in the

[1] A track here is simply a series of closely-spaced observations, usually associated with the same object.

catalog, it is labeled a UCT. Resolving UCTs is largely a manual process, but several methods have been proposed and tested to alleviate this [36]–[38].

Of particular relevance to this dissertation is breakup processing. There is little recent documentation in the literature about operational practices, although two older references describe the process as of the mid-1990s [39], [40]. Once a breakup is confirmed, unassociated observations are first filtered based on orbital plane limits to determine whether they are part of the breakup or not. Then, this (potentially very large) set of unassociated observations must be clustered and fit to candidate orbits. While there are some automated tools to help with this, the process is labor-intensive, which can be slow. Additionally, limitations in the distribution and amount of data often necessitates simply waiting for more observations in order to construct valid orbits.

2.1.1 Uncertainty Quantification

With the maturation of the orbit estimation field, the increasing number of objects in orbit, and increasing accuracy requirements for space surveillance and satellite missions, orbit estimation uncertainty began to attract a great deal of focus in the literature beginning in the late 1990s. The state uncertainty for satellites, if considered at all, has traditionally been treated as a multivariate Gaussian distribution. Methods for propagating the distribution have long relied on linear error theory. Briefly, linear error theory relies on the fact that the first two statistical moments of a distribution (which is a full description of a Gaussian distribution) can be trans-

formed exactly assuming a linear mapping [41].[2] In the case of orbit dynamics, the governing equations tend to be nonlinear in all but the simplest case of Keplerian motion (using the set of equinoctial orbital elements described in Section 3.1), and so the process relies on linearization, a technique that can be limited in its validity.

The seminal paper by Junkins, Akella, and Alfriend [42] is regarded as the canonical introduction to modern nonlinear uncertainty quantification in astrodynamics. The authors compare covariance propagation under linear error theory with Monte Carlo results for rectangular, cylindrical, and orbital element coordinates. They find that linearly-propagated covariance matrices, over time, do a poor job of describing the true uncertainty, due to the nonlinearity of the orbit propagation process. Notably, the authors' choice to use the classical orbital elements "because all but one of them are 'slow-varying' functions of time" [42, p. 541] worked out well: the measure of nonlinearity presented in the paper is in fact lowest for the orbital element coordinate set, and thus the uncertainty in these coordinates is found to most closely match the true distribution of Monte Carlo points. The concept of a nonlinearity index would be expanded on by Junkins and Singla [43] later in a comprehensive tutorial. Figure 2.1 neatly illustrates the issue that has been at the heart of nearly two decades of astrodynamics research.

Sabol *et al.* [44] revisited this problem several years later, while adding some valuable detail. Instead of only investigating the effects of propagation, here the authors begin with the orbit determination process itself. Their first major finding is that "the covariance is representative of the state error distribution at the esti-

[2] See Section 3.3 for more detail.

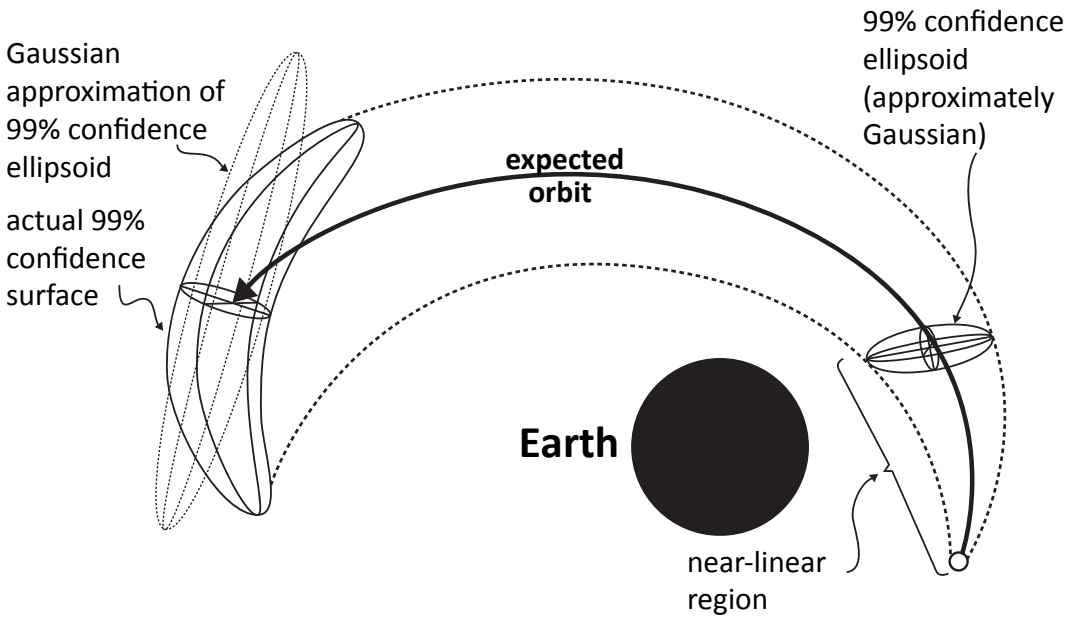

Figure 2.1: Effects of nonlinearity on covariance propagation [42, p. 553]. The "Gaussian approximation" of the confidence surface shown is a Cartesian representation.

mation epoch for all state representations" [44, p. 509]. They go on to look at the effects of propagation in various coordinate systems, again finding that Cartesian coordinates do a poor job capturing the true uncertainty distribution, while slow-varying element set representations do a much better job. The second key finding is that the uncertainty distribution in Cartesian coordinates becomes non-Gaussian largely "due to the limitations of the linear geometric reference frame and not the linearized dynamics used to propagate the covariance" [44, p. 524]. This suggests that linear error theory may be acceptable in many situations, as long as the chosen coordinate system more accurately reflects the physical nature of satellite orbits. These results heavily influence the choice of coordinate system for the work in this dissertation, as described in Chapters 3 and 4.

While linear error theory is a convenient way to handle uncertainty, it is by no

means the only way. In the early 2000s, researchers began exploring the application of more novel methods of nonlinear uncertainty propagation to the orbit estimation problem. Park and Scheeres [45] discussed using the Fokker-Planck equations for this along with what they term state transition tensors (STT), essentially including higher-order terms from the Taylor series expansion of the dynamics, accounting for some of the nonlinearity. Fujimoto, Scheeres, and Alfriend [46] extend this work by showing an analytical solution of the Fokker-Planck equations for the special case of deterministic Hamiltonian systems. STTs are also a main part of the work by Majji, Junkins, and Turner [47], which presents the Jth Moment Extended Kalman Filter (JMEKF), an extension of the first-order EKF that includes the higher-order nonlinear effects in a familiar Kalman filter framework. These methods allow for more accurate transformations of the uncertainty distribution by accounting for effects that are normally removed when relying on linearization alone.

Majji, Weisman, and Alfriend [48] present a similar approach, solving the Liouville equation (which describes distribution dynamics in phase space) for Keplerian motion through the use of the transformation of variables (TOV) technique. Application of the TOV technique, which yields an exact transformation of a probability distribution function, to orbit uncertainty problems was expanded by Weisman and Jah [49] and Weisman, Majji, and Alfriend [50].

Terejanu, Singla, Singh, and Scott [51] propose modeling the probability distribution function as a mixture of Gaussians, a technique that had been proposed some years before for the general estimation problem (see Alspach and Sorenson [52] and Ito and Xiong [53]). The benefit of this approach is that existing estima-

tion techniques such as Extended Kalman Filters can be easily adapted to handle an arbitrarily complex, non-Gaussian PDF simply by operating on each (Gaussian) component. The Gaussian mixture (also referred to as Gaussian sum) modeling method has proved to be very useful and flexible, and several groups have pursued this line of research, with significant contributions from Giza *et al.* [54], Horwood, Aragon, and Poore [55], DeMars, Bishop, and Jah [56], DeMars and Jah [57], and Vishwajeet, Singla, and Jah [58]

More recently, the polynomial chaos expansion technique for uncertainty propagation has been explored by Brandon Jones and Alireza Doostan. They have applied it to the general problem of uncertainty propagation [59], as well as the more specific problem of estimating probability of collision between satellites [60], [61], and have demonstrated good results and performance.

Horwood *et al.* [62] propose the Gauss von Mises (GVM) family of probability distribution functions. Essentially, for a set of orbital elements (see Section 3.1), uncertainty in the angular coordinate (or "anomaly") describing the instantaneous position along the orbit is defined on the circle, instead of as an unbounded Cartesian quantity. While other orbital elements are also angular variables, the argument is made that because they do not vary as quickly, and because typical uncertainties in these coordinates is small (compared to the anomaly), they may be treated as Cartesian values without much impact. Thus, the full orbital element state space is defined on a cylindrical manifold (five Cartesian dimensions and one circular). This preserves a good deal of uncertainty information for longer propagation times when compared to treating all six dimensions as unbounded Cartesian quantities. The

authors discuss the concept of covariance realism, which formalizes the concept of comparing the validity of covariance matrices, and they implement an estimator with the same computational cost as the Unscented Kalman Filter (UKF) that maintains a realistic covariance "for up to eight times as long" [62, p. 276]. Horwood and Poore [63] present a nice comparative study of different uncertainty propagation methods, including their GVM-based one, which outperforms the others chosen for the study. The authors find that GVM-based filters outperform the EKF and UKF in all cases, and the equinoctial orbital element set yields a more realistic covariance for much longer propagation times than Cartesian coordinates, as discussed previously.

The present work relies mainly on judicious selection of the coordinate system to alleviate many of the adverse effects of linearized uncertainty propagation. However, several more practical problems in space surveillance, such as probability of collision estimation, require a Cartesian representation. As the previous work has shown, uncertainty represented by a multivariate Gaussian in Cartesian space can be quite unrealistic, and so more exotic methods of representation and propagation become necessary. While this dissertation demonstrates that the cluster tracking problem is viable using linearization methods in equinoctial orbital element space, downstream analysis, which may be necessary to perform in Cartesian space, may require the more advanced uncertainty quantification techniques discussed above.

2.1.2 Multiple Target Tracking (MTT)

Multiple target tracking involves updating the dynamic states of more than one object using information derived from various types of sensors. While space surveillance has been a multiple target tracking problem since the beginning,[3] target ambiguity was rarely an issue, meaning the problem could be simplified to a collection of many single target tracking problems. Most of the literature in this section focuses on multiple target tracking in the nontrivial case, in which target ambiguity is a central issue.

A large portion of the MTT literature deals with the steps leading up to state estimation, namely observation and track association.[4] Given a scene with multiple targets and sensor noise, it is often difficult to assign sensor returns to the correct target. Some methods to address this use weighted combinations of multiple ambiguous measurements to update a target's state (e.g. joint probabilistic data association), while others assign a single observation to a single target, often hypothesizing many such assignments (e.g. multiple hypothesis tracking). Both these "hard" and "soft" assignment methods differ from the group tracking methodology, in which a collection of nearby observations is used collectively to update some set of parameters describing the entire group of targets.

Until fairly recently, multiple target tracking techniques have not been particularly relevant to the astrodynamics community due to the confluence of the historically low population density of satellites, the lack of cluster or formation flying

[3] *Sputnik-1*'s launch vehicle upper stage also attained orbit.

[4] A track in this context is simply a series of observations that are taken from a single target.

missions, and sensors that could not resolve small debris. The modern spaceflight environment is now changing such that all three of these are no longer as true as they once were (see Chapter 1), and there is a small but growing body of multiple target tracking literature specifically geared toward the space surveillance field.

Binz and Healy [64] surveyed several multiple target tracking techniques, including nearest neighbor, probabilistic data association (PDA) and multiple-hypothesis tracking (MHT), and applied them to a simple satellite tracking scenario. Aristoff *et al.* [65] and Singh *et al.* [34] presented an implementation of their MHT tracker applied to the space surveillance problem, and demonstrated excellent performance and results for catalog-level processing.

While some of the literature in applying these traditional MTT techniques to space surveillance problems demonstrates success, no work to date has demonstrated the ability to process and characterize a large-scale explosion or breakup on short timelines (single orbital periods after the event). This is less a failing of the techniques themselves, and more an indication of the fundamental limitations of precise target tracking as a whole. Sensor limitations, data sparsity, and surveillance coverage constraints all limit the effectiveness of target tracking methods for this particular problem.

2.1.3 Debris Modeling

Satellites may break up for any number of reasons, including collision with other objects, a loss of structural integrity, or stored energy sources (such as batteries or

propellant stores) suddenly failing [66]. In all of these cases, the imparted energy on the resulting fragments is a fraction of the total energy of the original orbit, meaning the fragments continue to travel along a similar path until differential perturbations cause them to drift farther apart from one another.

Most of the work involving orbital debris dynamics after a breakup or collision has focused on modeling the debris field given some prior assumptions about the nature of the event itself. Barrows, Swinerd, and Crowther [67] presents a comprehensive review of this work up to 1996, when it was published. While the authors of this work split the literature into event modeling and debris field propagation, they note that the former is "very much a black art", and the majority of the literature surveyed falls into the latter category.

Chobotov and Spencer [68] pioneered a simple short-term model based on a linear state transition matrix derived using the Clohessey-Wiltshire equations for relative motion, and described the cloud shape as a "pulsating ellipsoid." McKnight [69] separates the cloud evolution into phases. The first phase is immediately following a breakup, before the longer-term effects of in-track drift due to differences in semimajor axis or plane drift due to Earth's oblateness manifest, when the cloud can still be described as a pulsating ellipsoid. Tracking the cloud during this phase is the primary focus of the work in this dissertation. The next phase occurs after the change in orbit energy has caused the debris pieces to distribute along the orbit track (but before the orbit planes have diverged significantly), forming a torus. Finally, after several days or weeks, the torus is spread out around the Earth's rotational axis, forming a debris band, which is much longer-lasting than the previous phases.

Jehn [21] describes similar phases, with an added "transition phase" after the torus phase but before the debris field has spread uniformly around Earth's rotational axis. Figure 2.2 illustrates these four phases.

Ashenberg [70] derives several equations for the time spent in each of these phases, under the assumption of an isotropic breakup in low-Earth orbit, where the primary perturbing forces are due to Earth's oblateness and atmospheric drag. Typically, the time from fragmentation to full torus formation is about 10-15 orbit periods[5]. The typical time for the "band" phase to start is several hundred orbit periods.

As Hujsak [18] and Chobotov *et al.* [19] point out, there is another notable feature during the pulsating ellipsoid phase: the presence of pinch points located at the initial fragmentation point and half an orbit later (the latter is also referred to as the pinch wedge). In these regions, the pulsating ellipsoid collapses, leading to a spike in density, as Hujsak and others have shown. Later work by Healy *et al.* [17] revealed the existence of other features in the early phases of a fragmentation by exactly propagating the debris distribution.

Letizia, Colombo, and Lewis [15], extending the work of McInnes [16], take a different approach than the previous authors–solving the continuity equation for fragment density to obtain an analytical expression for the evolution of a debris cloud. This method is more accurate than the linearized propagation method compared to full numerical propagation of individual pieces, as demonstrated by Letizia,

[5]The author notes that the expression given for this is an approximation, and "mostly accurate for a weak impulse" [70].

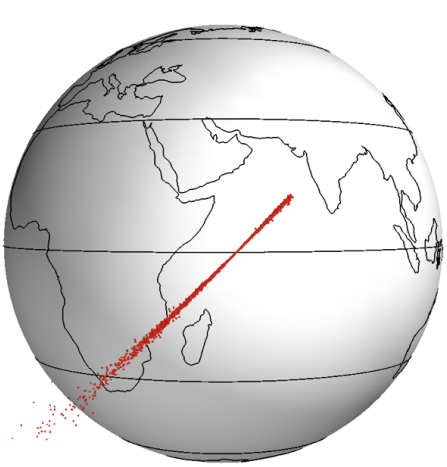

(a) Ellipsoid (90 minutes)

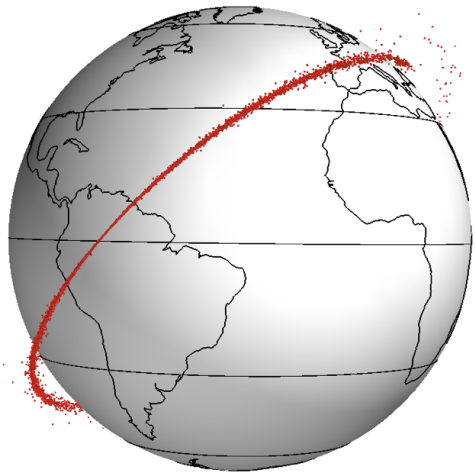

(b) Torus (9 hours)

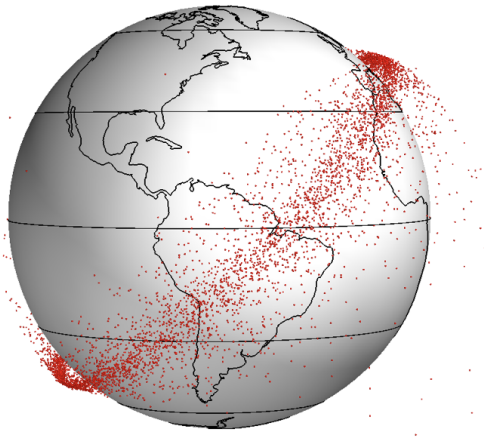

(c) Transition (50 days)

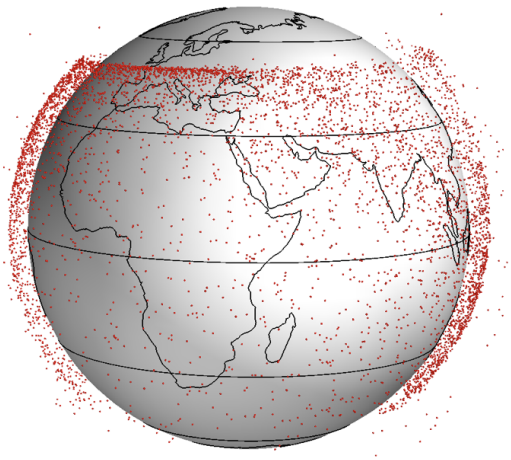

(d) Band (200 days)

Figure 2.2: Illustration of the four breakup phases described in Jehn [21]. The approximate elapsed time since the breakup for each phase is given in parentheses.

Colombo, and Lewis [15]. However, it is only valid once the band phase has been reached, and numerical propagation of the individual pieces is required during the earlier phases. The authors also use a Kolmogorov-Smirnoff test to more precisely define the start of the "band" phase, and find that the approximate band formation times given in, e.g., Ashenberg [70], underestimate the formation time significantly—by a factor of about 3.

2.2 Group & Extended Object Tracking

Compared with the more general field of multiple target tracking, *group tracking* as a subject has not received quite as much attention. This, however, is beginning to change, in part due to new challenges arising in related fields, such as computer vision and video surveillance, which are addressed later in this chapter.

2.2.1 Early Work

Samuel Blackman's widely-cited book *Multiple-Target Tracking with Radar Applications* [71, Ch. 11] contains a chapter outlining the concept of group tracking that is largely based on earlier work by Taenzer [72], [73], Flad [74], and others. The problem is introduced as one of conserving radar resources: tracking a collection of objects as a group requires less computational effort and tracking system bandwidth than tracking the objects individually. Blackman classifies this work into two categories: "centroid group tracking" methods and "formation group tracking" methods. Details for the centroid tracking method are given in an appendix of his

book, but Blackman admits these are sometimes "defined based upon *ad hoc* rules" [71, p. 325].

The main contribution of this early work is the definition of a "generalized residual covariance matrix" [71, p. 314]:

$$S_G = \hat{S}_D + R_G + HPH^T. \tag{2.1}$$

R_G is the measurement covariance matrix, HPH^T represents the uncertainty in the centroid state transformed into measurement space, and $\hat{S}_D$ is called the *dispersion matrix*. The dispersion matrix is a representation of the distribution of the objects in measurement space around the centroid. Importantly, this only describes the expected "spread" in measurement space, and thus does not directly represent the target distribution in state space.

S_G is used in determining which observations belong to the cluster. Once this determination has been made, the centroid of the group of observations is computed and used to update the centroid group target state. A potential drawback noted by Blackman is that the centroid can be quite sensitive to the detection (or non-detection) of one or more group members. This is especially problematic when the number of group members is relatively low. Blackman suggests that probabilistic weighting of observations could help to mitigate these effects. The principles from this early work inform a good deal of the research presented in this dissertation.

In contrast to the centroid group tracking method described above, formation group tracking methods track the states of all member objects as well as the group

state. This may be done either by tracking absolute states of each member object, or by tracking the group centroid state plus relative states of each member. Taenzer's paper [72] has more detail on formation group tracking, including logic for handling missing observations, new targets, merging and splitting targets, and targets that cross paths with each other. In general, Taenzer and Blackman both place most of their emphasis on the process of assigning and processing observations, and rely on slight modifications to typical (at the time) fixed-coefficient or linear Kalman filtering methods for the estimation part of the problem.

A paper by Drummond, Blackman, and Pretrisor [75] is the next major development in the group target tracking literature. The authors identify four "classes" describing the spectrum of group tracking techniques [75, p. 363]:

- Class 1: Cluster tracking only–no individual target information retained
- Class 2: "Simplified" individual target tracks in addition to cluster tracking
- Class 3: Individual target tracks "supplemented by cluster information"
- Class 4: Individual target tracks with no cluster tracking at all

In addition to the problem of tracking groups of targets, the problem is also posed in terms of tracking a single *extended object*, or an object that cannot be reasonably described as a dimensionless point target when compared to the resolution of the observations. In this paper, the authors represent the target group in three dimensions by an ellipsoid in Cartesian space, specified by a matrix. The estimation is done via two separate filters: one to estimate the centroid of the target group,

and one to estimate the ellipsoid extent parameters. The example given is that of a group of targets consisting of a missile reentry vehicle and several decoy targets deployed from a bus. Observations are assumed to exist from passive optical sensors only–an important point, as the measurement function is no longer a simple linear function of the target state. The separation velocity is assumed to be fairly small, and in time the objects separate enough to track individually. This is similar to the problem of a satellite deployment or a low-energy breakup event, and illustrates a situation in which one might move from Class 1 to Class 4-type tracking over time. The authors of [75] demonstrate that such a tracking methodology is feasible.

A good overview of some of the earlier work in the group tracking field is given by Waxman and Drummond [20], who exhaustively cover publications through July 2002. They explain that a large amount of this work was in response to the Strategic Defense Initiative (SDI) focus on reentering ballistic missiles. Some of the papers in this bibliography are only tangentially related to the topic at hand, while many others exist only as difficult-to-find technical reports written for various companies and agencies. However, this comprehensive paper represents the state of the art in group tracking at the time of its publication.

Zimmer and Tsai [76] also address the problem of tracking a missile reentry vehicle and decoys using multiple optical sensors. In contrast to Drummond, Blackman, and Pretrisor [75], the cluster extent was only tracked as the two-dimensional projection of the ellipsoid in the sensor's field of view. The authors do not use the full three-dimensional extent ellipsoid because they found it to be "almost unobservable using angles-only measurements". They use an extended Kalman filter to

track the Cartesian position and velocity of the centroid. Their approach for estimating the projected extent ellipse is similar to those described in Blackman's book [71]. As Blackman noted, the authors find that estimates of the centroid state of the target group tend to be sensitive to unresolved and undetected targets in the group. Additionally, as the relative configuration of the constituent targets changes on a relatively slow time scale, measurement error may be consistently biased. The authors find that the biases induced by this phenomenon are not resolved with the introduction of a second sensor, contrary to what they expected. The authors also note that "although the centroid does not follow the ballistic equation of motion exactly, [they] have found the deviation to be very small" [76, p. 270]. This indication that centroid motion should be carefully considered is what led to the work presented in Chapter 4 of this dissertation.

An insightful paper by Daum and Fitzgerald [77] is cited as motivation for several of these works in group target tracking. In it, the authors derive expressions for the probability of resolution P_R (or the probability that a sensor will be able to distinguish multiple separate objects from a single object) and the probability of correct data association P_{DA} as functions of target separation and sensor resolution. He shows that, in most cases, $P_R < P_{DA}$. In other words, sensor resolution is a more important problem than observation assignment, even though the majority of multiple target tracking literature focuses almost exclusively on the latter. By acknowledging the possibility of ambiguous target resolution, and addressing this through techniques which consider the dynamics of two or more targets collectively (i.e. group target tracking), we may address the general multiple target tracking

problem more effectively.

Yang *et al.* [78] focused on the specific problem of group initialization. Using geometric arguments, the authors present ideas for initializing a group target using Cartesian position vectors as observations. The problem of target grouping/assignment is discussed, as well as the issue of estimating the initial centroid velocity, using information from all observations within a group. Shea [79] discussed group tracking approaches for processing ground moving target indicator radar measurements. While little detailed technical content is presented, the overall process as described seems to use group information while maintaining individual target tracks. The author's two-threshold sequential algorithmic scheme for group partitioning is presented for addressing the "clustering" problem involved in group assignment.

D. J. Salmond and N. J. Gordon published several papers in the late 1990s and early 2000s on group tracking [80]–[82]. Their first papers [80], [81] described their process for applying particle filtering methods to the problem of tracking group (or extended) targets. In all of their papers, the target group was modeled as a set of individual states, with an overall "bulk" component describing the group properties. Both a linear-Gaussian (assuming linear models and Gaussian noise characteristics) multiple hypothesis tracking (MHT) solution and a more generally-applicable particle filtering implementation are shown in [81]. The later papers focus on the particle filtering implementation only. Multiple observation association hypotheses are constructed, though "awkward hypothesis management is avoided" [82, p. 532] through the use of a particle filter. A detailed Bayesian solution is presented in Salmond and Gordon [82] for a known number of targets.

An interesting paper by Connare *et al.* [83] incorporates several techniques for dealing with group target tracking. A variable-structure interactive multiple model (VS-IMM) is used to deal with maneuvering targets, while the joint belief-probability data association (JBPDA) approach is used for observation assignment. Briefly, the JBPDA essentially performs the logic of joint probabilistic data association (JPDA) [84] twice: once for measurement-to-track association ("observation association"), and once for track-to-group association ("track association"). This framework is applied to the problem of maneuvering ground targets being tracked with a high-resolution radar.

This section gives a brief history of early work in the group tracking field. While many of the concepts are still used in contemporary work, the field has largely moved on to alternate approaches to the problem. Two of these approaches are discussed in the following sections.

2.2.2 Finite Set Statistics

Beginning in 1997 [85], Ronald Mahler began publishing about his framework for using the concept of finite-set statistics (FISST) in the context of multi-target tracking problems. Mahler [86] presents an application of FISST to the group target tracking problem. Mahler [87] presents a broad overview of FISST, which he describes as "an 'engineering friendly' version of point process theory" [87, p. 54]. Within the FISST framework, targets, sensors, and measurements are all represented by random finite sets (RFS), and all estimation operations are done on these sets, instead of their

constituent members. Issues such as target death and birth, missed detections, and false detections may all be modeled. This is a potentially powerful method, as it is a more unified approach to the multi-target tracking problem.

Gilholm *et al.* [88] applies a very similar technique involving Poisson models to the group tracking problem. The authors note that a "key aspect of the [Bayesian] solution is the construction of feasible measurement-source assignment hypotheses and the evaluation of the probabilities of these hypotheses" [88, p. 1] and in doing so, they introduce a key aspect of set-based tracking approaches: the ability to avoid a combinatorial explosion by avoiding explicit assignments when there exist many measurements and many targets.

One of the most popular implementations of a FISST-based tracking scheme is the Probability Hypothesis Density (PHD) filter, introduced by Mahler [89]. The overall aim of the PHD filter is to track the first statistical moment (the "intensity") of the multi-target RFS, analogous to a constant-gain Kalman filter [89]. The PHD filter is a simplification of the full Bayesian FISST tracker (which is generally intractable [90, p. 225]), but is not so basic that it has no utility. Vo and Ma [91] addresses the issue of "state extraction" (producing target state estimates from the posterior intensity). Previously, this was done via clustering systems of particles, and Vo presents a way to perform state extraction in a more efficient way using Gaussian mixture models (GMM). Several implementations of FISST using variants of the PHD filter for extended object tracking [92], [93] and group tracking [94] have been reported in the literature, including a paper by Li and Yi [95] that examines the missile tracking problem. A recent paper by Lian *et al.* [96] examines

the performance bounds of suboptimal RFS-based group target trackers. McCabe and DeMars [97], LeGrand and DeMars [98], and DeMars *et al.* [99] have applied the PHD filter to problems ranging from relative multiple target tracking to satellite breakup processing, with mixed results.

Recent work by Hussein *et al.* [100] and Hussein *et al.* [101], [102] is of particular relevance to this dissertation. Hussein and his coauthors introduce the FISST framework to the astrodynamics community with this series of papers, and present implementations using Gaussian mixtures to approximate the problem, similar to the PHD approach. The resulting *AEGIS-FISST* tracker is applied to the problem of tracking modestly sized sets of targets with noisy or missing measurements, and initial results are promising. However, this work is largely proof-of-concept, using a very small number of targets (1-3) and two-body dynamics in two dimensions.

Gehly [103] also presents a study of space object tracking within the FISST framework, and expands on some of the limitations. Computational complexity of the method is the main limitation mentioned, as well as computationally expensive calculations associated with using GMM. This is likely the reason why most published studies present simulations with a relatively small number of objects to be tracked (e.g. 10 in Gehly [103]). These key limitations are also inherent to the method: GMMs are required because the problem is formulated in Cartesian space with large uncertainty regions. It is important to note that the overall goal of these studies is to distinguish and track all objects individually (Classes 3 and 4 as defined above), which is distinctly different from the goal of this dissertation of tracking a large cloud of objects as a single entity, without attempting to solve the association

problem.

2.2.3 Random Matrix Methods

Koch [104] presents another major development in the group tracking literature. The author considers the problem of estimating a group target's "extension", or shape, parameters in a rigorous, Bayesian framework. The extension of a group target is modeled as a symmetric, positive definite (SPD) matrix, which can be thought of as analogous to a single object's covariance matrix, but instead of representing state uncertainty, it represents the physical extension of the group. In this paper, Koch shows that this extension matrix may be estimated via a "natural extension of the standard Kalman filter" [104, p. 210], thus obviating the need for particle filters as in earlier work. He derives this filter with a considerable amount of detail. Performance of this new filter is demonstrated with an aircraft formation flying example, complete with a maneuvering target group and target splitting. This approach is expanded and documented in a seminal 2008 journal article [105]. Importantly, there are several key assumptions made by the authors. First, this method seems limited to a Cartesian state space representation, which restricts its use in the orbit estimation application. Second, measurements are assumed to be full Cartesian position vectors, which, again, is not ideal for this application. Nevertheless, considering the group target extent as a random matrix-variate state is another key innovation that enables the work presented in this dissertation.

Feldmann and Franken [106] build on this work, presenting some improve-

ments in the way sensor performance is handled. Particularly, the authors focus on an assumption in Koch [105] that the measurement error "is small and thus neglectable in comparison with the object extension that hence dominates the spread of the measurements" [106, p. 1030]. They then present an alternative to Koch's approach that allows for more robust tracking when sensor error does not meet this restrictive assumption. The authors also discuss the notion of a confidence region for the elliptical extension. Essentially, a lower and upper bound on the scale of the extension ellipse may be determined, which yields a sort of elliptical torus shape representing uncertainty in the extension itself. This work was later expanded and published as a journal article [107].

Following along this line of research, Lan and Li [108], [109] published a pair of papers addressing, in part, the same drawback to Koch's random matrix method that Feldmann, Franken, and Koch [107] addressed. The authors argue that the treatment of measurement noise presented in Feldmann, Franken, and Koch [107] does not lend itself to an easily-derivable Bayesian solution. Further, they point out that it does not consider the case in which the object extension changes with time, or distortion of the extension due to sensor-target geometry. To alleviate these issues, the authors introduce a method that handles these issues by transforming the extension via the similarity transform. With this modification, the Bayesian solution is derived and nicely summarized. The second of these papers [109] focuses on more accurate modeling of an extended object, such as an airplane, through the use of multiple overlapping ellipsoids.

Granström and Orguner [110] presents a new prediction step for the random

matrix method of extended object tracking based on the Kullback-Leibler divergence, a measure of similarity between two probability distribution functions [111]. As the author points out, Koch's original random matrix work predicts the extension parameters "based on simple heuristics [that] increase the extension's covariance, while keeping the expected value constant" [110, p. 1577]. Granström and Orguner show that their method matches or exceeds prior methods for extension prediction via several simulations, including a multiple target tracking problem that utilizes a PHD filter.

Another series of papers by Baum and Hanebeck [112]–[114] considers the problem of estimating the extent of a target group more accurately through the use of random hypersurface models. The earliest paper [112] uses elliptic random hypersurface models, but the later papers [113], [114] allow for more accurate shape estimation through the use of star-convex random hypersurface models. Sun, Li, and Lan [115] have similar objectives, and the authors propose the use of support functions and extended Gaussian images to more accurately model complex target geometries.

Wieneke and Koch [116] presents an extension to the random matrix method applicable to multiple extended objects. For this, a probabilistic multi-hypothesis (PMHT) tracker is proposed, which "simultaneously estimates the ellipsoidal shape and the kinematics of each object using expectation-maximization (EM)" [116, p. 2349].

Pang, Li, and Godsill [117], [118] presented work focused on modeling intra-group interactions using stochastic differential equations with repulsive forces. In

doing so, the authors also developed an efficient filtering algorithm using Markov Chain Monte Carlo (which they dub "MCMC-particles") in order to efficiently draw samples from high-dimensional probability distributions, avoiding the inefficiency of using a sequential Monte Carlo method (e.g. particle filtering), as with earlier work. This was adopted for use in other problems, including a space object tracking problem [119]. In this paper, the authors tracked all individual objects, as well as the group centroid and an existence variable, using the MCMC-particles algorithm for filtering and JPDA for measurement association. They noted better tracker performance when tracking the group centroid. Preliminary work along the same line was presented by Binz and Healy [120], although this focused more on modeling the centroid itself, and did not address the issue of estimation.

A very good overview of the various modern schools of thought on the subject of group tracking is given by Mihaylova *et al.* [121]. The authors break the subject down into sub-problems, and summarize the solution methods for each, with emphasis on sequential Monte Carlo (SMC, also known as particle filtering) techniques. For small groups with non-interacting members, standard Bayesian tracking techniques suffice; interacting members of small groups require joint estimation of the group structure and individual states. Conversely, and more relevant to the debris tracking problem, large groups (defined as having "many" components that cannot be easily distinguished) rely on either set-theoretic approaches or other Bayesian methods, such as random matrix methods, which are discussed elsewhere in this chapter.

2.2.4 Computer Vision

Some applications of computer vision involve tracking objects or targets in crowded scenes, such as tracking individual people in a crowd. As such, the computer vision/video surveillance field has spurred several advances in group tracking, and thus deserves special mention here. Gennari and Hager [122] use a Kalman filter to estimate the group centroid state, cardinality, and extension, similar to early group target tracking work. The study uses probabilistic data association (PDA) to associate measurements with groups. Duan *et al.* [123] examine the extended object tracking problem from a robotics/machine vision perspective. Bazzani, Cristani, and Murino [124] use a particle filtering approach to perform joint tracking of individuals and groups in surveillance video. A good discussion of multiple extended object tracking is presented by Cancela *et al.* [125], along with an interesting application of the Hellinger distance as a metric between histograms in order to track "appearance" between video frames, which could be adapted for the more general group target tracking problem. In fact, several of the metrics explored in Chapter 5 originate from the computer vision field.

2.3 Summary

While the space object tracking field has many mature methods for tracking individual point targets, these methods fail in cases where many targets exist in similar orbits. Indeed, few studies have even attempted to address this problem, with the exception of the FISST-based methods mentioned–and the limitations of these are

discussed above.

Instead, we can draw inspiration from the broader multi-target tracking field, in which several attempts have been made to address this problem. Group tracking is chosen as the most promising approach, as it eliminates the computationally-expensive association problem and conceptually matches the stated motivations. Unfortunately, most studies use Cartesian coordinates in their state space representations, which is ill-suited for the orbit problem due to the nonlinearity of many of the operations involved. As such, this dissertation will use the findings that more natural orbital element state space representations reduce the effects of this nonlinearity. Additionally, we can extend the use of random matrix methods to represent distributions across ellipsoids in such state space.

Chapter 3: Extent Modeling

This chapter presents the background necessary to model a cluster of objects in orbit. The various relevant coordinate systems are discussed, followed by a description of the method used for satellite propagation, including the relevant orbit perturbations. Finally, the ellipsoid representation of the cluster is discussed, along with various tools for working with such a representation.

3.1 Coordinate Systems

As discussed in Chapter 2, there are several options for describing the motion of a satellite in orbit, and careful coordinate system selection is key.

The most common coordinate system used to describe an orbit is the set of classical orbital elements [126, p. 104]. Details can be found in any orbital mechanics reference, including the aforementioned reference, and the elements are listed here for completeness:

a semimajor axis

e eccentricity

i inclination

Ω right ascension of the ascending node

ω argument of periapsis

ν true anomaly

The elements corresponding to orientation of the orbital plane are illustrated in Figure 3.1.

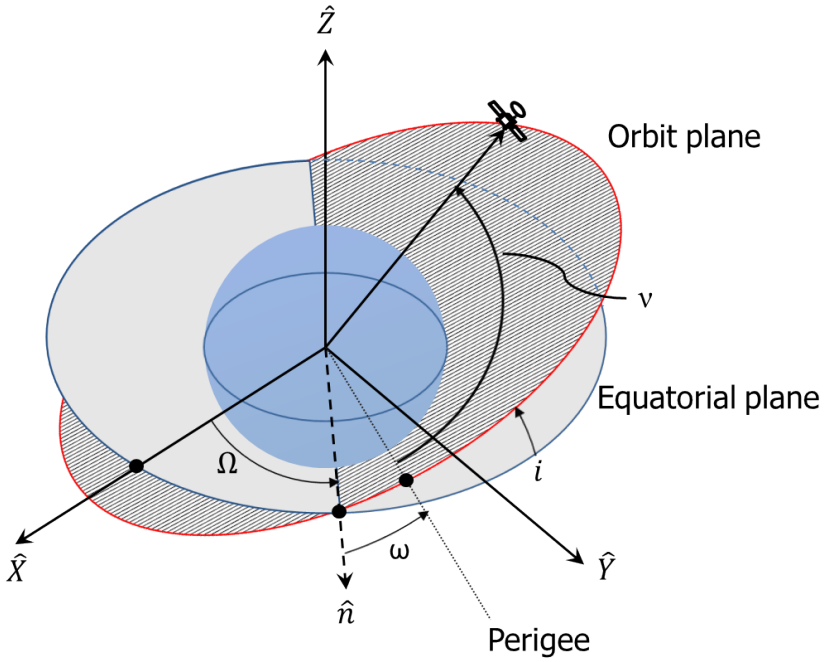

Figure 3.1: Illustration of the classical orbital elements.

Another representation, and the coordinate system used for the majority of this dissertation, is the equinoctial element set $(n, a_f, a_g, \chi, \psi, \lambda)$ [127], [128], which is nonsingular except for the case of retrograde equatorial orbits. These elements are defined in terms of the better-known classical orbital elements as follows:

$$
\begin{aligned}
n &= \sqrt{\frac{\mu}{a^3}} \\
a_f &= e\cos(\omega + \Omega) \\
a_g &= e\sin(\omega + \Omega) \\
\chi &= \tan\frac{i}{2}\sin\Omega \\
\psi &= \tan\frac{i}{2}\cos\Omega \\
\lambda &= M + \omega + \Omega.
\end{aligned}
\tag{3.1}
$$

There is significant variation in the literature concerning equinoctial elements. The set presented here is identical to that published by Arsenault, Ford, and Koskela [127]. Often, the mean motion n will be replaced with either semimajor axis a (as in Montenbruck and Gill [129, p. 30]), or the semilatus rectum, as in Walker, Ireland, and Owens [130]. Another common difference is in the variable names themselves: (a_f, a_g) are often denoted (k, h), and (χ, ψ) as (p, q). Additionally, the set presented by Walker, Ireland, and Owens [130] introduces the "retrograde factor" I, which removes the singularity for retrograde equatorial orbits. The history of the different equinoctial element sets is nicely summarized in an appendix of Wright [128]. Figure 3.2 shows the definition of the equinoctial reference axes $(\hat{X}_{eq}, \hat{Y}, \hat{Z}_{eq})$ with respect to the inertial frame $(\hat{X}, \hat{Y}, \hat{Z})$ and the node vector $\hat{n}$.

As discussed in Chapter 2, distributions described in equinoctial elements have been shown to remain valid for longer periods of time than other descriptions such as Cartesian state vectors or classical orbital elements [63], this is the coordinate

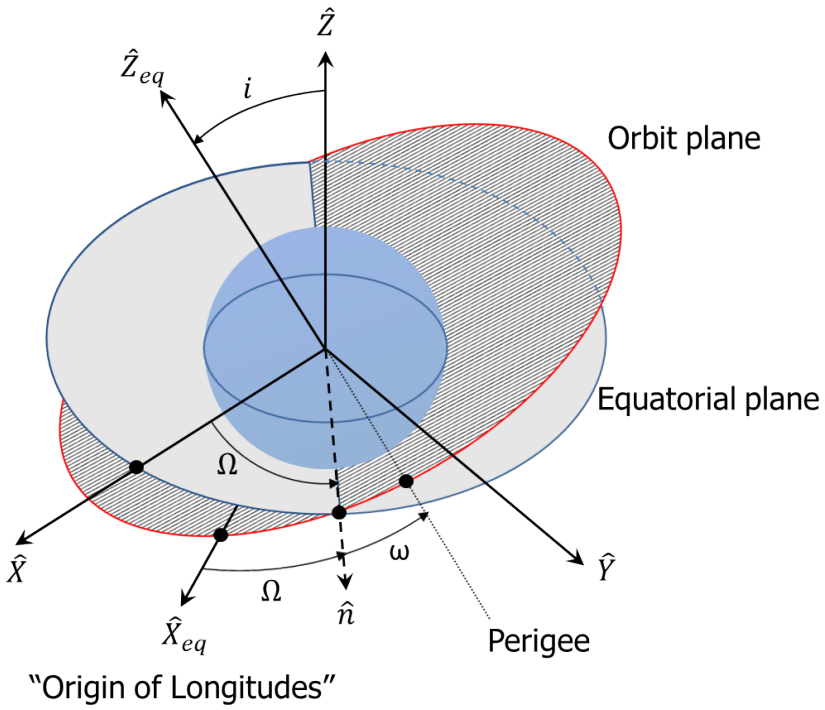

Figure 3.2: The equinoctial frame (recreated from Long *et al.* [28, p. 3-9]).

system used in this dissertation. The following section describes a method for high-accuracy orbit propagation using the equinoctial element set.

3.2 Dynamics

Two-body motion is simple to describe in equinoctial element space. All elements remain constant, except for the mean longitude λ, which changes as a linear function

of the element n:

$$\begin{aligned} n(t) &= n_0 \\ a_f(t) &= a_{f0} \\ a_g(t) &= a_{g0} \\ \chi(t) &= \chi_0 \\ \psi(t) &= \psi_0 \\ \lambda(t) &= \lambda_0 + n_0(t - t_0). \end{aligned} \tag{3.2}$$

3.2.1 Variation of Parameters

In order to incorporate perturbing accelerations into the equations of motion, expressions for the time rate of change of the equinoctial elements are needed. Gauss's variational equations provide such a framework, and the expressions are presented here. Using these equations has the benefit of allowing us to integrate the equations of motion directly using our preferred coordinate system, without needing to perform computationally-expensive and potentially error-prone coordinate transformations at every time step. Following Long *et al.* [28], we can define the time rate of change of a given "slow" element α as

$$\frac{d\alpha}{dt} = \frac{\partial \alpha}{\partial \boldsymbol{\dot{x}}} \overline{P}, \tag{3.3}$$

where $\dot{\boldsymbol{x}}$ is the inertial velocity vector, and $\overline{P}$ is the perturbing acceleration vector in inertial coordinates. For the "fast" coordinate (also called the time element) β, the time derivative is defined as

$$\frac{d\beta}{dt} = \beta' + \frac{\partial\beta}{\partial\dot{\boldsymbol{x}}}\overline{P}, \tag{3.4}$$

where β' is the time derivative of β for Keplerian motion. Long et al. gives the time rate of change of semimajor axis [28, p. 5-32]:

$$\frac{da}{dt} = \frac{2\dot{\boldsymbol{x}}}{n^2 a} \cdot \overline{P}. \tag{3.5}$$

In order to get this in terms of the mean motion n, we invoke the chain rule. Combining Equations (3.3) and (3.5), we see that

$$\frac{\partial a}{\partial\dot{\boldsymbol{x}}} = \frac{2\dot{\boldsymbol{x}}}{n^2 a}. \tag{3.6}$$

Then from the definition of mean motion, we see that

$$\begin{aligned} n &= \sqrt{\frac{\mu}{a^3}} \\ \frac{\partial n}{\partial a} &= \frac{\partial}{\partial a}\left[\sqrt{\frac{\mu}{a^3}}\right] = \frac{-3n}{2a}. \end{aligned} \tag{3.7}$$

Then, using the chain rule,

$$\begin{aligned}\frac{\partial n}{\partial \dot{\boldsymbol{x}}} &= \frac{\partial a}{\partial \dot{\boldsymbol{x}}}\frac{\partial n}{\partial a} \\ &= \frac{2\dot{\boldsymbol{x}}}{n^2 a}\cdot\frac{-3n}{2a} = \frac{-3\dot{\boldsymbol{x}}}{na^2},\end{aligned} \tag{3.8}$$

and finally

$$\frac{dn}{dt} = \frac{-3\dot{\boldsymbol{x}}}{na^2}\cdot\overline{P}. \tag{3.9}$$

So, the full set of equations from Long *et al.* [28, Section 5.7.2] (with the replacement of semimajor axis with mean motion) is then:

$$\begin{aligned}\frac{dn}{dt} &= \frac{-3\dot{\boldsymbol{x}}}{na^2}\cdot\overline{P} \\ \frac{da_g}{dt} &= \left[\frac{1}{\mu}\left[(2\dot{X}_1 Y_1 - X_1\dot{Y}_1)\hat{f} - X_1\dot{X}_1\hat{g}\right] + \frac{a_f}{G}(\psi Y_1 - \chi X_1)\hat{w}\right]\cdot\overline{P} \\ \frac{da_f}{dt} &= \left[-\frac{1}{\mu}\left[Y_1\dot{Y}_1\hat{f} - (2X_1\dot{Y}_1 - \dot{X}_1 Y_1)\hat{g}\right] - \frac{a_g}{G}(\psi Y_1 - \chi X_1)\hat{w}\right]\cdot\overline{P} \\ \frac{d\chi}{dt} &= \left[\frac{1+\chi^2+\psi^2}{2G}Y_1\hat{w}\right]\cdot\overline{P} \\ \frac{d\psi}{dt} &= \left[\frac{(1+\chi^2+\psi^2)}{2G}X_1\hat{w}\right]\cdot\overline{P} \\ \frac{d\lambda}{dt} &= \left[n - \frac{2}{na^2}\boldsymbol{x} + \beta\left(a_f\frac{\partial a_g}{\partial\dot{\boldsymbol{x}}} - a_g\frac{\partial a_f}{\partial\dot{\boldsymbol{x}}}\right) + \frac{1}{na^2}(\psi Y_1 - \chi X_1)\hat{w}\right]\cdot\overline{P}.\end{aligned} \tag{3.10}$$

Here, $X_1, \dot{X}_1, Y_1, \dot{Y}_1$ are the components of the position and velocity of the satellite in the equinoctial frame, which is defined by orthogonal axes $(\hat{f}, \hat{g}, \hat{w})$ (also

called $(\hat{X}_{eq}, \hat{Y}_{eq}, \hat{Z}_{eq})$, see Figure 3.2). The factors G and β are defined as

$$G = na^2\sqrt{1 - a_f^2 - a_g^2} \tag{3.11}$$

$$\beta = \frac{1}{1 + \sqrt{1 - a_f^2 - a_g^2}}. \tag{3.12}$$

$\boldsymbol{x}$ and $\dot{\boldsymbol{x}}$ are the Cartesian position and velocity, respectively, expressed in an inertial frame.

These equations need to be numerically integrated using an ordinary differential equation (ODE) solver. In this work, a fourth-order Runge-Kutta method with fifth-order error control and Dormand Prince coefficients is used, mostly for its flexibility and good general performance [131][1].

3.2.2 Perturbation Forces Modeled

In addition to the force of gravity exerted on a satellite by the central body, there are other, smaller forces, known as perturbations. For Earth-orbiting objects, the most important perturbations come from Earth's non-spherical gravity field, atmospheric drag, solar radiation pressure, and gravity effects from other massive bodies. Smaller effects, such as Earth's tides, relativity, and Earth albedo, are typically only useful in very high-precision applications, such as satellite geodesy.

For the extent estimation application, in which the times of interest are only several hours or days after an event, only the stronger perturbative forces are likely to matter: the lower-order spherical harmonics in Earth's gravity field, and atmospheric

[1]Implemented in the DifferentialEquations.jl Julia library[132].

drag for low Earth orbits. The simulations in Chapter 6 include all dynamic effects that are appropriate for the given orbit regime, and while the dynamics model could be simplified to improve performance for an operational implementation, that study is considered future work.

3.2.2.1 Central Body Gravity

Besides atmospheric drag for low-altitude satellites, the dominant perturbation for Earth-orbiting satellites is due to the fact that Earth is not a perfect sphere, but instead is more pear-shaped. King-Hele gives an exaggerated cross-sectional view of the Earth's geoid, reproduced in Figure 3.3. This represents the mass distribution within Earth, not the actual surface shape.

The result of this imperfect mass distribution is that the magnitude and direction of the acceleration due to gravity varies over the satellite's orbit. We can model this complicated gravity field using a spherical harmonic expansion [126, p. 545]. The potential function may be written as [126, p. 543]:

$$U = \frac{\mu}{r}\left[1 + \sum_{l=2}^{\infty}\sum_{m=0}^{l}\left(\frac{R_\oplus}{r}\right)^l P_{l,m}[\sin\phi_{gc}]\left\{C_{l,m}\cos(m\lambda) + S_{l,m}\sin(m\lambda)\right\}\right]. \quad (3.13)$$

In this equation, $P_{l,m}$ is the associated Legendre polynomial of degree l and order m [134, pp. 247–248], $C_{l,m}$ and $S_{l,m}$ are un-normalized gravity field coefficients [135], λ and ϕ_{gc} are the longitude and geocentric latitude of the satellite respectively, $R_\oplus$ is the Earth's mean radius, r is the distance of the satellite from the center of the

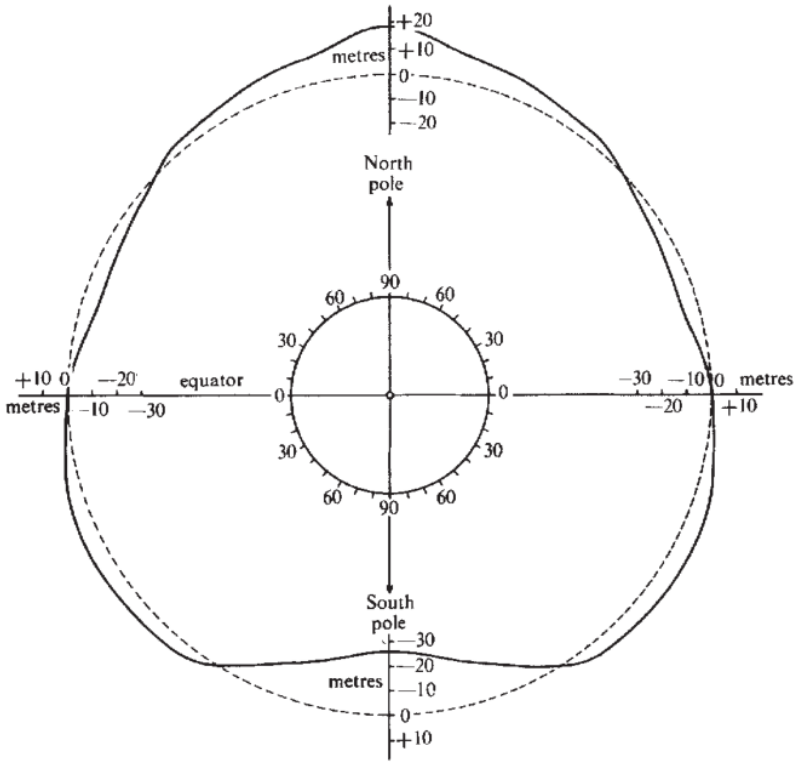

Figure 3.3: Earth's geoid (solid line), relative to the simpler model of a flattened ellipsoid (dashed line) [133].

Earth, and μ is the Earth's gravitational parameter. The equations for the inertial acceleration components are given in Appendix A.

Of highest significance are the so-called zonal harmonics, defined by $m = 0$ [126, p.545], and usually denoted by[2] $J_n = -C_{l,0}$. These are symmetric around the Earth's z-axis. The J_2 zonal harmonic is most significant by several orders of magnitude, and causes secular changes in an orbit's right ascension of the ascending node, argument of perigee, and mean anomaly [126, p. 643]. In fact, a reasonable approximation for an orbit's true perturbed motion can be obtained by modeling only the secular effects of J_2, and this approximation is used often, e.g. when computing the partial derivatives required for orbit determination [127].

3.2.2.2 Atmospheric Drag

For low Earth orbits, the other perturbation that must be taken into account is atmospheric drag. The acceleration on a satellite due to drag is [126, p. 549]

$$\boldsymbol{a}_{drag} = -\frac{1}{2}\frac{c_D A}{m}\rho v_{rel}^2 \frac{\boldsymbol{v}_{rel}}{|\boldsymbol{v}_{rel}|}, \tag{3.14}$$

where ρ is the density of the atmosphere, $\boldsymbol{v}_{rel}$ is the relative velocity of the atmosphere with respect to the satellite, c_D is the drag coefficient, and A/m is the satellite's area to mass ratio. The quantity $c_D A/m$ is often referred to as the satellite's ballistic coefficient. For most satellites, the ballistic coefficient has an attitude dependence and perhaps explicit time dependence, but often, including in this dis-

[2] As King-Hele describes in his book *A Tapestry of Orbits*, the J_n notation was chosen as the standard as an homage to the geophysicist Sir Harold Jeffreys [1, p. 40].

sertation, a constant value is used.

In this work, we use a simple exponential model of atmospheric density [126, p. 562]:

$$\rho = \rho_0 \exp\left[-\frac{h - h_0}{H}\right]. \tag{3.15}$$

Here, ρ_0 is the reference density at the reference height h_0, h is the actual satellite altitude above the ellipsoid, and H is the scale height. Values for H, h_0, and ρ_0 have been tabulated as in, e.g., NASA, NOAA, and USAF [136].

3.2.2.3 Solar Radiation Pressure

Solar radiation pressure results from photons from the Sun reflecting off of a satellite's surface and transferring momentum. For higher orbits, this can have a significant effect on the orbit. The acceleration due to solar radiation pressure can be approximated by the formula given in Kubo-oka and Sengoku [137]:

$$\boldsymbol{a}_{SRP} = -c_R \frac{A}{m} \frac{\Phi}{c} \hat{\boldsymbol{s}}, \tag{3.16}$$

where Φ is the solar constant ($1367\,\mathrm{W/m^2}$ [126, p. 575]), c is the speed of light, $\hat{\boldsymbol{s}}$ is the unit vector from the Sun to the satellite, and c_R is the radiation pressure coefficient:

$$c_R = 1 + \frac{4}{9}\delta, \tag{3.17}$$

where δ is the diffuse reflectivity coefficient [137]. This is the cannonball model for solar radiation pressure, and is an approximate model that is not always applicable,

for instance when the satellite has a very high area to mass ratio (see, e.g., DeMars, Bishop, and Jah [138]). The actual acceleration due to solar radiation pressure is sensitive to many parameters, including variations in the solar flux, the attitude of the satellite, the area to mass ratio, and the surface properties of the material.

3.3 Covariance Transformations

While the previous section describes the motion of individual point masses in orbit, it is often necessary to be able to propagate a distribution. In the case of an arbitrary distribution, this is generally not possible without resorting to Monte Carlo techniques or other numerical approximations, however, if a Gaussian distribution is assumed, a number of simplifications can be made that makes this analytically solvable.

As discussed in Chapter 2, the state uncertainty is often represented as a multivariate Gaussian distribution for convenience–an assumption that is often good near the estimation epoch, but breaks down as the nonlinear effects of propagation begin to manifest themselves. The need to manipulate covariance matrices (transform them from one coordinate system to another, or propagate them in time) often arises, and there are several techniques for doing so.

If we define P_x as the covariance matrix of the Gaussian random vector $\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{\mu}_x, P_x)$, and the linear transformation (e.g., a reference frame rotation) from $\boldsymbol{x}$

to $\boldsymbol{y}$ as $\boldsymbol{y} = M\boldsymbol{x}$, it is a well-known result [139, pp. 75–76] that $\boldsymbol{\mu}_y = M\boldsymbol{\mu}_x$ and

$$\begin{aligned} P_y &\equiv \mathrm{E}\left[(\boldsymbol{y} - \boldsymbol{\mu}_y)(\boldsymbol{y} - \boldsymbol{\mu}_y)^T\right] \\ &= \mathrm{E}\left[(M\boldsymbol{x} - M\boldsymbol{\mu}_x)(M\boldsymbol{x} - M\boldsymbol{\mu}_x)^T\right] \\ &= \mathrm{E}\left[M(\boldsymbol{x} - \boldsymbol{\mu}_x)(\boldsymbol{x} - \boldsymbol{\mu}_x)^T M^T\right] \\ &= M\mathrm{E}\left[(\boldsymbol{x} - \boldsymbol{\mu}_x)(\boldsymbol{x} - \boldsymbol{\mu}_x)^T\right] M^T \\ &= MP_xM^T. \end{aligned} \tag{3.18}$$

Note that this is valid for any linear transformation M. In the more general case where $\boldsymbol{y} = m(\boldsymbol{x})$, we can linearize the transformation with a Taylor series expansion around the mean (keeping in mind that $\boldsymbol{\mu}_y = m(\boldsymbol{\mu}_x)$):

$$\begin{aligned} \boldsymbol{y}|_{\boldsymbol{\mu}_x} &\approx m(\boldsymbol{\mu}_x) + \left.\frac{\partial m}{\partial \boldsymbol{x}}\right|_{\boldsymbol{\mu}_x} (\boldsymbol{x} - \boldsymbol{\mu}_x) + ... \\ &= m(\boldsymbol{\mu}_x) + J(\boldsymbol{x} - \boldsymbol{\mu}_x) + ... \\ P_y &\equiv \mathrm{E}\left[(\boldsymbol{y} - \boldsymbol{\mu}_y)(\boldsymbol{y} - \boldsymbol{\mu}_y)^T\right] \\ &= \mathrm{E}\left[\left(m(\boldsymbol{\mu}_x) + J(\boldsymbol{x} - \boldsymbol{\mu}_x) - \boldsymbol{\mu}_y\right)\left(m(\boldsymbol{\mu}_x) + J(\boldsymbol{x} - \boldsymbol{\mu}_x) - \boldsymbol{\mu}_y\right)^T\right] \\ &= \mathrm{E}\left[\left(J(\boldsymbol{x} - \boldsymbol{\mu}_x\right)\left(J(\boldsymbol{x} - \boldsymbol{\mu}_x\right)^T\right] \\ &= \mathrm{E}\left[J(\boldsymbol{x} - \boldsymbol{\mu}_x)(\boldsymbol{x} - \boldsymbol{\mu}_x)^T J^T\right] \\ &= JP_xJ^T. \end{aligned} \tag{3.19}$$

$J = \frac{\partial m}{\partial \boldsymbol{x}}$ is the *Jacobian matrix*. This process is often referred to as the similarity transformation [140], and that term is used throughout this dissertation, however it

seems to be more commonly applied to a transformation of the form $B = M^{-1}AM$ that preserves the eigenvalues of A [141, p. 232]. Note that this process can be applied to reference frame transformations (which are linear, e.g., quasi-inertial to Earth-fixed), coordinate system transformations (which are generally not linear, e.g., Cartesian to Keplerian elements), and propagation.

An alternative method is to use the unscented transform to propagate or transform the covariance matrix. This method provides better accuracy for long-term propagation since there is no explicit linearization. Full details are given in Chapter 5, but briefly, the unscented transform propagates/transforms the covariance by computing the weighted sample covariance of a set of individually-transformed, deterministically-chosen sigma points. This method has the advantage of not requiring explicit linearization or formation of Jacobian matrices, and has been shown to result in a more realistic covariance when compared to a fully nonlinear transformation via Monte Carlo [142].

Besides the use of the unscented transform in the centroid filter (see Section 5.3.1), the similarity transform is used throughout this dissertation when propagating and transforming covariance and extent matrices.

3.4 Extent Modeling

This section describes the method chosen for this dissertation to describe the size, shape, and orientation of a cluster of objects: the extent.[3] In two dimensions,

[3]Portions of this section have been published previously by the author in Binz and Healy [120], [143].

constant-probability contours of a bivariate Gaussian distribution take the form of ellipses. In three dimensions, these become ellipsoidal constant-probability surfaces. Generalizing to N dimensions, the constant-probability surfaces become hyperellipsoids. Analogous to these constant-probability hyperellipsoidal surfaces, the cluster extent is represented here by a symmetric, positive-definite matrix in equinoctial element space. This matrix $\boldsymbol{X}$ defines the orientation and dimensions of the extent, while the center is the estimated state of the centroid of the cluster $\bar{\boldsymbol{x}}$, such that the cluster surface consists of the level set of points $\boldsymbol{x}$ such that [144, p. 692]

$$(\boldsymbol{x} - \bar{\boldsymbol{x}})^T \boldsymbol{X}^{-1} (\boldsymbol{x} - \bar{\boldsymbol{x}}) = 1. \tag{3.20}$$

No assumption about the distribution of objects within the extent is required, nor is one made. [4]

Conceptually, the ellipsoid described by $\boldsymbol{X}$ is similar to the covariance of a given state estimate. However, there are several key differences. Primarily, the covariance describes the uncertainty of the state estimate, whereas the extent matrix describes a "surface" containing all members of the satellite cluster. Further, because we decouple the problem of estimating the centroid of the cluster from the problem of estimating its extent, both the covariance estimate and the extent matrix estimate may be used, with the expectation that the covariance will be significantly different from the extent. This enables precision tracking of the centroid with uncertainty

[4] As an aside, note that because a symmetric, positive-definite matrix can be represented as an ellipsoid, we will occasionally use the terms "extent matrix" and "extent ellipsoid" interchangeably, particularly when discussing random matrices in Chapter 5.

similar to that of a single tracked object, while maintaining information about the spread of objects for sensor tasking and conjunction warnings.

Prior work has shown that a covariance represented in equinoctial elements remains realistic over longer propagation spans than one represented in other coordinate systems, such as Cartesian [44], [63]. As shown in Section 3.2, the two-body problem is linear in the equinoctial elements, with the first five elements undergoing no change, and the sixth element, mean longitude, changing as $\lambda(t) = \lambda_0 + n(t - t_0)$. As Horwood, Aragon, and Poore point out [55], the fact that, for two-body dynamics, $\lambda(t)$ can be expressed as a linear combination of the state variables themselves, a Gaussian PDF in equinoctial elements will remain Gaussian indefinitely, i.e., if $\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ and $\boldsymbol{y} = F\boldsymbol{x}$, then $\boldsymbol{y} \sim \mathcal{N}(F\boldsymbol{\mu}, F\boldsymbol{\Sigma}F^T)$ [145, p. 285]. No such statements about linearity can be made for the Keplerian or Cartesian satellite state representations (or even for an equinoctial element set that uses semimajor axis instead of mean motion), and so the Gaussian assumption about PDFs in these spaces does not remain valid for very long [44].[5]

3.4.1 Computing Extent

While the previous section introduces an unambiguous definition for *extent*, there are multiple ways to compute extent given an actual set of points. This section discusses the two options considered for this work.

First, the unbiased sample covariance of the set of N points, given the sample

[5] It should be noted here that the fact that λ is an angle means that it requires special handling in a statistical sense. Horwood and Poore [63] introduces the Gauss von Mises distribution to address this problem specifically, and techniques for computing statistics involving angular quantities are discussed in Chapter 4.

mean $\bar{\boldsymbol{y}}$, is defined as:

$$\boldsymbol{Y}_s = \frac{1}{N-1}\sum_{i=1}^{N}(\boldsymbol{y}_i - \bar{\boldsymbol{y}})(\boldsymbol{y}_i - \bar{\boldsymbol{y}})^T. \tag{3.21}$$

The ellipsoid described by this matrix, however, corresponds to only a 1σ bound. In order to enclose the entire set of points, consistent with our definition of extent, we then scale $\boldsymbol{Y}_s$ by the largest Mahalanobis distance of the set of points [143]. The Mahalanobis distance of the ith point $\boldsymbol{y}_i$ is defined as [146]

$$k_i = \sqrt{(\boldsymbol{y}_i - \bar{\boldsymbol{y}})^T \boldsymbol{Y}_s^{-1} (\boldsymbol{y}_i - \bar{\boldsymbol{y}})}. \tag{3.22}$$

The extent is then

$$\boldsymbol{Y} = k_{\max}^2 \boldsymbol{Y}_s, \tag{3.23}$$

where

$$k_{\max} = \max\{k_i\}, i \in \{1, 2, ..., N\}. \tag{3.24}$$

An alternative method of calculating extent is to use the minimum-volume enclosing ellipsoid (MVEE) for the set of points. Techniques for doing this have been studied extensively, as MVEEs (also called Loewner or Loewner-John ellipsoids [147]) are useful tools for problems in statistics and optimization [148].

The method for solving for the MVEE used here is the Fedorov-Wynn-Frank-Wolfe method, described in Todd [148]. It requires a set of *centered* points, thus the first step is to transform the entire set $\{\boldsymbol{y}\}$ to $\{\boldsymbol{y}_{\text{centered}}\} = \{\boldsymbol{y}_i - \bar{\boldsymbol{y}}\}\ \forall\ \boldsymbol{y}_i \in \{\boldsymbol{y}\}$, where

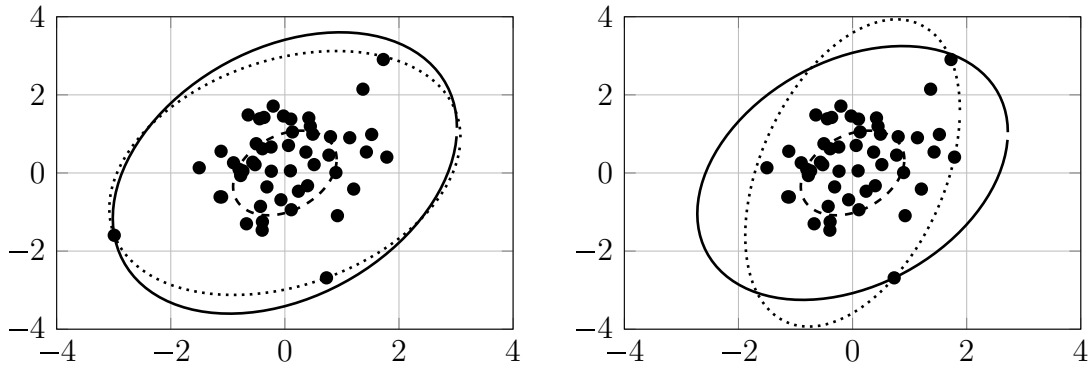

Figure 3.4: Comparison of the MVEE extent (dotted line) and the scaled covariance (solid line) for a set of points. On the right, the farthest point (using Mahalanobis distance) has been removed and the MVEE and scaled covariance ellipses recalculated.

special consideration is taken for angular quantities, as discussed in Chapter 4.

Figure 3.4 shows a set of points from the bivariate unit normal distribution $\mathcal{N}(\mathbf{0}, \mathbb{1}_2)$, along with the 1σ sample covariance ellipse (dashed line), the scaled covariance ellipse (solid line), and the MVEE (dotted line). On the right, the farthest point (using Mahalanobis distance) is removed from the set, and the scaled covariance is re-scaled, while the MVEE is recomputed. Unlike the scaled covariance ellipsoid, the MVEE is dependent only on the outer hull of points, and not the overall structure of the cluster. As such, the orientation of the MVEE is extremely sensitive to individual points, which can be undesirable when cluster members are being dynamically included or excluded from the cluster (see Section 5.2.1).

Due to this sensitivity of the MVEE method to individual points, the estimation process uses the scaled sample covariance, while the output of the filter is compared against a truth extent in state space generated as a MVEE. The relative smoothness of transitions (when individual objects get included or excluded) allows

for a better-performing estimator, while a comparison against a MVEE truth extent ensures the description is realistic and accurate with respect to this definition of extent.

When computing extent in equinoctial element space, the tolerance of this iterative method is chosen to be 10^{-5}, as this yields consistent results for this problem while being fast to compute. At this tolerance level, convergence usually takes several hundred iterations, however the overall computation time is still acceptably low (on the order of 50 milliseconds on a single 1.7 GHz processor core).

3.4.2 Sampling the Extent

It is often desirable to be able to generate points uniformly within a given extent matrix, e.g. for visualization. This section outlines the process for doing so.

First, we generate a set of points uniformly distributed on the surface of a n-dimensional hypersphere, where n is the dimension of the extent matrix. This is a straightforward process: generate n random draws from the unit normal distribution: $x_i \sim \mathcal{N}(0,1)$, $i = 1, 2, ..., n$. Then, the vector

$$\boldsymbol{x} = \frac{1}{\sqrt{\sum_{i=1}^{n} x_i^2}} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \tag{3.25}$$

is distributed normally over the surface of the unit n-dimensional hypersphere [149].

This essentially chooses a uniformly-distributed random direction for a given point.

In order to sample from within the volume of the hypersphere, Harman and Lacko [150] shows that multiplying the above vector by a factor of $u^{1/n}$, where $u \sim \mathcal{U}(0,1)$ is the uniform distribution, we get a uniformly-distributed point inside the hypersphere: $u^{1/n}\boldsymbol{x}$.

Next, the points must be scaled and rotated according to the eigenvalues and eigenvectors of the extent matrix. The vector $u^{1/n}\boldsymbol{x}$ may be scaled by multiplying each component by the square root of the corresponding eigenvalue v_i:

$$\boldsymbol{x}_{\text{scaled}} = u^{1/n} \begin{bmatrix} \sqrt{v_1}x_1 \\ \sqrt{v_2}x_2 \\ \vdots \\ \sqrt{v_n}x_n \end{bmatrix} \tag{3.26}$$

The rotation is dictated by the eigenvectors of the extent matrix $\boldsymbol{d}_i$. If $D = [\boldsymbol{d}_1|\boldsymbol{d}_2|\ldots|\boldsymbol{d}_n]$ is the matrix of eigenvectors, this is essentially a rotation matrix from our original frame to the frame defined by the principal axes of the extent, and the rotation can be written as

$$\boldsymbol{x}_{\text{ellipsoid}} = D\boldsymbol{x}_{\text{scaled}}. \tag{3.27}$$

The end result is a point from a uniform distribution over the hyperellipsoid defined by the extent matrix centered at the origin. All that remains is to translate each point: $\boldsymbol{x}_{\text{extent}} = \boldsymbol{x}_{\text{ellipsoid}} + \boldsymbol{x}_{\text{centroid}}$.

As an example, this process is used to generate points over the volume of the

three-dimensional ellipsoid defined by $\boldsymbol{X}^T\boldsymbol{A}\boldsymbol{X} = 1$, where

$$\boldsymbol{A} = \begin{bmatrix} 1/16 & 1/32 & 1/16 \\ 1/32 & 1/4 & 1/16 \\ 1/16 & 1/16 & 1/4 \end{bmatrix}. \tag{3.28}$$

Figure 3.5 shows a set of 500 samples as black dots, along with the ellipsoid surface.

This chapter presented the preliminaries necessary to simulate the dynamics of a point object in space, methods for generalizing this to distributions of particles, and techniques for representing and manipulating the extent of a cluster of particles. These tools will be applied to the cluster estimation problem in the following two chapters.

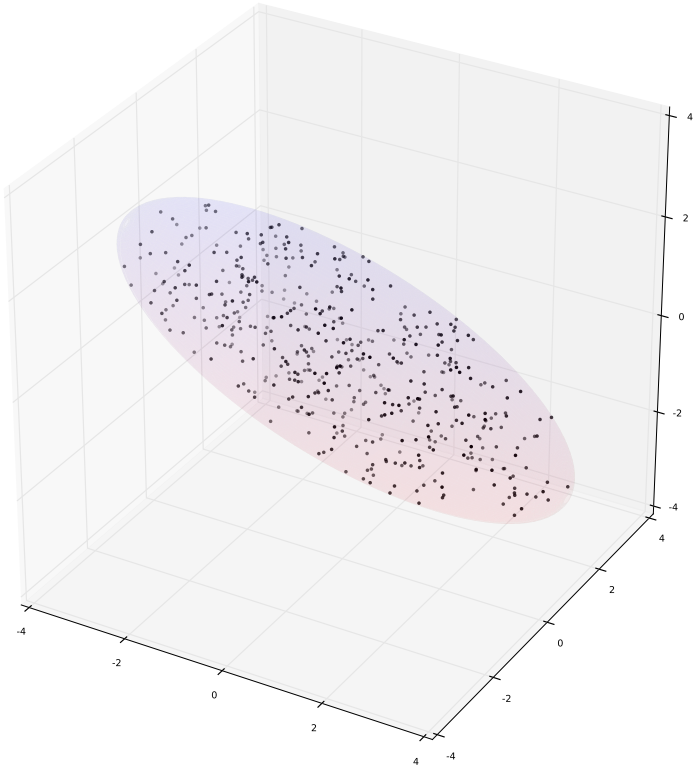

Figure 3.5: Uniform sampling over the volume of an ellipsoid.

Chapter 4: Centroid Dynamics

As discussed in the previous chapters, this work decouples the cluster estimation problem into two parts: estimation of the centroid of the cluster, and estimation of its extent.[1] This chapter discusses the former: the definition of the cluster centroid, and its dynamic behavior as compared to a single orbiting object.

The centroid of a cluster of satellites is defined as the "average orbit" over all of the satellites in the cluster. Of course, there are multiple ways of defining this based on which parameterization of the orbit is used, and the first part of this chapter compares several.

Centroid estimation involves finding the geometric center of all gated, or included, observations, then using this "virtual observation" to update the centroid state of the cluster. Near the limit of a sensor's resolution, the centroiding may be done implicitly–two or more objects could be so close together that they present a single merged target.

However, the centroid of a group of satellites will not necessarily exhibit the same dynamical behavior of the satellites themselves. One aim of this chapter is to quantify how closely the dynamics of a centroid match those of a physical, orbiting

[1] Portions of this chapter have been published previously by the author in Binz and Healy [120].

object.

This chapter begins with a simple 2D example to illustrate the effects of coordinate system selection. Then, more realistic comparisons between a measurement space centroid and a state space one are made by simulating observations of a satellite cluster from a radar. The last section explores the effects of cluster size and spread on centroid error. The results from this chapter are intended to characterize the motion of a composite centroid satellite state and justify the choice of coordinate system.

4.1 Instantaneous Centroid

As a simple illustration of this process, we begin by examining the centroid in two dimensions. Consider Keplerian motion of two satellites in circular orbits, one with a radius of 7000 km and the other 50 km lower. At $t = 0$ they share the same angular position, that is, the lower object is directly below the higher one. In the plane of the orbit [151, Section 2.10], it is simple to represent the motion of each of these objects in two-dimensional Cartesian space:

$$x(t) = r\cos(nt) \tag{4.1}$$

$$y(t) = r\sin(nt), \tag{4.2}$$

where $n = \sqrt{\mu/r^3}$ is the angular velocity, or mean motion. The representation in polar form is:

$$r = r_0 = \text{constant} \tag{4.3}$$

$$M(t) = nt + M_0 \tag{4.4}$$

where M is the polar angle (mean anomaly) and r is radius.[2]

Taking the components of the instantaneous state of each of the two objects, the Cartesian centroid of the combined set is simply

$$\overline{x} = \frac{x_1 + x_2}{2} \tag{4.5}$$

$$\overline{y} = \frac{y_1 + y_2}{2}, \tag{4.6}$$

and the centroid in polar form is similarly

$$\overline{r} = \frac{r_1 + r_2}{2} \tag{4.7}$$

$$\overline{M} = \frac{M_1 + M_2}{2}, \tag{4.8}$$

ignoring, for now, the fact that M is an angular variable.

Both objects, along with the Cartesian centroid (calculated at each time step), are plotted in Figure 4.1 for $t =$ 400–500 seconds, with markers every ten seconds. The faster angular rate of the lower object is readily apparent. Plotting the magnitude of the Cartesian centroid position (Figure 4.2), however, identifies a key

[2]For this example, we define M to be the angle from the $+x$ axis to the position vector.

problem with centroiding in Cartesian coordinates: the radius of the centroided orbit decreases in time. In the extreme case, when the two objects are 180° out of phase, the centroid would be nearly at the focus of the orbits. Clearly, a different treatment is required, especially for large separation distances.

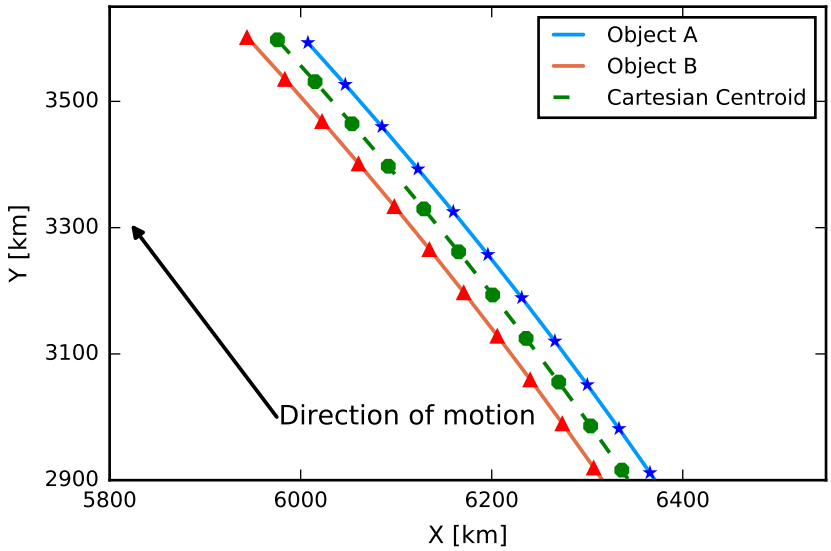

Figure 4.1: Cartesian centroid test case. Object A is in a higher orbit than Object B, resulting in a secular in-track drift.

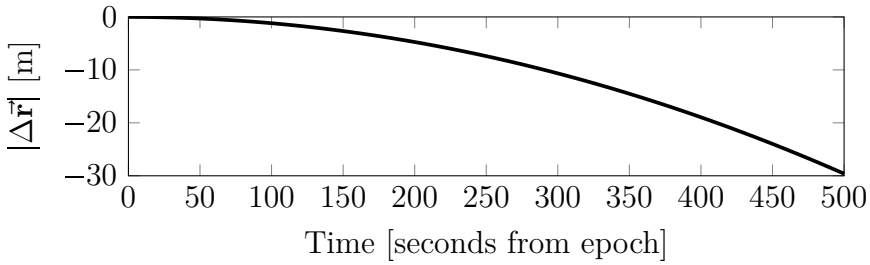

Figure 4.2: Change in position magnitude of the Cartesian centroid in time. Values are relative to the initial centroid radius (6975 km).

The polar form of the centroid will have a constant radius, which is the intuitive result of averaging two circular orbits. However, the angular component must be more carefully analyzed. To get the average direction, we use the tools of *directional statistics*[152, Ch. 2]. The mean direction of a set of angles $\{M_k\}$ is then calculated

as

$$\overline{z} = \frac{1}{N}\sum_{k=1}^{N} \exp(iM_k) \tag{4.9}$$

$$\overline{M} = \text{Arg}(\overline{z}). \tag{4.10}$$

When working with only two objects, the benefit of this is only evident when the angular difference between the two objects is greater than 180° when measured on the range $[0°, 360°)$ (for instance, when crossing the 360° boundary). However, when considering larger numbers and distributions of objects, treating the angular coordinate properly becomes more important, and the difference between the above method and the arithmetic mean become more pronounced. This method of centroiding angular quantities is used throughout the work presented here, both for angular state variables and angle measurements.

Because the angular rate n in Equation (4.4) is dependent on the radius of the orbit, calculating the centroid in polar form has two possible interpretations. In the first, the radius centroid is calculated as in Equation (4.7), and the resulting polar angle, now as a function of time, is

$$M(t) = \sqrt{\frac{\mu}{\overline{r}^3}}t + \overline{M}_0, \tag{4.11}$$

where $\overline{M}_0$ is the average of the mean anomaly at epoch of each constituent satellite, calculated with Equations (4.9) and (4.10). This formulation results in a circular orbit whose radius is the average of the two satellites'.

Alternatively, the average *rate* (mean motion) could be used, in which case the centroid would be calculated as

$$\overline{n} = \frac{1}{2}\sum_{i=1}^{2}\sqrt{\frac{\mu}{r_i^3}} \tag{4.12}$$

$$\overline{r} = \sqrt[3]{\frac{\mu}{\overline{n}^2}}, \tag{4.13}$$

with the polar angle now calculated as

$$M(t) = \overline{n}t + \overline{M}_0. \tag{4.14}$$

Note that because $\overline{n} \neq \sqrt{\mu/\overline{r}^3}$ (when $r_1 \neq r_2$), these two methods yield different orbits. In effect, we can either average the orbit rate, or the orbit radius (or, more generally, the semimajor axis).

It is not immediately obvious which of the averaging methods above yields a more realistic description of the centroid dynamics. In order to compare the two, we first find the average radius centroid using Equations (4.7) and (4.11). This yields a constant orbit radius of $\overline{r} = 6975\,\text{km}$, and a mean anomaly defined by Equation (4.11), with $\overline{M}_0 = 0$. Then, we can use the mean motion centroid as defined in Equation (4.13). The resulting orbit radius from this method is 110 m lower. Each of these is used to generate a centroid orbit state at epoch that may be used for prediction and propagation. The distance between the two centroids can be expressed in closed form as a function in time from the formula for the distance between two points in polar coordinates:

$$d(t) = \sqrt{\left(\frac{r_1 + r_2}{2}\right)^2 + \left(\frac{\mu}{\bar{n}^2}\right)^{2/3} - (r_1 + r_2)\sqrt[3]{\frac{\mu}{\bar{n}^2}}\cos\left[t\left(\bar{n} - \sqrt{\frac{8\mu}{(r_1 + r_2)^3}}\right)\right]}. \tag{4.15}$$

Plotting this out for 1000 seconds, we see in Figure 4.3 that the distance grows from its initial value of about 110 meters to over 200 meters. This distance is periodic with period equal to the synodic period of the two orbits, that is, the reciprocal of the sum of the reciprocals of the individual periods.

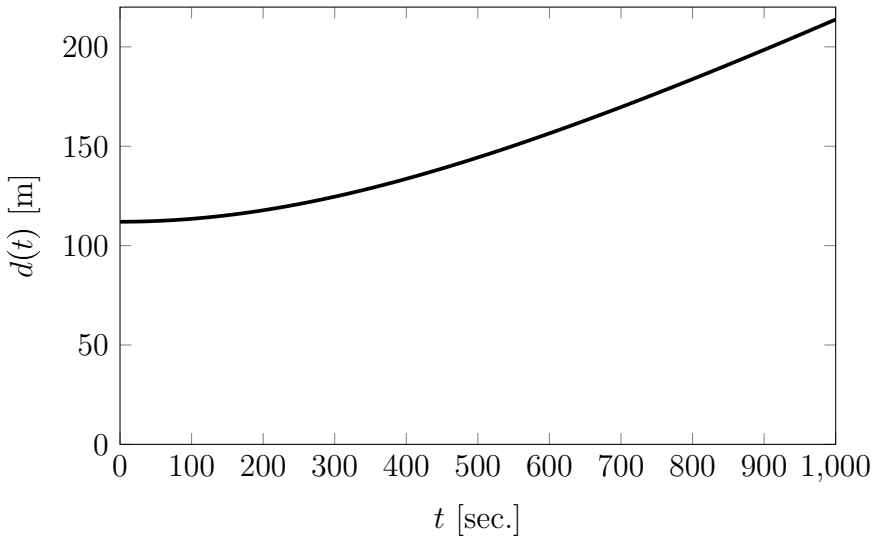

Figure 4.3: Distance between satellites constructed with two different centroiding methods–one with a mean radius, the other with a mean rate.

The next step is to investigate each method's accuracy when compared to the centroid of all observations in a given frame of measurements. Here, centroid measurements are calculated by averaging the range and polar angle to each satellite at each time step.[3] Then, each state-space centroid is propagated, and range and polar angle to the centroid position are calculated. The series of centroided mea-

[3]For simplicity, "measurements" here are taken from the center of the Earth.

surements, in effect, represents an object with average semimajor axis *and* average mean motion, which is not possible for a non-maneuvering satellite. Thus, there will be a difference between the position derived from centroided measurements and the centroided states. The position difference between the mean measurements and each element space centroid for 1000 seconds after epoch is plotted in Figure 4.4.

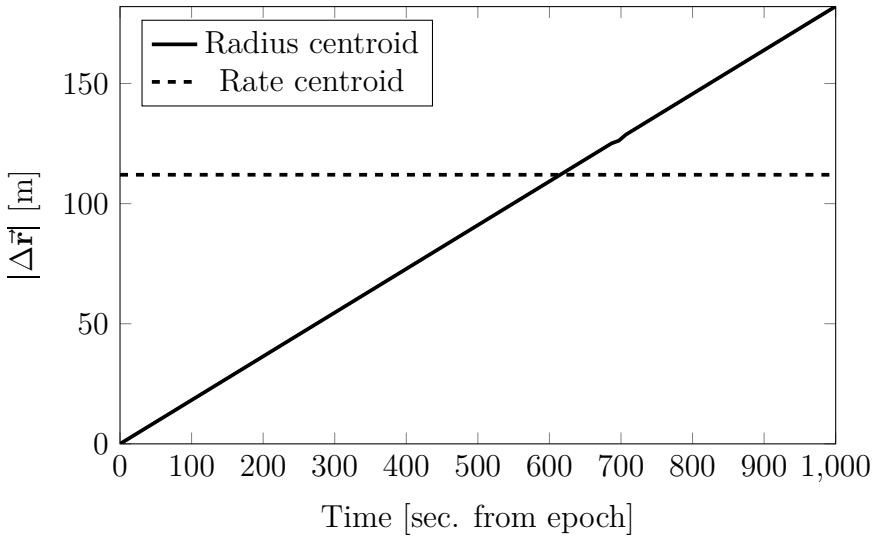

Figure 4.4: Position difference between measurement-space and each state-space centroid technique.

When averaging the angular rate, the resulting polar angle is, in fact, the mean polar angle of the two satellites. Meanwhile the radius (derived from the mean motion) is slightly offset from the mean radius of the two satellites, resulting in a constant bias between the measurement-space and state-space centroids (dashed line in Figure 4.4). When the mean radius is calculated and used to define the mean motion of the centroid, however, the result is a constant drift away from the mean polar angle value, because the rates do not match exactly (solid line in Figure 4.4). This is a consequence of the nonlinear relationship between semimajor axis and

mean motion.

It is important to note that, in this example, the measurements are geocentric range and angle, not topocentric (from the Earth's surface). Later examples use a more realistic scenario with topocentric measurements. Despite the simplicity of this example, it indicates that an orbit with an average mean motion, not semimajor axis, is likely preferable due to the absence of a secularly increasing difference.

We can also address this problem using a relative motion approach. The reference frame used here is centered on one of the satellites, dubbed the primary, and constructed by aligning the first axis with the position vector of the primary ($\hat{R}$), the third along the orbit normal ($\hat{W}$), and the second ($\hat{S}$) in the along-track direction, completing the triad[126, pp. 163–164] (see Figure 4.5). As the orbits are circular, we can use the solutions to Hill's equations for linearized relative motion[153, pp. 84–90]. Designating the higher satellite as the primary (an arbitrary choice), the motion of the lower satellite will simply be a linear drift along the in-track axis. Centroiding in Cartesian space in the RSW frame is essentially finding the midpoint between the second satellite and the origin, and so there are no issues in this frame due to the linearized motion. Further, when the satellites are coelliptic, the in-track velocity term is a linear function of radial displacement ($\dot{y} = -3/2 x_0 n_{\text{primary}}$) [126, p. 403] and so the issues of averaging rate versus averaging semimajor axis do not appear here. As such, determining the centroid relative state at epoch and propagating it independently (using the solutions to Hill's equations) yields the same motion as centroiding the relative state at each time step.

It is important to remember the assumptions inherent to the solutions of Hill's

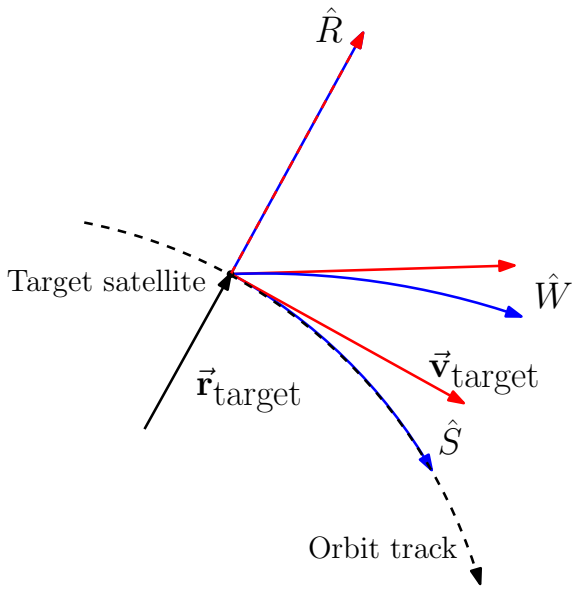

Figure 4.5: *RSW* (red) and curvilinear *RSW* (blue) frames.

equations, however: this is only valid for near-circular orbits, and for small separation distances. After converting the centroid ephemeris back to an inertial frame, we see a departure from the initial centroid orbit radius of 6975 km, as shown in Figure 4.6. This is due to the linear nature of the in-track axis, and is significantly improved through the use of a curvilinear frame (see Figure 4.5),[153, pp. 99–103] in which the $\hat{S}$ and $\hat{W}$ axes follow the curvature of the orbit. Figure 4.7 shows the error in the curvilinear *RSW* frame. The maximum drift in radius of about 10 cm occurs when the satellites have already separated by tens of kilometers. While the centroid orbit radius does drift in time, the error grows at a much slower rate than with the linear *RSW* frame.

This section indicates that, considering the different centroid methods, averaging the angular rate variable (mean motion, here) may be preferable. While the curvilinear relative centroid results in very low drift in the centroid orbit radius,

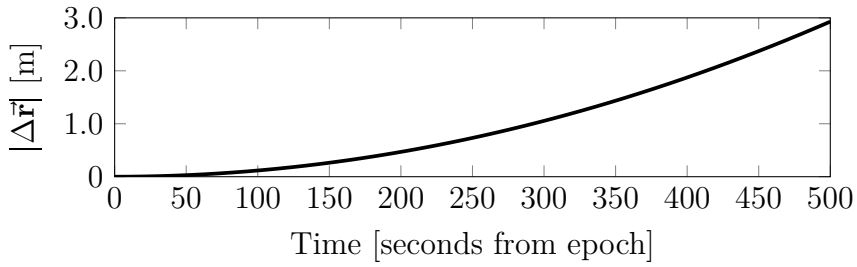

Figure 4.6: Drift in radius resulting from calculating the Cartesian centroid in the RSW frame.

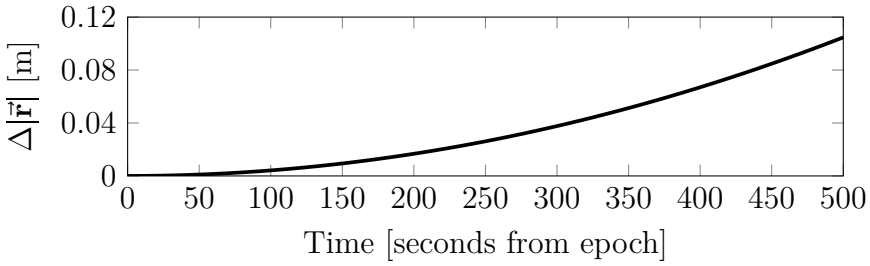

Figure 4.7: Drift in radius resulting from centroiding in a curvilinear relative frame.

the restrictive assumptions required are problematic for the general case. The next section takes this further by analyzing a set of more realistic simulations than the simple illustrative models used in this section.

4.2 Simulation

In order to further investigate the dynamics of the centroid, a scenario with a single ground station and pairs of satellites was constructed. Each scenario begins with a reference orbit with Keplerian elements $(a, e, i, \Omega, u_0) = (7000\,\text{km}, 0, 60^\circ, 175^\circ, 0^\circ)$ where $u_0 = \omega + \nu_0$ is the argument of latitude, used because the orbit is circular at epoch. The simulations include only two-body dynamics.[4] Each satellite is propagated in time, and range, azimuth, and elevation measurements are taken from the

[4]Comparisons with a propagator including the secular effects of the J_2 zonal harmonic show only very slight differences over the time periods considered here.

ground station to both satellites simultaneously. Measurements are assumed to be perfect for our purposes (including perfect observation association), so that we do not conflate error due to the centroid method with error due to measurement noise.

Measurements are taken for both satellites, and they are combined using separate techniques: Equation (4.7) for range values, and Equations (4.9) and (4.10) for the angular components. These averaged measurements are then converted back to an inertial state vector. The combination of range and two angles allows for an exact transformation from centroided measurement to inertial position vector.

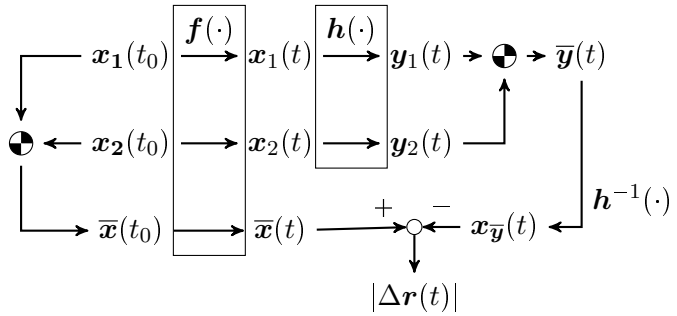

Figure 4.8: Diagram of the comparison process. $\boldsymbol{x}$ are object state vectors, $\boldsymbol{y}$ are measurement vectors, the ◓ symbol represents the centroiding process, $\boldsymbol{f}(\cdot)$ represents the dynamics function (i.e. propagation), and $\boldsymbol{h}(\cdot)$ is the measurement function.

This averaged position vector is compared to the state vector resulting from averaging the full satellite states directly at epoch, using the three methods described above: average semimajor axis, average mean motion, and average initial relative state. The element set used for the first two is the equinoctial coordinate system $(n, a_f, a_g, \chi, \psi, \lambda)$ (discussed in Section 3.1). The elements $(n, a_f, a_g, \chi, \psi)$ are not angular coordinates, and may be averaged as in Equation (4.7). The mean longitude λ is an angular coordinate and must be handled as such by using Equations (4.9)

and (4.10). As explained previously, the state-space centroid representation will be different than the one obtained by averaging in measurement space, since the measurement-space centroid effectively represents an object that does not quite follow Keplerian dynamics. The overall process for comparing the state space centroids with measurement space ones is shown in Figure 4.8.

The first example involves a deployment of one satellite from another shortly prior to the start of the pass, in which the secondary satellite is separated from the primary satellite along the $-\hat{R}$ axis at 1 m/s roughly 6 minutes prior to the ground pass. The magnitude of the position difference between the measurement space centroid and state-space centroids is plotted in Figure 4.9.

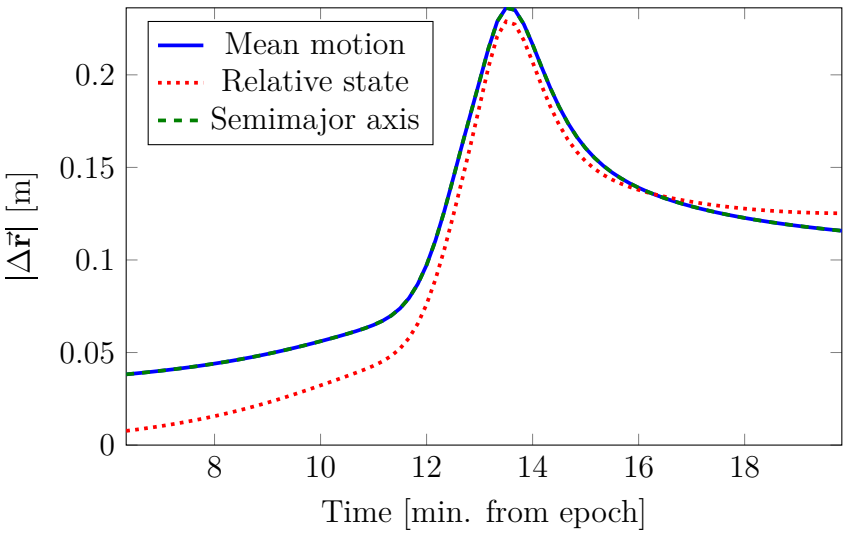

Figure 4.9: Difference between measurement-space and state-space centroid representations for the radial deployment scenario. The curves for average mean motion and average semimajor axis are on top of one another.

As the simulated deployment is in the radial direction, there is no change to the semimajor axis, and so no difference between the average mean motion and average semimajor axis methods. In all cases, the range difference begins by increasing

slowly, until the satellites near peak elevation (around 86°) with respect to the ground station. Near this peak elevation, the measured azimuthal angle for each satellite is changing at a high rate, resulting in a larger azimuth angle separation. The maximum difference between the centroids is about 24 cm. Over the same pass, the true separation between satellites is about 1.6 km. In general, the error seems to grow with increasing angular separation (see Figure 4.10). While the relative state centroid exhibits the same pattern, it actually performs slightly better over most of the pass.

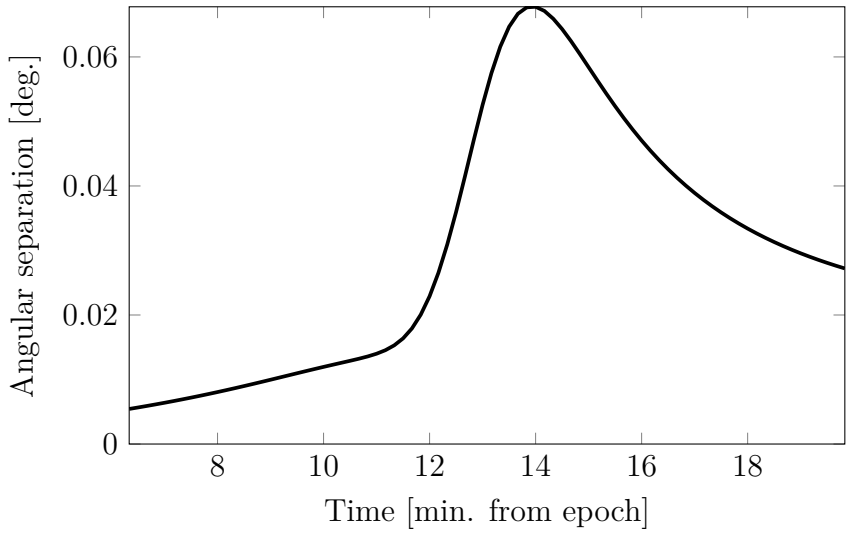

Figure 4.10: Angular separation between the two satellites as measured from the ground station.

A deployment in the in-track direction, on the other hand, will yield the largest difference between the two element space centroid techniques. A 1 m/s in-track deployment is simulated, and the position difference is plotted in Figure 4.11. As expected, there is a pronounced difference between the techniques, and the difference grows with time. The difference from the relative state centroid is the lowest at the beginning of the pass, but as the satellites drift apart the average mean motion

method outperforms it. A similar dependence on angular separation is seen when

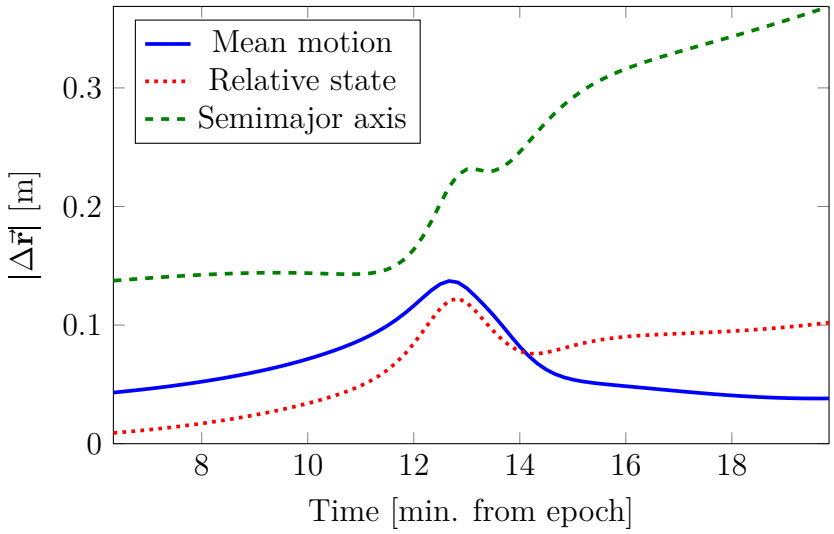

Figure 4.11: Magnitude of the difference between the position vector derived from centroided measurements and those derived from the state-space centroids for the in-track deployment scenario.

averaging in mean motion.

In the next test case, 1000 deployment vectors are drawn from a uniform distribution on the unit sphere (see Section 3.4.2), while the magnitude of the initial separation remains constant at 1 m/s. The same centroiding techniques described above are used, and the maximum range error (between measurement-space and element-space centroids) is mapped to color on the unit sphere in Figures 4.12, 4.13, and 4.15. As in the previous sections, all simulations were again performed using a two-body propagator.

First, we look at the error resulting from averaging the semimajor axis of the two satellites at epoch (Figure 4.12). The reference axes in the figure correspond to the *RSW* frame. The maximum position difference (about 29 cm) occurs when the deployment vector is roughly along the $\hat{S}$ axis, or along-track. This corresponds to a

greater increase in semimajor axis, and a higher difference between the two centroid methods. Referring back to Figure 4.4, we can see that averaging semimajor axis leads to a secular drift (in-track), away from the measurement-space centroid.

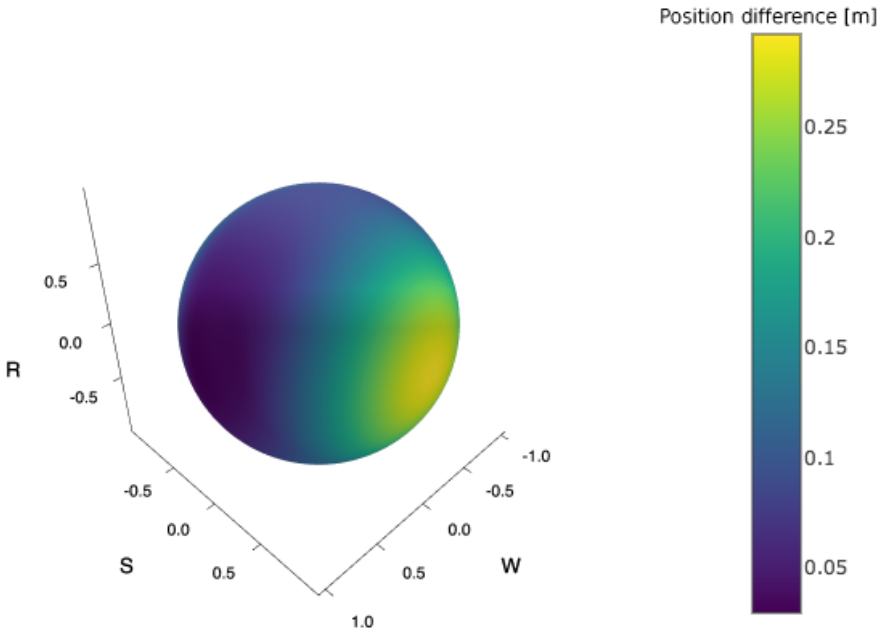

Figure 4.12: Maximum position difference between measurement-space and state-space (using semimajor axis) centroids for a randomly-chosen deployment vector.

Figure 4.13 shows the result of averaging the mean motion of the two satellites at epoch. The maximum position error here (8.93 cm) is much lower than that in Figure 4.12. Here, error is highest when the deployment vector is along the $\hat{R}$ axis. To understand this, first note that radial deployment vectors result in the smallest change in semimajor axis. Thus, the difference between the two centroid methods is smallest for these directions. The error in this case, then, is simply driven by the angular separation (as shown in Figure 4.10). Figure 4.14 illustrates this for the randomly sampled deployment vector case. Larger angular separations over the pass occur near the $\hat{R}$ axis, corresponding to the higher error at these deployment

directions.

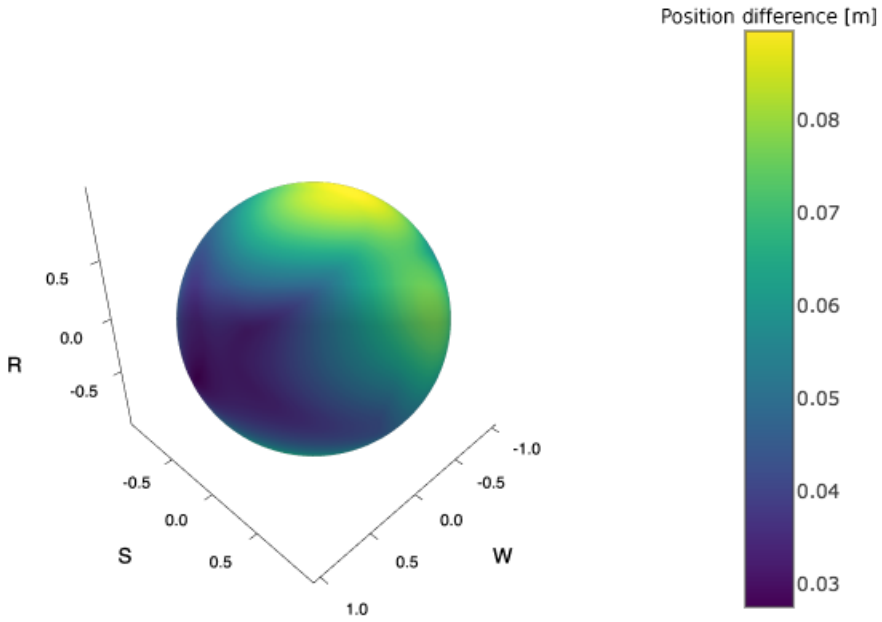

Figure 4.13: Maximum position difference between measurement-space and state-space (using mean motion) centroids for a randomly-chosen deployment vector.

Finally, Figure 4.15 shows the result of averaging in relative state-space at epoch. The magnitude of the position difference is similar to the average mean motion centroid, with a maximum difference of 8.85 cm. However, the deployment vector at which this difference is maximal is at about 45° off-zenith in the orbit plane. Referring to Figure 4.14, this does not quite correspond to peak angular separation. Plotting the separation distance between the two satellites in Figure 4.16, the same trend as in Figure 4.15 is evident. This indicates the error when using relative state-space centroiding is driven by the Euclidean distance between satellites. Note that this distance also affects the validity of the assumptions that go into using the solutions to Hill's equations for propagation.

As an interesting side note, the results presented in Figures 4.14 and 4.16

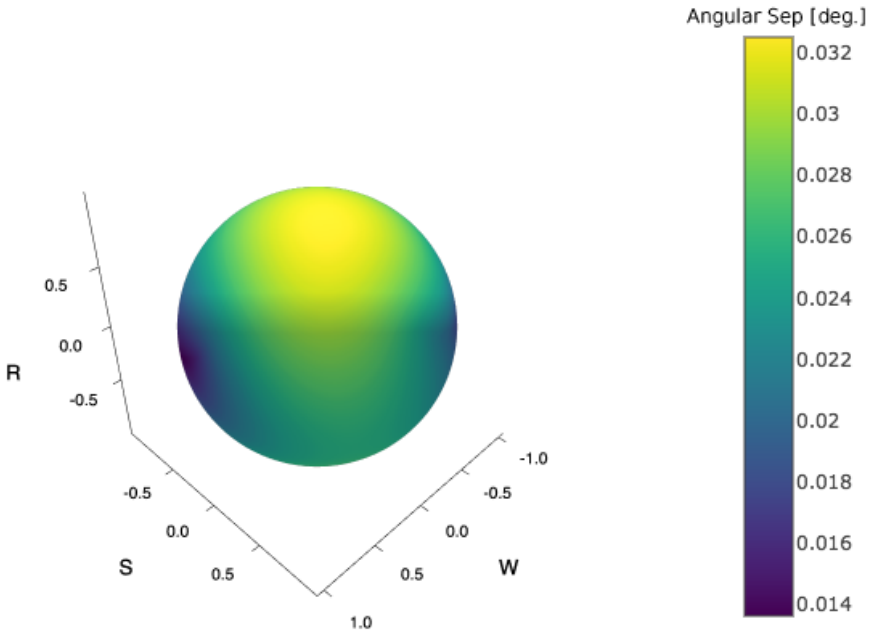

Figure 4.14: Maximum angular separation between satellites for a randomly-chosen deployment vector.

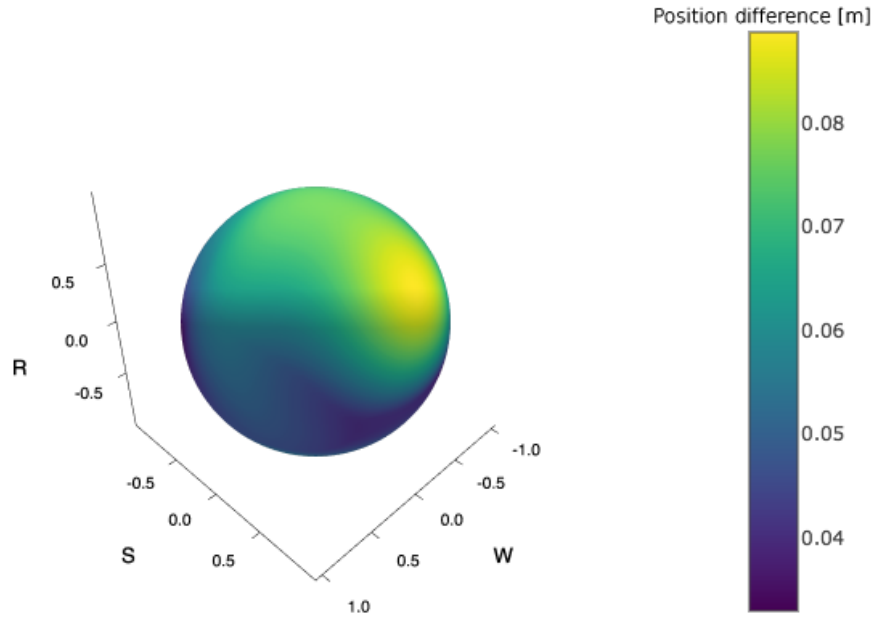

Figure 4.15: Position difference between measurement-space and relative state-space centroids for a randomly-chosen deployment vector.

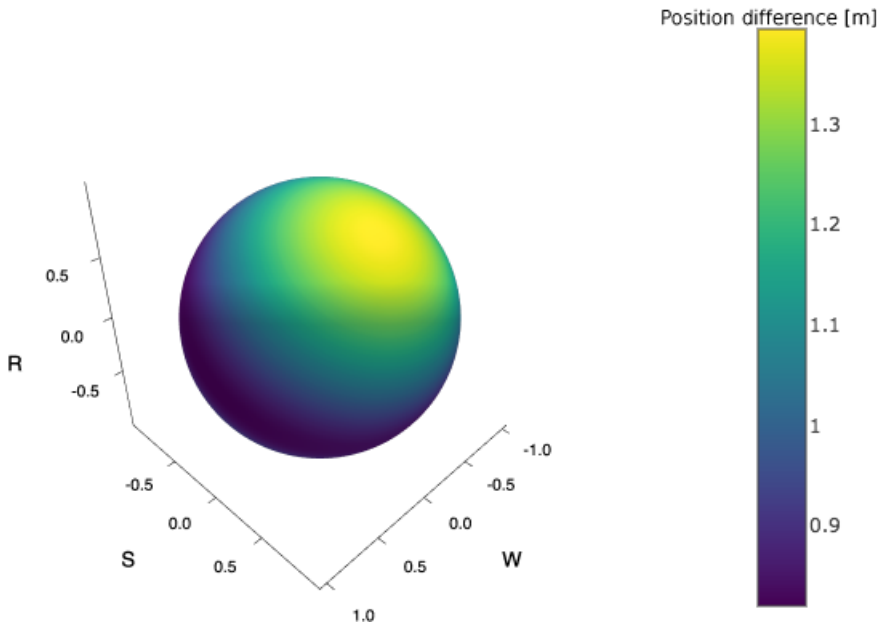

Figure 4.16: Maximum distance between the two satellites for a randomly-chosen deployment vector.

indicate preferred deployment vectors for mission designers who wish to induce maximum separation during a satellite deployment. Angular separation is likely most desirable from the standpoint of discrimination and identification, although of course the trajectory resulting in the maximum angular separation immediately post-deployment (along the $\hat{R}$ axis) will return to a point near the original deployment one orbit later. The preferred direction in Figure 4.16 changes the energy of the orbit, ensuring a secular drift along-track and still resulting in a modest angular separation immediately post-deployment. See Atchison and Rogers [7] for further discussion of this.

These simulations confirm the findings in the first section, and further quantify the advantage of using the average mean motion technique. The remainder of this chapter expands on the effects of the cluster makeup itself, including size and spread.

4.3 Effect of Cluster Size

Here, we examine the effect of cluster size (number of members) on centroid error. A low-discrepancy sequence[5] of 1000 unit vectors is generated according to the technique in Brauchart, Dick, and Fang [155], summarized here. First, a three-dimensional Sobol sequence is generated (using a *Julia* implementation[6]), resulting in a set of uniformly-distributed points in the three-dimensional unit cube. Then, each point $\boldsymbol{x}$ is transformed to the surface of a sphere by:

$$\boldsymbol{x}_{\text{sphere}} = \begin{bmatrix} \sqrt{1-(1-2x_2)^2}\cos(2\pi x_1) \\ \sqrt{1-(1-2x_2)^2}\sin(2\pi x_1) \\ 1-2x_2 \end{bmatrix}. \tag{4.16}$$

The result is a low-discrepancy sequence on the surface of the unit sphere [155]. The key property of a low-discrepancy sequence is that all smaller subsets of the sequence have the same underlying distribution as larger subsets. This allows the number of objects to scale without affecting the characteristics of the distribution, thus minimizing the effects of skewed sets of samples (particularly when the number of samples is small) on the overall result. Figure 4.17 shows the first 200 points from the two-dimensional Sobol sequence and 200 draws from a pseudorandom number generator (PRNG).

The samples are centroided and compared with observation-space centroids, as in the previous section. The ΔV magnitude for every sample is 100 m/s. Figure 4.18

[5] Here, the Sobol sequence is used, see Sobol [154].
[6] `https://github.com/stevengj/Sobol.jl`, version 0.2.0.

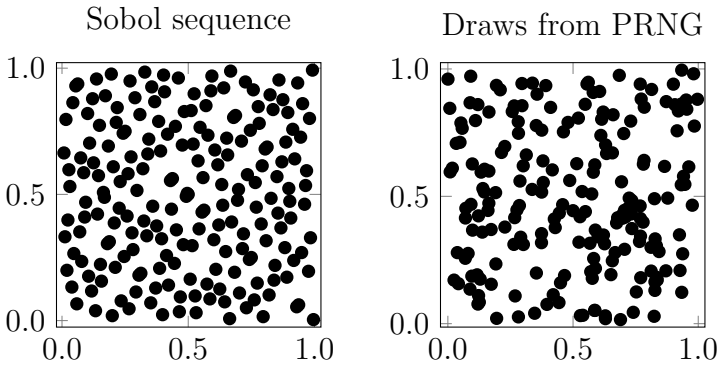

Figure 4.17: Comparison of a two-dimensional Sobol sequence with a sequence drawn from a pseudorandom number generator.

shows the effect of cluster size on maximum centroid error. When the cluster size grows larger than about 200 members, the centroid error converges to a steady state value.

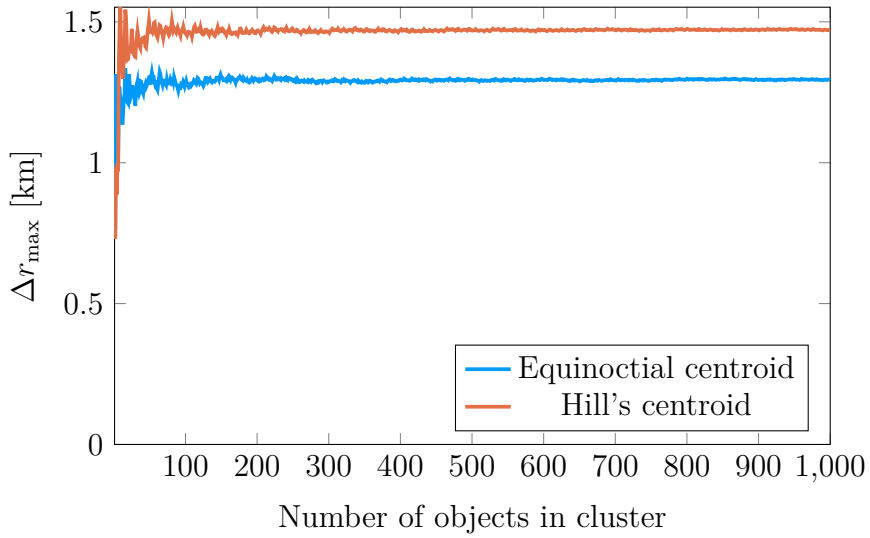

Figure 4.18: Maximum centroid error as a function of cluster size.

The equinoctial state space centroid error converges to 1.3 km, while the relative state space centroid error converges to 1.5 km. This is consistent with the results shown earlier: the relative state space centroid remains reasonably accurate only when the underlying assumptions of small relative distances hold. For the large

ΔV magnitude used in this simulation, that assumption breaks down more quickly, resulting in a higher error than the equinoctial state space centroid. The trend of lower variation in the maximum error with larger sample sets makes intuitive sense: the effect of any single sample is diminished in a larger set. The uniform spatial sampling results in a clearly-defined steady-state maximum value.

4.4 Effect of Separation Velocity

Next, we investigate the effects of the magnitude of the separation velocity. This will result in a larger spread among cluster members. 100 samples are generated with random separation vectors. Then, the ΔV magnitude is varied from 1 m/s to 100 m/s. The same process for centroiding observations and propagating centroids in the previous section is used here.

Figure 4.19 shows the maximum separation distance (over the ground station pass) between the cluster members and the primary satellite, which has the same starting state as the cluster members but undergoes no velocity change. Over the length of the pass (about 16 minutes), the farthest cluster member is about 120 km away from the primary for the highest ΔV case.

Figure 4.20 shows how the error between the different centroid methods and the centroided observation changes as a function of ΔV magnitude. Both centroid methods exhibit similar growth behavior, with the relative centroid growing slightly faster with increasing ΔV. This is a result of the larger separation distance shown in Figure 4.19.

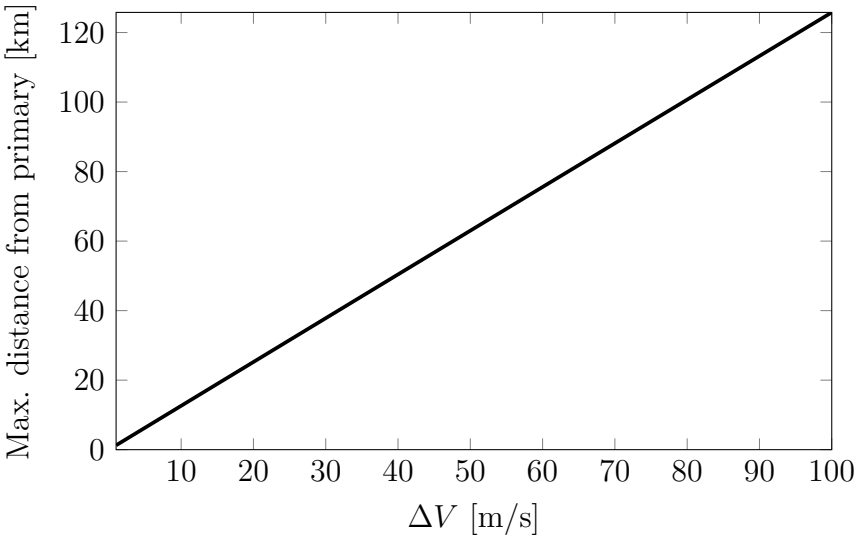

Figure 4.19: Maximum separation distance between the primary and cluster members as a function of ΔV magnitude.

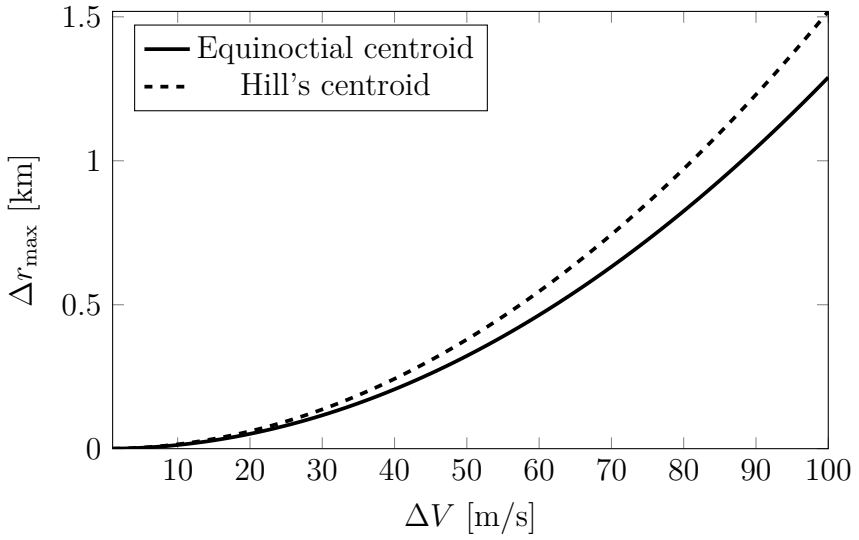

Figure 4.20: Maximum centroid error as a function of ΔV magnitude.

The maximum centroid error for the equinoctial centroid is about 1.3 km, which is only about 1% of the maximum separation distance, while the relative state space centroid performs slightly worse, with about 1.5 km maximum error. Figure 4.21 shows normalized error (defined as the ratio of centroid error to maximum separation distance) as a function of ΔV magnitude.

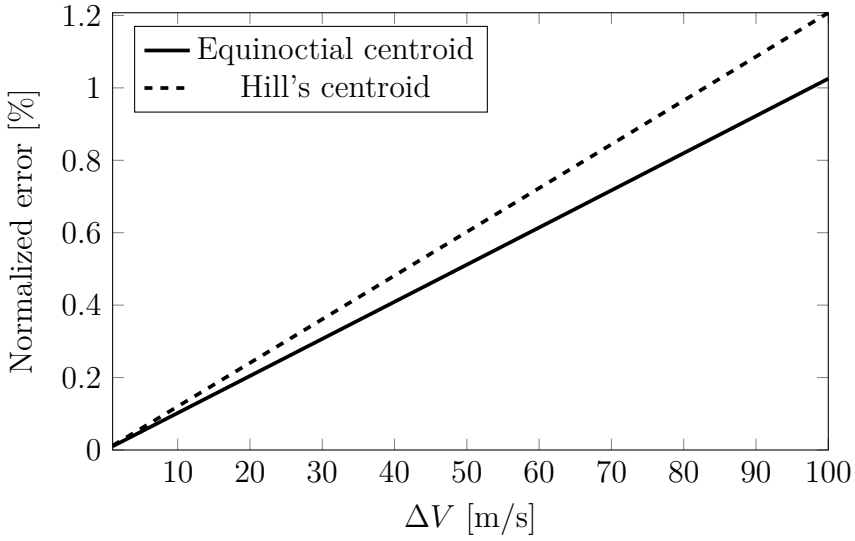

Figure 4.21: Centroid error normalized with respect to maximum separation distance, as a function of ΔV magnitude.

4.5 Summary

This chapter outlined four different ways of computing the centroid of a set of satellites: using inertial Cartesian coordinates, using Keplerian elements, using equinoctial elements, and using a set of Cartesian coordinates relative to one of the satellites. Through basic two-dimensional examples it was shown that inertial Cartesian coordinates are unsuitable for this problem, and that the choice between using Keplerian elements and equinoctial elements is non-obvious. A simulation in which a ground

observer takes measurements of a satellite cluster showed that using the equinoctial element set introduces the least error in the centroiding process, when compared with centroided ground observations.

These comparisons show that a centroid with an average mean motion matches centroided measurements more closely than one with an average semimajor axis, which provides a good basis for using the equinoctial element set. However, another key point is that the difference between the state-space and measurement space centroids is not negligible, even over relatively short time periods. This is important to keep in mind when the measurement centroid is used to update the state-space centroid during estimation–essentially, there is error inherent to this process that limits the attainable accuracy of the cluster centroid. During centroid estimation, this will be compensated for through the use of process noise. Although process noise is meant to simulate noise in the state evolution, it is often used to account for known or unknown deficiencies in the dynamics model, which is precisely the case here.

Part of the cluster estimation technique developed in this work involves centroid estimation, and although existing filtering methods may be used, it is important to ensure that the parameterization itself results in a low deviation from orbital motion. While the relative state space centroid performs reasonably well, it is typically outperformed by the equinoctial state space centroid technique and has more restrictive assumptions, and so there is no compelling reason for using it. As such, the rest of this work will use the equinoctial state space centroid technique.

Chapter 5: Extent Estimation

The main contribution of this dissertation, a method for estimating the extent of a satellite cluster, is developed in this chapter.[1] First, a simplified problem is presented in order to explore the unique aspects of the extent estimation problem outside of the astrodynamics context. Then, observation preprocessing, including gating and forming extent measurements, is discussed. Following that, the extent particle filter is developed. A key aspect of this filter, the likelihood function, must be carefully chosen, and a study of several options is presented. A series of simulations is used to evaluate these options, and a discussion of the consistency of these measures is given at the end of the chapter.

5.1 Simplified Problem

To motivate the rest of this chapter, we first consider a simple problem that incorporates many of the essential elements of the extent estimation problem. While the problem is not worked through completely, it is intended to orient the reader to the extent estimation problem while avoiding some of the complicating, domain-specific details.

[1] Portions of this chapter have been published previously by the author in Binz and Healy [120], [143], [156].

Consider a simple system (e.g. a spring-mass system) whose dynamics are described by the differential equation

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}), \tag{5.1}$$

where $\boldsymbol{x} = (x, \dot{x})$ is the full state of the mass (position and velocity). Consider that the initial state is uncertain, i.e. it is described by a probability density function $\boldsymbol{x}_0 \sim p(\boldsymbol{x})$. Next, consider N samples drawn from $p(\boldsymbol{x})$. Each sample $\boldsymbol{x}_i$ propagates according to the same dynamics function: $\dot{\boldsymbol{x}}_i = f(\boldsymbol{x}_i)$.

In order to approximate the evolution of the probability density function in $\boldsymbol{x}$ over time, we estimate the collection of samples, parameterized by the mean $\bar{\boldsymbol{x}} = \frac{1}{N}\sum_i \boldsymbol{x}_i$ and covariance $P = \frac{1}{N-1}\sum_i^N (\boldsymbol{x}_i - \bar{\boldsymbol{x}})(\boldsymbol{x}_i - \bar{\boldsymbol{x}})^T$ of the set.

The mean state does not evolve according to the same dynamics function f but instead a different, unknown function which depends on each individual sample, i.e. $\dot{\bar{\boldsymbol{x}}} = g(\boldsymbol{x}_1, \boldsymbol{x}_2, ...\boldsymbol{x}_N)$ (see Chapter 4). We may propagate the covariance using the similarity transform, i.e. $P_k = FP_{k-1}F^T$, where P_k is the covariance at the kth time step and F is the discretized dynamics matrix.

Assume observations are made of the mean and standard deviation of position only. Thus each observation is dependent on all samples, i.e. $\boldsymbol{y} = h(\boldsymbol{x}_1, \boldsymbol{x}_2, ...\boldsymbol{x}_N)$. Updates to the mean and covariance are subject to the additional constraint that P be a positive semidefinite matrix.

Importantly, the individual samples are not being estimated–only the mean and covariance are. Additionally, the noise characteristics of the measurements

are unknown, meaning the measurement likelihood function is also unknown. The satellite extent estimation problem has several further complicating factors–e.g., instead of covariance being estimated, the extent, or bounding surface, is.

5.2 Measurement Processing

5.2.1 Measurement Gating

In a noisy, cluttered environment, some observations collected in a given frame might have a source that is different than the target of interest, or might simply be extraneous (i.e. sensor noise). As a result, one of the first processing steps for target tracking is *gating*. Gating refers to the process of choosing the observation (or set of observations) to consider when attempting to associate observations to targets (it does *not* refer to the association itself). In the cluster tracking problem, this means choosing the set of observations to consider as originating from members of the target cluster.

This work uses the statistical gating process first described by Drummond, Blackman, and Pretrisor [75]. First, an ellipsoidal gate S_G is constructed by combining—in measurement space–the predicted extent at time t_k (denoted $\boldsymbol{X}_k^-$), the expected measurement covariance of a single given observation R (typically derived from the properties of the sensor), and the covariance of the centroid state estimate P:

$$S_G = H\boldsymbol{X}_k^- H^T + R + HPH^T, \tag{5.2}$$

where H is the partial derivatives matrix for the transformation from equinoctial elements to measurement space. Typically, the size of the gate S_G is dominated by the extent contribution $H\boldsymbol{X}_k^- H^T$, although this could change with very tight clusters or very noisy sensors.

Once S_G is computed, the Mahalanobis distance [146] between each observation $\boldsymbol{y}_i$ and the predicted centroid observation $\boldsymbol{y}^-$ is calculated and compared to the chosen gate size g:

$$\sqrt{(\boldsymbol{y}_i - \boldsymbol{y}^-)^T S_G^{-1} (\boldsymbol{y}_i - \boldsymbol{y}^-)} \leqslant g, \tag{5.3}$$

Observations that do not satisfy Equation (5.3) are rejected, and not considered further. The value chosen for g depends on the problem at hand, but is typically near one in this work. A larger value for g has the effect of inflating the acceptance criteria.

5.2.2 Extent Measurements

Given a set of N gated observations of a cluster of satellites $\{\boldsymbol{y}\}$, we must compute a measurement of both the centroid and the extent of the cluster. The centroid calculation is discussed at length in Chapter 4, though a measure of its uncertainty is also required. For any individual observation, this would be described by the measurement covariance matrix R; however this does not hold for the derived centroid measurement. Instead, the sample covariance scaled by the number of elements in the cluster, as discussed in Drummond, Blackman, and Pretrisor [75], is found to yield good results when used in the filter. An alternative measurement covariance

is proposed in Section 5.3.7.

Section 3.4.1 discusses the various methods for computing the extent from a set of points. For extent in observation space, the scaled covariance method is used for the reasons described previously, namely that it is less sensitive than the MVEE to the inclusion (or exclusion) of individual points. To reiterate, the extent is computed as the unbiased sample covariance (Equation (3.21)) scaled by the maximum Mahalanobis distance between each point measurement and the sample mean and covariance of the set (Equations (3.22) and (3.23)).

5.3 Extent Particle Filter

In order to introduce the extent particle filter, a brief summary of Bayesian estimation is given first, which motivates the explanation of the particle filter. The unscented Kalman filter, used for centroid estimation, is described later in this section.

Bayesian estimation involves updating a prior belief based on observed data, and as such, target tracking is a natural application. The Bayesian tracking and estimation problem can be expressed generally in a set of two equations. Typically, these involve complicated, potentially intractable integrals, which is one reason that so much work goes into formulating alternative approximate solutions to this problem.

We begin with Bayes' theorem[144, p. 684]. If $\boldsymbol{A}$ and $\boldsymbol{B}$ are generic "events",

then Bayes theorem is written as:

$$p(\boldsymbol{A}|\boldsymbol{B}) = \frac{p(\boldsymbol{B}|\boldsymbol{A})p(\boldsymbol{A})}{p(\boldsymbol{B})}, \tag{5.4}$$

where

$p(\boldsymbol{A}|\boldsymbol{B})$ is the "posterior" probability of $\boldsymbol{A}$, given $\boldsymbol{B}$ occurred.

$p(\boldsymbol{B}|\boldsymbol{A})$ is the "likelihood" of $\boldsymbol{B}$, given $\boldsymbol{A}$ occurred.

$p(\boldsymbol{A})$ is the "prior probability" of $\boldsymbol{A}$ occurring, or simply the "prior".

Next, we introduce some target tracking-specific terminology. Start with the following discretized dynamic system:

$$\begin{aligned} \boldsymbol{x}_k &= f(\boldsymbol{x}_{k-1}, t_k) + \boldsymbol{q}_{k-1} \\ \boldsymbol{y}_k &= h(\boldsymbol{x}_k) + \boldsymbol{r}_k, \end{aligned} \tag{5.5}$$

where $\boldsymbol{x}$ is the target state vector, $\boldsymbol{y}$ is an observation vector, k denotes the kth time step, $f(\cdot)$ is the dynamics function (which maps the state forward in time), $h(\cdot)$ is the measurement function (which maps state space to measurement space), $\boldsymbol{q}$ and $\boldsymbol{r}$ are process and measurement noise characteristics, respectively.

The objective of the target tracking problem is to get the optimal estimate of the target state at t_k, given an estimate of the state at the previous time and an observation of the target. Because no physical quantity may be known or measured without some degree of uncertainty, each of these quantities may be represented by a *probability density function* (PDF). So, the objective is to optimally estimate

$p(\boldsymbol{x}_k)$. This may be written in a Bayesian sense as

$$p(\boldsymbol{x}_k|\boldsymbol{y}_{1:k}) = \frac{p(\boldsymbol{y}_k|\boldsymbol{x}_k)p(\boldsymbol{x}_k|\boldsymbol{y}_{1:k-1})}{p(\boldsymbol{y}_k|\boldsymbol{y}_{1:k-1})}. \tag{5.6}$$

The prior $p(\boldsymbol{x}_k|\boldsymbol{y}_{1:k-1})$ is given by the Chapman-Kolmogorov equation[157, p. 30]:

$$p(\boldsymbol{x}_k|\boldsymbol{y}_{1:k-1}) = \int p(\boldsymbol{x}_k|\boldsymbol{x}_{k-1})p(\boldsymbol{x}_{k-1}|\boldsymbol{y}_{1:k-1})d\boldsymbol{x}_{k-1}. \tag{5.7}$$

The total probability [158]

$$p(\boldsymbol{y}_{1:k-1}) = \int p(\boldsymbol{y}_k|\boldsymbol{x}_k)p(\boldsymbol{x}_k|\boldsymbol{y}_{1:k-1})d\boldsymbol{x}_k \tag{5.8}$$

is a normalizing term. Equation (5.7) is often intractable, though there are a number of methods to approximate it, including a numerical approximation via Monte Carlo methods.

The extent estimation component (recall that centroid and extent estimation are decoupled) developed in this work is based on the particle filtering method. Particle filtering is a Monte Carlo technique for approximating intractable integrals that may appear in the tracking equations (see Arulampalam *et al.* [158] for a good overview of particle filters). Given some prior distribution in state space, a large number of samples, also called particles, are randomly generated. Each of these samples can be propagated through the dynamics function, and converted to measurement space via the measurement function. Then, each sample is compared with the observed measurement using a measurement likelihood function. This

comparison is used to weight each sample such that samples that closely match the observation are weighted more heavily than those that do not. The resulting set of weighted particles can then be used to approximate the posterior PDF.

Revisiting Equation (5.6), a particle filter approximates the posterior distribution as the weighted sum of a set of particles [158]:

$$p(\boldsymbol{x}_k|\boldsymbol{y}_k) \approx \sum_{i=1}^{N} w_k^i \delta(\boldsymbol{x}_k - \boldsymbol{x}_k^i). \tag{5.9}$$

In a *sampling importance resampling* (SIR) particle filter, the sampling of the posterior occurs at every time step. As such, the weights may be chosen to be [158]:

$$w_{k+1}^i \propto p(\boldsymbol{y}_{k+1}|\boldsymbol{x}_{k+1}^i), \tag{5.10}$$

that is, the weight of a particle is proportional only to the likelihood of the particle.

For the present problem, the steps may be summarized as:

1. Sample the prior distribution $\boldsymbol{X}_k^i \sim p(\boldsymbol{X}_k)$.
2. Propagate the samples: $\boldsymbol{X}_{k+1}^{i-} = f(\boldsymbol{X}_k^i)$.
3. Compute the expected measurement for each sample: $\boldsymbol{Y}_{k+1}^i = h(\boldsymbol{X}_{k+1}^{i-})$.
4. Compute the weights for each sample using some measure of the likelihood $w_i \propto p(\boldsymbol{Y}_{k+1}|\boldsymbol{X}_{k+1}^{i-})$ and normalize.
5. Compute the mean of the posterior distribution: $\boldsymbol{X}_{k+1} \approx \sum_{i=1}^{N} w_i \boldsymbol{X}_{k+1}^{i-}$.

Once the posterior is computed, the process is repeated for the next time step.

5.3.1 Unscented Kalman Filter

Although any nonlinear sequential estimator may be used for the centroid estimation component, this work uses the unscented Kalman filter (UKF). Details of the UKF can by found in any of a number of sources (see Julier and Uhlmann [159] and Särkkä [160] for two particularly good examples), but essentially it relies on a set of deterministically-chosen weighted samples, called sigma points, to characterize the first two moments of the PDF of the target state. The sigma points are propagated individually through the dynamics function for the filter's prediction step, and the same process is used in the update step to represent the measurement PDF. Särkkä presents a compact version of the UKF equations, shown below [161]. The process for choosing the set of sigma points is as follows.

Given a state and covariance at some time $(\boldsymbol{x}, P)$, the n-dimensional multivariate Gaussian PDF may be represented by the set of sigma points $\{\boldsymbol{\mathcal{X}}\}$. The $n \times n{+}1$ matrix $[\boldsymbol{\mathcal{X}}]$, whose columns comprise the sigma points, is given by

$$[\boldsymbol{\mathcal{X}}] = \begin{bmatrix} \boldsymbol{x} & \cdots & \boldsymbol{x} \end{bmatrix} + \sqrt{n+\lambda} \begin{bmatrix} \mathbf{0} & \sqrt{P} & -\sqrt{P} \end{bmatrix}, \tag{5.11}$$

where $\mathbf{0}$ represents an $n \times 1$ zero vector and $\sqrt{P}$ denotes the Cholesky decomposition of the covariance matrix P[141, p. 36]. The constant $\lambda = \alpha^2(n + \kappa) - n$ is a function of the dimension of the state space ($n = 6$) and two positive constants α and κ, which are parameters that affect the spread of the sigma points and allow for additional tuning based on prior knowledge of the underlying distribution of $\boldsymbol{x}$,

respectively (in this work, $\alpha = 0.75$ and $\kappa = 3$ are used) [32], [162], [163]. The weights are calculated as

$$\begin{aligned} w_0^{(m)} &= \tfrac{\lambda}{n+\lambda} \\ w_0^{(c)} &= \tfrac{\lambda}{n+\lambda} + 1 - \alpha^2 + \beta \\ w_i^{(c,m)} &= \tfrac{1}{2(n+\lambda)}, \qquad\qquad i = 1, \ldots, 2n. \end{aligned} \tag{5.12}$$

β may be introduced to incorporate prior information about the distribution of $\boldsymbol{x}$ [163]. For the present problem, the value $\beta = .5$ is chosen heuristically–it results in a converged estimate in all of the test cases, while larger values tend to diverge in one or more of the cases. A more systematic way of tuning the free parameters α, β, and κ is a potential subject for future investigation.

The prediction and update steps are given below.

Prediction:

$$[\hat{\boldsymbol{\mathcal{X}}}]_k = f([\boldsymbol{\mathcal{X}}]_{k-1}, k-1) \tag{5.13a}$$

$$\boldsymbol{x}_k^- = [\hat{\boldsymbol{\mathcal{X}}}]_k \boldsymbol{w} \tag{5.13b}$$

$$P_k^- = [\hat{\boldsymbol{\mathcal{X}}}]_k \boldsymbol{W}_c [\hat{\boldsymbol{\mathcal{X}}}]_k^T + Q_{k-1} \tag{5.13c}$$

Update:

$$[\boldsymbol{\mathcal{Y}}]_k^- = h([\boldsymbol{\mathcal{X}}]_k^-, k) \tag{5.13d}$$

$$\boldsymbol{\mu}_k = [\boldsymbol{\mathcal{Y}}]_k^- \boldsymbol{w} \tag{5.13e}$$

$$S_k = [\boldsymbol{\mathcal{Y}}]_k^- W_c [\boldsymbol{\mathcal{Y}}]_k^{-T} + R_k \tag{5.13f}$$

$$C_k = [\boldsymbol{\mathcal{X}}]_k^- W_c [\boldsymbol{\mathcal{Y}}]_k^{-T} \tag{5.13g}$$

$$K_k = C_k S_k^{-1} \tag{5.13h}$$

$$\boldsymbol{x}_k = \boldsymbol{x}_k^- + K_k \left[\boldsymbol{y}_k - \boldsymbol{\mu}_k\right] \tag{5.13i}$$

$$P_k = P_k^- - K_k S_k K_k^T, \tag{5.13j}$$

$[\boldsymbol{\mathcal{X}}]$ and $[\boldsymbol{\mathcal{Y}}]$ are matrices whose columns are the state-space and measurement space sigma points, respectively. A "hat" (e.g. $[\hat{\boldsymbol{\mathcal{X}}}]$) indicates a predicted value, $f()$ is the dynamics function, $h()$ is the measurement function, P is the state covariance, and R is the measurement covariance. The k subscripts represent the time step. The $2n+1 \times 2n+1$ vector $\boldsymbol{w}$ is the weight vector (Equation (5.12)), and the matrix W_c is constructed as

$$W_c = (\mathbb{1}_{2n+1\times 2n+1} - [\boldsymbol{w}\cdots\boldsymbol{w}]) \tag{5.14}$$
$$\times \operatorname{diag}(w_0\cdots w_{2n})$$
$$\times (\mathbb{1}_{2n+1\times 2n+1} - [\boldsymbol{w}\cdots\boldsymbol{w}])^T .$$

5.3.2 Prior Distribution

Following Koch [105], we chose the Inverse Wishart distribution for the prior distribution of extent matrices. It is often used as the conjugate prior for unknown covariance matrices of multivariate normal distributions (meaning the distribution of these unknown matrices is itself Inverse Wishart) [164, p. 111], and has been used to good effect in the group tracking literature [105], [107].

The probability density function of an Inverse Wishart-distributed random matrix $Y \sim \mathcal{IW}(\nu, \Psi)$ is [165, p. 272]

$$f(Y) = \frac{|\Psi|^{\frac{\nu}{2}}}{2^{\frac{\nu p}{2}}\Gamma_p\left(\frac{\nu}{2}\right)}|Y|^{-\frac{\nu+p+1}{2}}\exp -\frac{1}{2}\operatorname{tr}(\Psi Y^{-1}), \tag{5.15}$$

where

Y is a $p \times p$ positive definite matrix,

Ψ is the scale matrix (also $p \times p$ and positive definite),

$|\cdot|$ denotes the matrix determinant,

ν is the number of degrees of freedom of the distribution, and

Γ_p is the p-variate gamma function, defined as[165, p. 257]

$$\Gamma_p(t) = \pi^{p(p-1)/4} \prod_{i=1}^{p} \Gamma\left[t - \frac{1}{2}(i-1)\right]. \tag{5.16}$$

To illustrate the nature of the Inverse Wishart distribution, Figure 5.1 plots 20 ellipses sampled from the distribution $\mathcal{IW}(4, \mathbb{1}_2)$, where $\mathbb{1}_2$ is a two-dimensional identity matrix. The ellipse specified by the scale matrix $\mathbb{1}_2$ is plotted with a heavy red line. Note that both scale (corresponding to eigenvalue magnitude) and axial orientation (eigenvectors) are varied.

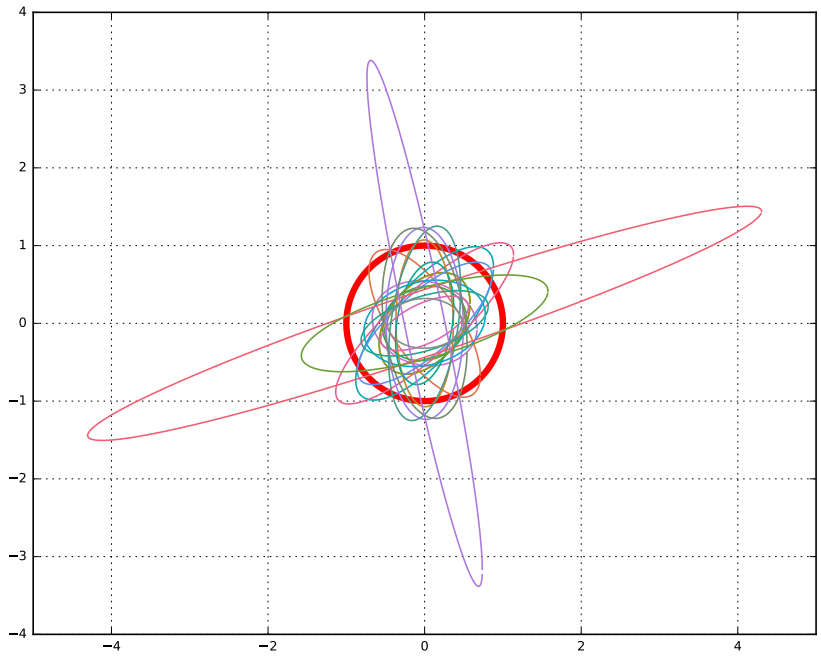

Figure 5.1: Samples from the two-dimensional Inverse Wishart distribution $\mathcal{IW}(4, \mathbb{1}_2)$.

One other important property of the Inverse Wishart distribution that will be used here is the expected value of an Inverse Wishart-distributed random matrix[165,

p. 273]:

$$\mathrm{E}[Y] = \frac{\Psi}{\nu - p - 1}. \tag{5.17}$$

Here, the constraint that ν be greater than $p + 1$ is apparent.

5.3.3 Implementation

Assume we start with an estimate of the cluster centroid, its covariance, and its extent at time t_k, collectively denoted $\{\boldsymbol{x}, P, \boldsymbol{X}\}_k$. If we assume that $\boldsymbol{X} \sim \mathcal{IW}(\nu, \Psi)$, we can compute $\Psi_k = (\nu - p - 1)\boldsymbol{X}_k$. This defines the prior distribution at time t_k, allowing us to draw N random samples in order to approximate the transition density and measurement likelihood. In the SIR framework, resampling occurs at every time step, and so all samples begin with equal weights: $w_i = 1/N$.

Each particle $\boldsymbol{X}_k^i$ is propagated using the linearized state transition matrix F_{k+1} around the centroid solution $\boldsymbol{x}_k$:

$$\boldsymbol{X}_{k+1}^{i-} = F_{k+1}\boldsymbol{X}_k^i F_{k+1}^T + Q_k, \tag{5.18}$$

where Q_k is the process noise covariance matrix (see Appendix D). Note that this is analogous to propagating the centroid covariance via the similarity transform P, and could feasibly be done using, e.g., the unscented transform (see Section 3.3). The result of this is an approximation of the extent distribution at time t_{k+1}, prior to the updating step.

Observations of the extent matrix are projections of the extent matrix into

measurement space involving the nonlinear measurement function $h(\boldsymbol{X})$. This projection may be linearized around the current centroid state, yielding the matrix of partial derivatives H_{k+1} for transforming between state space (equinoctial elements) and measurement space (e.g., range, azimuth, and elevation). Then, each matrix-variate particle $\boldsymbol{X}^i$ may be transformed into measurement space using the similarity transformation (see Section 3.3):

$$\boldsymbol{Y}_{k+1}^i = H_{k+1}\boldsymbol{X}_{k+1}^{i-}H_{k+1}^T. \tag{5.19}$$

While the density of $\boldsymbol{X}_k$ is known *a priori*, the measurement likelihood is not, and so we must choose a function that is proportional to the measurement likelihood that we can evaluate. Denoted $s_i = s(\boldsymbol{Y}_{k+1}, \boldsymbol{Y}_{k+1}^i)$, this function is a measure of similarity between the observed measurement extent $\boldsymbol{Y}_{k+1}$ and the predicted measurement extent of the ith particle $\boldsymbol{Y}_{k+1}^i$ (see Section 5.4 for a discussion of the options). The set of similarities s_i is then used to construct the set of normalized particle weights:

$$w_i = \frac{s_i}{\sum_{j=1}^N s_j}. \tag{5.20}$$

With the weighted set of propagated extent particles, we may now compute an approximation of the *a posteriori* density:

$$\boldsymbol{X}_{k+1} = \sum_{i=1}^N w_i \boldsymbol{X}_{k+1}^{i-}. \tag{5.21}$$

Importantly, because samples of the Inverse Wishart matrix are positive definite,

this sum is guaranteed to be positive definite[141, p. 250]. The centroid state and covariance $(\boldsymbol{x}, P)_k$ may be updated separately (see Section 5.3.1), giving a complete solution at time t_{k+1}: $(\boldsymbol{x}, P, \boldsymbol{X})_{k+1}$. This process is then repeated for the next frame of measurements.

5.3.4 Degrees of Freedom Evolution Model

As discussed in Section 5.3.2, ν is a measure of the degrees of freedom in the Inverse Wishart distribution. A lower value of ν essentially results in a larger "spread" of matrix samples. There is a lower bound of $\nu > p + 1$ (see Equation 5.17), and no mathematical upper bound (though a practical upper bound is introduced in the model below). Figure 5.2 shows several random draws from $\mathcal{IW}(\nu, \mathbb{1}_2)$ for increasing values of ν.

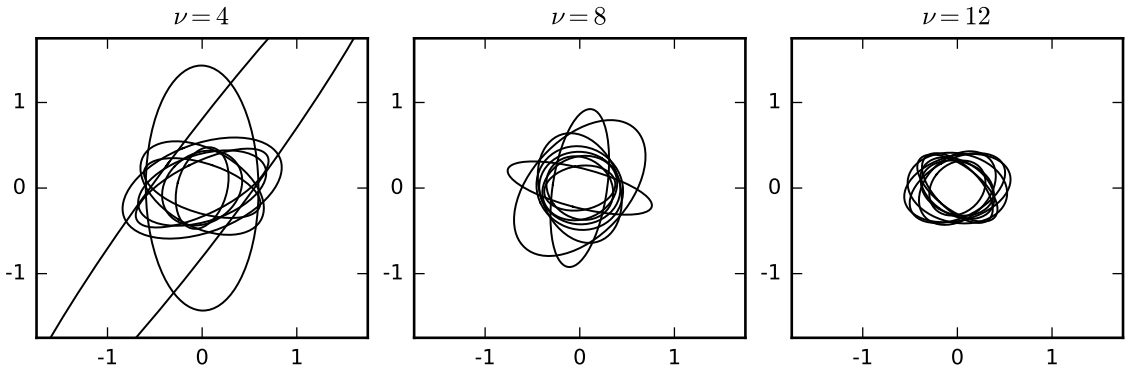

Figure 5.2: Ellipses drawn from $\mathcal{IW}(\nu, \mathbb{1}_2)$ for different values of ν.

During the filtering process, the value of ν should change to reflect the confidence of the current extent estimate. Koch [105] discusses a heuristic approach to this, which we use as a starting point in this work. Essentially, we desire the distribution to generate samples with a larger variance when the time between updates is

long, and each update should reduce that variance, representing an increase in the level of confidence in the extent distribution.

During the prediction step, a longer time since update should result in a decreasing measure of confidence in the extent estimate, corresponding to a decreasing value of ν. Koch suggests an exponential decay model for prediction:

$$\nu_k^- = e^{-\Delta t/\tau}\nu_{k-1}, \tag{5.22}$$

where τ is some time constant, and Δt is the time between the kth update and the previous update $(k-1)$. Here, we use a similar model, with the addition of a "floor", $\nu_{\min}$, to satisfy the constraint that $\nu > p+1$ (see Section 5.3.2):

$$\nu_k^- = e^{-\Delta t/\tau}(\nu_{k-1} - \nu_{\min}) + \nu_{\min}. \tag{5.23}$$

The tuning parameter τ is taken to be equal to one orbit period in this work.

Each successful update step increases our confidence in the estimate of the extent matrix, reflected by an increased value of ν. For the update step of the filter, a linear increase, up to some maximum value $\nu_{\max}$, is used:

$$\nu_k = \begin{cases} \frac{(\nu_{\max}-\nu_{\min})}{\beta} + \nu_k^-, & \text{if } \nu_k^- < \nu_{\max} \\ \nu_{\max}, & \text{otherwise.} \end{cases} \tag{5.24}$$

β is another tuning parameter that controls the rate at which ν grows during updates, and is set to some fraction of an orbit period. While these models allow the

extent matrix distribution to evolve based on the update rate of the filter, in practice they offer little benefit over a constant-ν model. Nevertheless, this dissertation incorporates these evolution models. Figure 5.3 shows the evolution of ν over a typical pass for a LEO satellite.

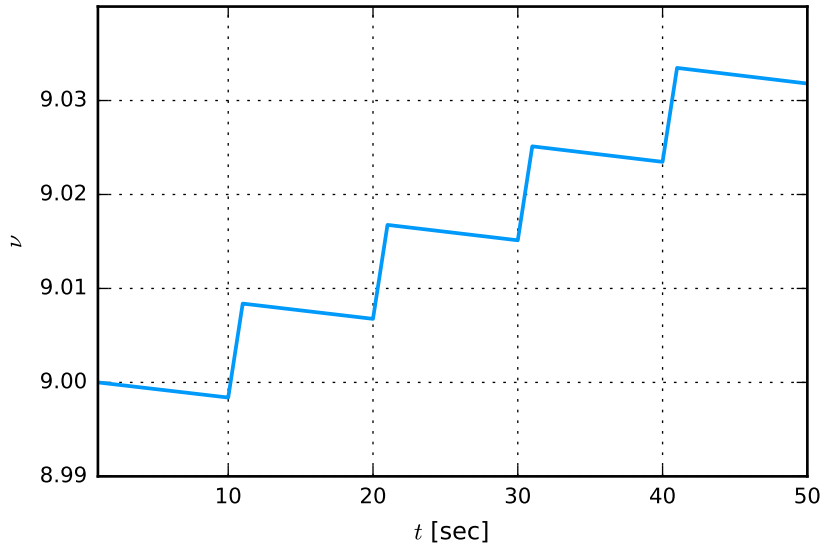

Figure 5.3: Evolution of ν in a scenario with updates every 10 seconds, $\tau = 5400$, and $\beta = 600$.

5.3.5 Initialization

When a breakup or other event is detected, and the need for extent estimation is recognized, the extent filter must be initialized using some *a priori* data. The centroid state and covariance, as well as an initial estimate of the extent matrix, are required. A good initial guess for all of these can be found from the nominal trajectory estimate of the single, pre-breakup object. Using this state and covariance estimate as the initial centroid state and covariance estimate works well, and is the way the simulations in this work are initialized. A propagated state and covariance

estimate is also acceptable–the estimation epoch of the pre-event object does not need to be at the event time.

The initial extent estimate deserves more attention. Ideally, the extent estimation process would begin with a minimally restrictive and maximally uncertain starting point, allowing the filter to converge without, for example, unnecessarily excluding objects from the cluster. Bernardo calls this a *noninmofrmative prior*, which he describes as "a mathematical description of that situation where data would best dominate prior knowledge about the quantity of interest" [166, p. 17]. In Cartesian space, this could be a large hypersphere around the nominal centroid state. Converting this into equinoctial element space results in a reasonable initial extent estimate. One key difference between the initialization of the extent estimate and the state and covariance estimate initialization is that a large extent is not necessarily more uncertain than a small extent. Therefore, the uncertainty in the initial extent estimate (quantified by the parameter ν) is more important than the actual initial extent estimate, although a larger initial extent will allow for more inclusive gating early on.

Another option is to use a multiple of the initial centroid covariance estimate. This encodes some information about the estimated orbit immediately, resulting in, perhaps, a more restrictive initial estimate. In such a case, a very large multiplier, such as $k = 10$ or 20, should be used (as in Section 5.4.3).

Minimizing restrictiveness in equinoctial orbit space requires paying special attention to degenerate or near-degenerate axes of the 6D hyperellipsoid. That is, when the extent matrix is projected in all 15 possible planes, the ratio of the semi-

minor axis to the semi-major axis of the ellipse (or, the determinant of the 2×2 sub-matrix corresponding to the projection) should not be nearly zero[2]. Treating the extent matrix as a covariance matrix, another way of stating this is that the *correlation coefficient* ρ between two variables should not be unity. ρ is defined as [139, p. 74]

$$\rho_{xy} = \frac{\sigma_{xy}}{\sigma_x \sigma_y}, \tag{5.25}$$

where σ_{xy} is the cross-covariance term between two variables x and y, and σ_x is the standard deviation in x. The correlation coefficient between all 15 pairs of variables for the *a priori* extent estimate should be computed and flagged if less than some value $1 + \varepsilon$, where ε is small (e.g. 10^{-4}).

For extent matrices with near-degenerate axes, the following method is used to improve the conditioning of the matrix. Let $\boldsymbol{X}_{ij}$ be the 2×2 sub-matrix for the degenerate (i, j) projection of the extent matrix $\boldsymbol{X}$. $\boldsymbol{X}_{ij}$ can be represented by the eigenvalue decomposition

$$\boldsymbol{X}_{ij} = [\boldsymbol{v}_1 \boldsymbol{v}_2] \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} [\boldsymbol{v}_1 \boldsymbol{v}_2]^{-1}. \tag{5.26}$$

In the degenerate case, the smaller eigenvalue will be several orders of magnitude smaller than the larger eigenvalue. Increasing this by one or more orders of

[2]Note that, in some cases, the actual distribution of cluster states may fall nearly on a line in some projections, as in the (n, a_f) plane in the GEO test case presented in Section 5.4.3.

magnitude ($\lambda_1' = k\lambda_1$), then reconstructing the sub-matrix as

$$\boldsymbol{X}_{ij}' = [\boldsymbol{v}_1 \boldsymbol{v}_2] \begin{bmatrix} \lambda_1' & 0 \\ 0 & \lambda_2 \end{bmatrix} [\boldsymbol{v}_1 \boldsymbol{v}_2]^{-1} \tag{5.27}$$

has the effect of expanding the minor axis of the projected ellipse. λ_1' should be chosen such that the resulting correlation coefficient of $\boldsymbol{X}_{ij}'$ satisfies the lower bound of $1 + \varepsilon$. A recent paper by Hall, et al. discusses other methods for correcting ill-conditioned covariance matrices [167].

5.3.6 Partially Obscured Cluster

As cluster pieces spread apart and the cluster gets larger, there will be passes over ground observation stations during which the entire cluster is not visible at once. This causes an issue with the current method, because each step of the particle filter estimates the entirety of the cluster extent. Further, without observing the entire cluster, we do not have sufficient information to even estimate the centroid position. Several articles have explored this problem, mostly as it pertains to video surveillance and targets that get occluded by some obstruction in the scene [168], [169]. These techniques do not fit our purposes: either they are too domain-specific (scene modeling), or simplistic (assuming symmetric shapes). Further, there is a body of work pertaining to target tracking with missing observations (e.g. Zhang *et al.* [170]), but they assume data is randomly missing due to sensor dropouts, not because of some known obstruction, and so they address a different problem than

ours. In order to update the cluster extent estimate with the information that is available (i.e., with the portion of the extent that is currently visible over the ground station), we borrow and adapt the statistical technique of *imputation* [171].

Imputation is the process of handling missing data in some data set [172], [173]. There are several imputation methods, including techniques as simple as replacing the missing data point with the sample mean of the set. In our case, measurements that are obscured by the Earth or masked by the sensor are imputed using the last extent estimate. This allows us to assume the non-visible portions of the cluster are simply propagated forward from the previous estimate using the cluster evolution model. While not true imputation in the strict sense, it enables an extent estimate using the available observations. The process is described below.

First, the predicted cluster estimate is used to determine whether the entire cluster will be visible at the next time step. If it is not, a sample population of virtual cluster members is generated over the entire extent, using uniform sampling (see Appendix B.2) over the predicted extent matrix. The number of samples is an adjustable parameter, and in this work we used 100. Then, we determine whether there is a line of sight from the ground station to the satellite, i.e.:

$$\arcsin\left(\frac{\rho_z}{\rho}\right) > el_{\min}, \tag{5.28}$$

where $\boldsymbol{\rho}$ is the position vector from the ground station to the satellite in the topocentric horizon coordinate system[126, p. 161], and $el_{\min}$ is the minimum elevation at which the ground station can observe the satellite, set to zero unless there is some

other obscuration (e.g., nearby mountains).

Those virtual satellites that are in the field of view of the ground station are discarded, leaving only the unobserved cluster portion sampled. This sample population is then combined with the set of actual observations, and the observation reduction and filtering process proceeds as normal, with one exception: the centroid estimate is not updated, since the measurement centroid cannot be observed at once. This process is illustrated in Figures 5.4–5.7.

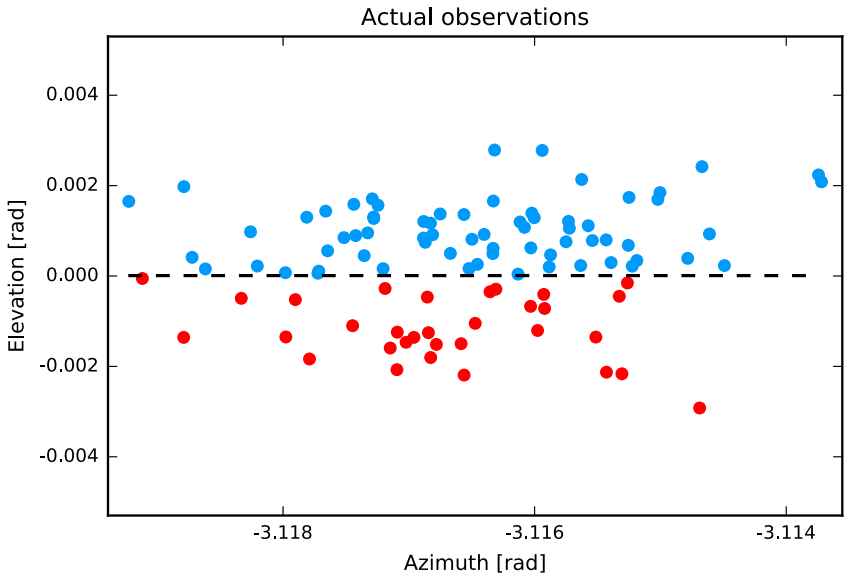

Figure 5.4: Observations of an entire cluster, without taking Earth obscuration into account. Red markers indicate observations that are below the horizon (black dashed line).

5.3.7 Adaptive Measurement Covariance

As discussed above, the scaled measurement extent is used to weight the centroid observations used in the centroid filter. The measurement space extent is sometimes quite sensitive to the exclusion of a single observation by the gating routine. In

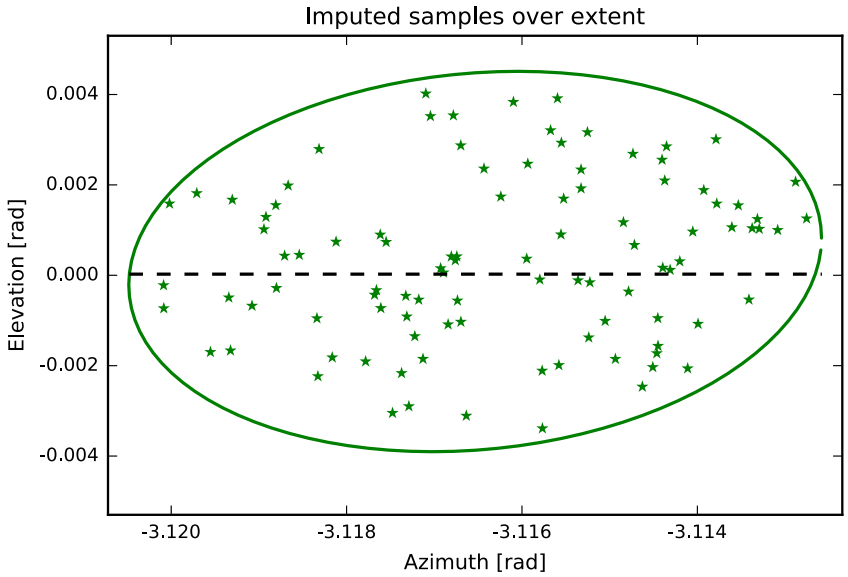

Figure 5.5: The predicted extent is then populated with uniform random samples.

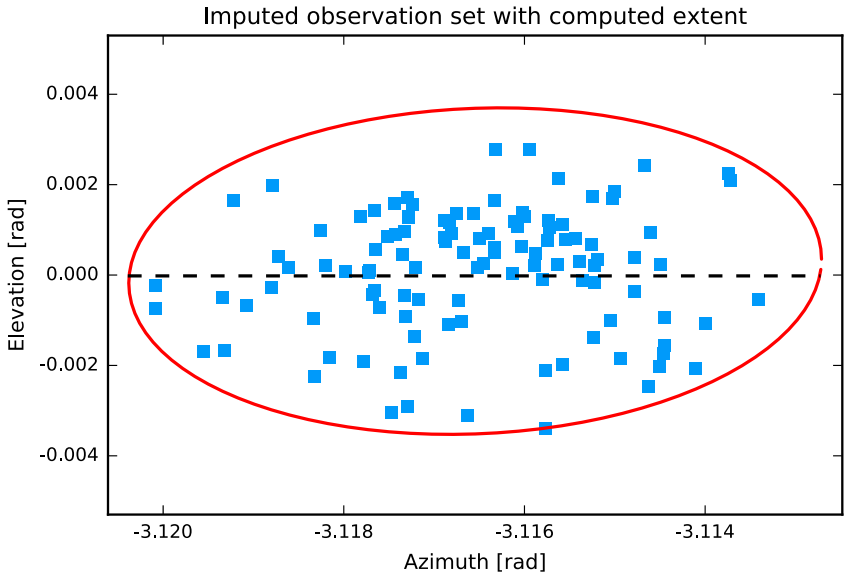

Figure 5.6: The imputed samples in the obscured part of the cluster are then combined with the unobscured observations, and this set is used to compute the measurement extent.

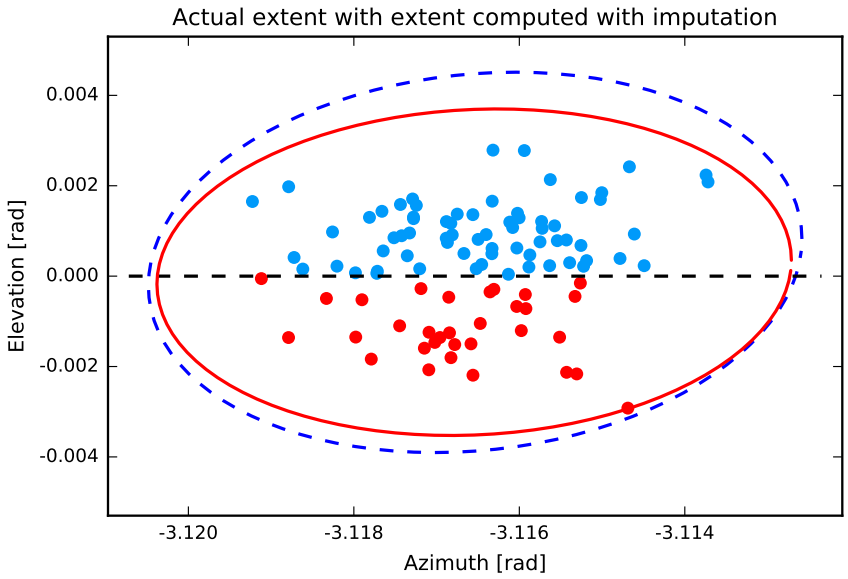

Figure 5.7: The imputed extent (solid red line) plotted with the original cluster observations and the true (unobservable) extent (dashed blue line).

order to account for this when estimating the centroid, an adaptive measurement covariance is proposed here.

First, we compute the difference in the centroid of the set of measurements due to the inclusion (or exclusion) of the most distant measurement. For the set of samples $\boldsymbol{y}$, the centroid is simply

$$\bar{\boldsymbol{y}} = \frac{1}{N}\sum_{i=1}^{N} \boldsymbol{y}^i. \tag{5.29}$$

Now, let the samples be indexed such that $\boldsymbol{y}^N$ is the farthest (using the L^2 distance as a metric) sample from the mean $\bar{\boldsymbol{y}}$, and as such, the most likely to be dropped during the gating process. The centroid for this modified set not including $\boldsymbol{y}^N$ is then

$$\bar{\boldsymbol{y}}_{\text{mod}} = \frac{1}{N-1}\sum_{j=1}^{N-1}\boldsymbol{y}^j. \tag{5.30}$$

Now, we can solve for the sensitivity of the centroid to removal of the Nth measurement:

$$\Delta\bar{\boldsymbol{y}} = |\bar{\boldsymbol{y}} - \bar{\boldsymbol{y}}_{\text{mod}}| = \left|\frac{1}{N}\sum_{i=1}^{N}\boldsymbol{y}^i - \frac{1}{N-1}\sum_{j=1}^{N-1}\boldsymbol{y}^j\right| \tag{5.31}$$

$$= \left|\frac{1}{N}(\boldsymbol{y}^1 + \ldots + \boldsymbol{y}^N) - \frac{1}{N-1}(\boldsymbol{y}^1 + \ldots + \boldsymbol{y}^{N-1})\right| \tag{5.32}$$

$$= \left|\frac{\sum_i^{N-1}\boldsymbol{y}^i}{N(N-1)} - \frac{\boldsymbol{y}^N}{N}\right| \tag{5.33}$$

$$= \left|\frac{\bar{\boldsymbol{y}}_{\text{mod}} - \boldsymbol{y}^N}{N}\right|. \tag{5.34}$$

We define $R_{\text{adapt}} = \text{diag}\left[\Delta\bar{\boldsymbol{y}}_1^2, \ldots, \Delta\bar{\boldsymbol{y}}_m^2\right]$ (for an m-dimensional measurement vector) as the *adaptive measurement covariance*, which is used in the update step of the centroid filter. For a sizable cluster, this term will be larger than the measurement covariance for an individual observation, which has the effect of decreasing sensitivity of the centroid estimate to the exclusion of individual observations during the gating step. For tighter clusters, the sensitivity to individual observations will be minimal.

5.4 Ellipsoid Similarity Measures

Given a set of full-state observations, it is known that the maximum likelihood estimate of the covariance matrix is given by Equation (3.21), otherwise known as the sample covariance [174]. Other Bayesian estimators of unknown covariance matrices also exist [165, p.273]. There are several reasons why none of these is appropriate for the current problem. At no point in the extent estimation process are observations expressed as full state vectors: the observed quantities are, at best, range and bearing, which are related to the equinoctial elements (state space) via a chain of nonlinear transformations (see Appendix C). Additionally, the sample covariance can yield unstable results when the number of samples is small [175], a situation which may arise here due to the potentially limited observability of individual cluster members (e.g. in the case of merged observations). Finally, the extent estimation problem as posed in this dissertation does not actually involve estimating the extent given observations of individual cluster members–instead, extent matrices in observation space are used to update the state space extent. The extent estimation problem involves distributions over positive definite matrices, given observations that themselves consist of positive definite matrices, with a nonlinear measurement function relating the two.

Koch [105] presents an example of this specific to extended target tracking, showing that the measurement likelihood is proportional to a Wishart density. However, the derivation, like others found in the literature, requires full-state measurements (resulting in a full-state measurement extent), making it unsuitable to the

satellite cluster tracking problem.

The task then becomes identifying a suitable approximation of the measurement likelihood function that may be used in this context. As Equation (5.10) shows, we must have some measure of the likelihood of a particle in order to properly calculate the weights. The approach taken here is to define some function that acts as a metric between the computed measurement-space extent given a particle $\boldsymbol{X}_{k+1}^{i}$ and the actual measurement extent $\boldsymbol{Y}_{k+1}$ to act as a proxy for the likelihood function. This function should yield greater values when the two matrices are similar (using definitions of similarity given in Section 5.4), and smaller values when the matrices are very different.

5.4.1 Similarity Measures

There exist several options for measuring the distance between–or similarity of—two ellipsoids. First, we consider the Frobenius norm. For a $m \times n$ matrix A, the Frobenius norm is defined as [176, p. 195]:

$$\|A\|_F = \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n}|a_{ij}|^2}. \tag{5.35}$$

The distance between matrices A and B can then be defined as $\|A - B\|_F$. Two similar matrices will have a low Frobenius distance between them, so the measure of similarity is defined as the reciprocal of $\|A_F\|$. While this is singular for identical matrices, that does not pose a practical issue in the extent estimation problem.

An alternative similarity measure called *compound similarity* is presented by

Moshtaghi *et al.* [177]. It is simply a combined measure of the distance between the center, orientation, and shape of two ellipsoids:

$$s_c = e^{-\|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|} e^{-\|\sin \boldsymbol{\theta}\|} e^{-\|\boldsymbol{\alpha}^* - \boldsymbol{\beta}^*\|}, \tag{5.36}$$

where $\boldsymbol{\mu}$ represents the center of the ellipsoid, $\boldsymbol{\theta}$ is a vector of angles between each pair of eigenvectors from the two ellipsoids, and $\boldsymbol{\alpha}^*$ and $\boldsymbol{\beta}^*$ are essentially the magnitude of each axis of the two ellipsoids. Note that when the ellipsoids share a center (as is the case for our application), this reduces to

$$s_c = e^{-\|\sin \boldsymbol{\theta}\|} e^{-\|\boldsymbol{\alpha}^* - \boldsymbol{\beta}^*\|}. \tag{5.37}$$

The Kullback-Leibler (KL) divergence is well-known as a measure between two arbitrary probability distributions. While it is not a specific measure of matrix or ellipsoid similarity, it is often used in conjunction with the assumption of multivariate normal distributions to compare covariance matrices, particularly in the context of image processing [178]–[181]. Additionally, Granström and Orguner [110] uses a method for extended target tracking based on minimization of the KL divergence. For two k-dimensional multivariate normal distributions $p_1 \sim \mathcal{N}(\boldsymbol{\mu}_1, \Sigma_1)$ and $p_2 \sim \mathcal{N}(\boldsymbol{\mu}_2, \Sigma_2)$, the KL divergence is defined as

$$d_{\mathrm{KL}} = \frac{1}{2}\left(\mathrm{tr}\left(\Sigma_2^{-1}\Sigma_1\right) + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \Sigma_2^{-1}(\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) - k + \ln\left(\frac{\det \Sigma_2}{\det \Sigma_1}\right)\right). \tag{5.38}$$

When $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$, this reduces to

$$d_{\mathrm{KL}} = \frac{1}{2}\left(\mathrm{tr}\left(\Sigma_2^{-1}\Sigma_1\right) - k + \ln\left(\frac{\det\Sigma_2}{\det\Sigma_1}\right)\right). \tag{5.39}$$

As with the Frobenius norm, the measure of similarity is taken to be the reciprocal of d_{KL}.

Förstner and Moonen [182] proposes a metric for covariance matrix similarity that depends on the eigenvalues $\lambda_i(\Sigma_1, \Sigma_2)$ from $|\lambda\Sigma_1 - \Sigma_2| = 0$. The distance measure is defined as

$$d_{\mathrm{Förstner}} = \sqrt{\sum_{i=1}^{n} \ln^2\lambda_i(\Sigma_1, \Sigma_2)}. \tag{5.40}$$

The authors note that this distance accounts for the difference in variance in each direction weighted by the ratio of the variances. Note that this measure does not account for differences in the mean, or center, of the covariance matrix.

Like the Kullback-Leibler divergence, the Bhattacharyya distance is another well-known measure between probability distributions [183], and has been used in clustering and image processing applications [184], [185]. For two multivariate normal distributions as defined above, the Bhattacharyya distance is defined as

$$d_{\mathrm{B}} = \frac{1}{8}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T\Sigma_m^{-1}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \frac{1}{2}\ln\left(\frac{\det\Sigma_m}{\sqrt{\det\Sigma_1\det\Sigma_2}}\right), \tag{5.41}$$

where $\Sigma_m = \frac{\Sigma_1+\Sigma_2}{2}$. When the distributions share a mean, this reduces to

$$d_{\mathrm{B}} = \frac{1}{2}\ln\left(\frac{\det\Sigma_m}{\sqrt{\det\Sigma_1\det\Sigma_2}}\right). \tag{5.42}$$

Note the similarities between the Bhattacharyya distance and the KL divergence.

Finally, the Hellinger distance is another measure between probability distributions, related to the Bhattacharyya distance in its construction [186], [187]. It is defined for two multivariate normal distributions as

$$d_{\mathrm{H}} = \sqrt{1 - \frac{\det \Sigma_1^{1/4} \det \Sigma_2^{1/4}}{\det \Sigma_m^{1/2}} \exp\left[-\frac{1}{8}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T \Sigma_m^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)\right]}, \tag{5.43}$$

where Σ_m is defined above. When $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2$, this reduces to

$$d_{\mathrm{H}} = \sqrt{1 - \frac{\det \Sigma_1^{1/4} \det \Sigma_2^{1/4}}{\det \Sigma_m^{1/2}}}. \tag{5.44}$$

The inverse of this distance is used to measure similarity.

5.4.2 Simple Comparisons

Because similarity between ellipsoids depends on the metric used, it is helpful to look at the qualitative behavior of these different metrics for a set of example ellipsoids. First, each measure's sensitivity to scale can be illustrated by comparing a set of concentric circles to the unit circle. The Frobenius norm, compound similarity, Kullback-Leibler distance, Förstner distance, Bhattacharyya distance, and Hellinger distance (inverting where appropriate to get a measure of similarity) are computed for each circle compared to the original unit circle, normalized, and plotted against the scale factor in Figure 5.8. Normalized values are used because they are most relevant to this problem: the similarity measures are used for particle weighting,

and in a particle filter, the weights must sum to one.

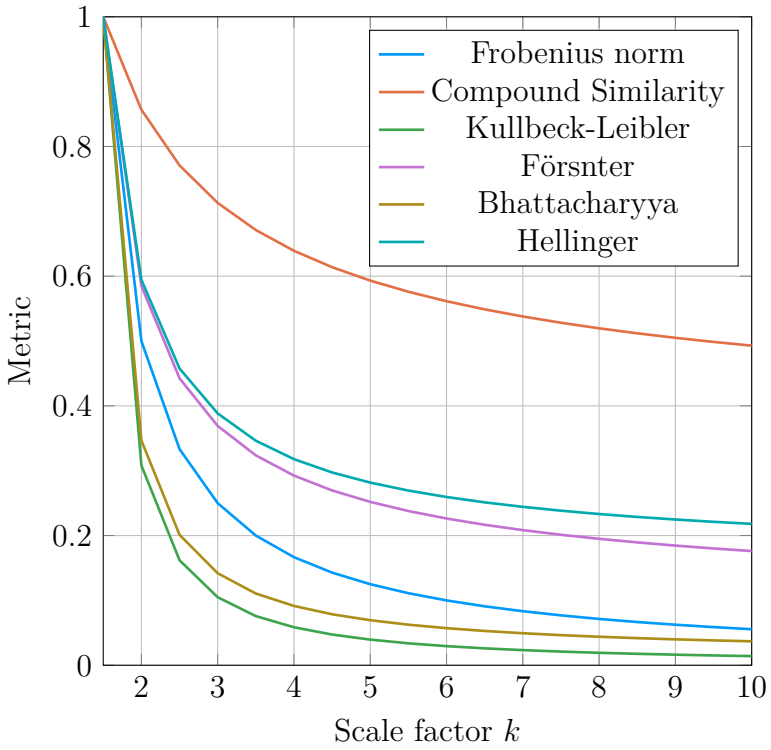

Figure 5.8: Similarity measures for a set of concentric circles.

All metrics display the expected trend: starting at one when the scale factor is close to unity, and decreasing as the scale factor increases. Compound similarity stands out in that it decays more slowly than the others, and the KL distance decays the fastest, indicating a higher sensitivity to the difference in scale between two ellipsoids.

Next, a set of random ellipsoids is drawn from the distribution $\mathcal{IW}(4, \mathbb{1}_2)$, and plotted along with the unit circle in Figure 5.9. The corresponding normalized similarity measures between each ellipsoid and the unit circle are plotted in Figure 5.10.

The pair of ellipses at c are most similar according to all metrics but the

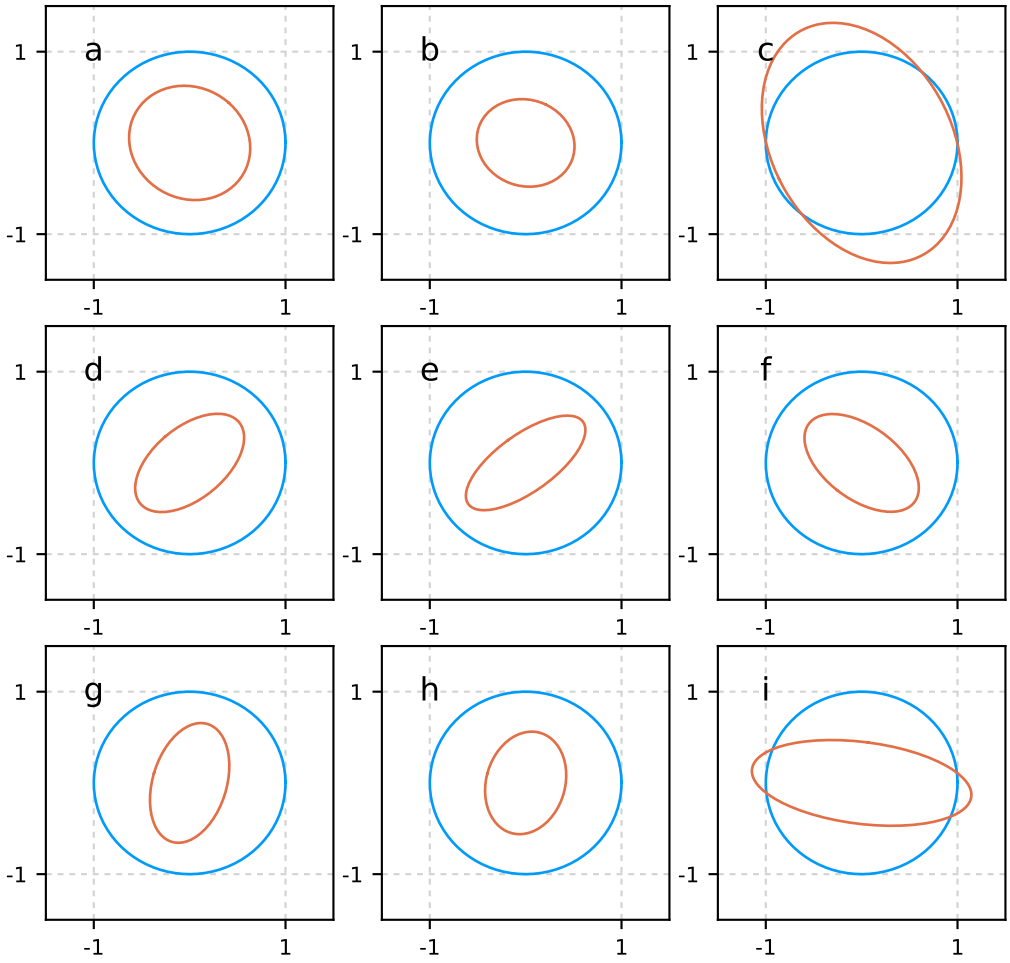

Figure 5.9: Random ellipses drawn from an Inverse Wishart distribution (orange), along with the unit circle (blue).

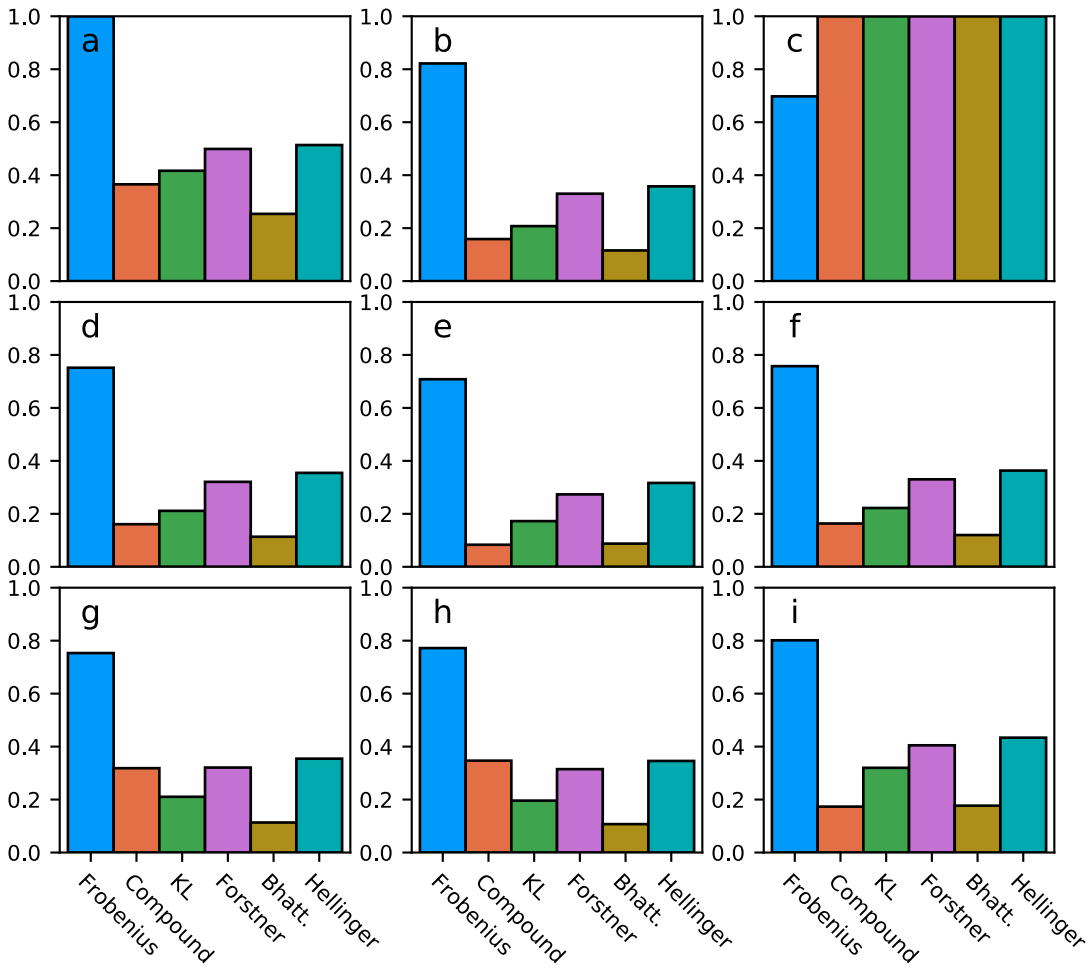

Figure 5.10: Similarity measures for the ellipse pairs in Figure 5.9.

Frobenius norm, which indicates the ellipses at a are most similar. Qualitatively, the latter pair have the most similar shape, while the former pair have a more similar size. In fact, the pair at c are the least similar of the set according to the Frobenius norm, while the other five metrics indicate the ellipses at e are the least similar. In all cases, the Förstner and Hellinger metrics have similar values.

Finally, although it is difficult to visualize an ellipsoid with more than three dimensions, it's important to look at the behavior of the similarity measures for higher-dimensional ellipsoids. For these cases, a set of 500 random 6-dimensional ellipsoids are drawn from the distribution $\mathcal{IW}(9, \mathbb{1}_6)$ (see Section B.2 for details), and compared with the unit hypersphere described by the 6-dimensional identity matrix $\mathbb{1}_6$. The normalized similarity measure values are plotted in the histogram in Figure 5.11, and the sample mean and variance of each distribution is listed in Table 5.1.

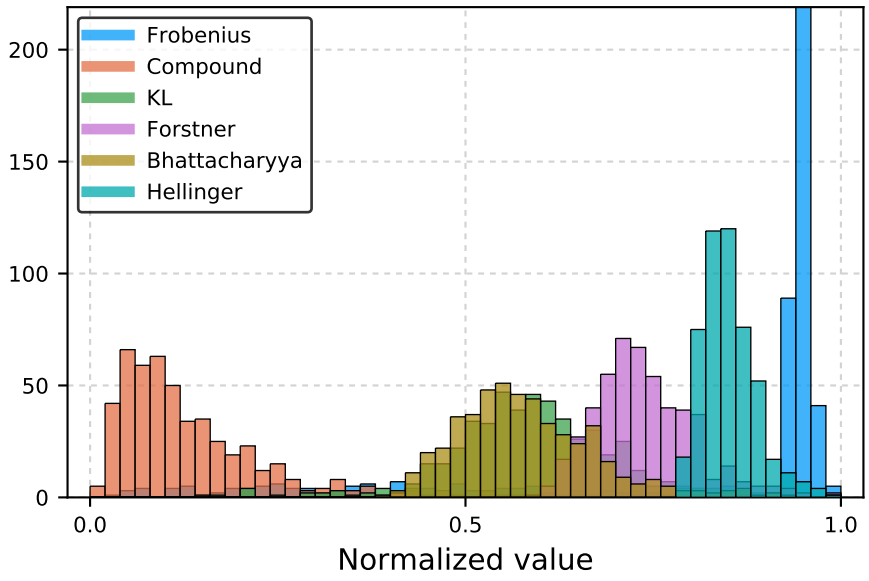

Figure 5.11: Distribution of similarity measure values for a set of 500 random 6-dimensional hyperellipsoids.

Measure	Sample mean	Sample variance
Frobenius	0.067746	0.80344
Compound sim.	0.013068	0.13920
KL	0.012237	0.58894
Förstner	0.004428	0.73853
Bhattacharyya	0.009180	0.58280
Hellinger	0.001231	0.85051

Table 5.1: Sample mean and variance for each ellipsoid similarity measure distribution plotted in Figure 5.11.

These distributions are informative. First, the inverse of the Frobenius norm is highly peaked near 1, with very little spread, a trend that is also evident in Figure 5.10. The other five have a larger spread, which indicates a more nuanced measure of similarity. The Bhattacharyya and KL divergence metrics are nearly overlapping, a result of the similar construction of the two metrics. The distribution of compound similarities is heavily skewed toward zero. Due to the use of normalization, the mean value of each measure is not particularly informative, but a higher variance in the distribution may indicate a more descriptive measure.

While few conclusions can be drawn from these simple comparisons (except perhaps that the Frobenius norm between two matrices is not a consistent reflection of the similarity between the two resulting ellipses) they demonstrate the qualitative characteristics of each measure. In the next section, these measures are compared using an actual extent estimation problem.

5.4.3 Simulation Results

A test case is constructed here for the purpose of investigating which similarity measure provides the most consistent results. A cluster is initialized at geosynchronous orbit (GEO) by sampling from the Cartesian covariance matrix $P_{\text{cart},0} = \text{diag}(20\ \text{km}^2, 20\ \text{km}^2, 20\ \text{km}^2, 0.02\ \text{km}^2/\text{s}^2, 0.02\ \text{km}^2/\text{s}^2, 0.02\ \text{km}^2/\text{s}^2)$. The cluster is generated from this distribution–rather than from, for example, a breakup model—in order to reduce outliers from the bulk of the cluster. Coupled with a very wide measurement gate, this allows us to observe the extent estimation process without conflating the results with the effects of gating.

Observations are simulated every 30 seconds from a single ground station for a period of 15 minutes. Figure 5.12 shows the longitudinal position of the cluster and the ground station position. Observation noise is modeled as zero-mean Gaussian distributed, with the noise covariance matrix $R = \text{diag}(225\ \text{m}^2, 225\ \mu\text{deg}^2, 225\ \mu\text{deg}^2)$. The effects of a 4×4 spherical harmonic gravity field and solar radiation pressure are modeled during propagation, and the same dynamics model is used in the filter.

Six tests are run, each one using a different matrix similarity measure as the likelihood function. The initial parameters are tuned slightly between runs to yield the best final extent estimate for the chosen likelihood function. The initial extent estimate is $kP_{\text{cart},0}$, where k is an integer multiplier (between 10 and 20 for these simulations), and $P_{\text{cart},0}$ is the initial Cartesian covariance matrix given above. It's worth noting that since we use the same covariance matrix to initialize the extent matrix as we do to generate the cluster, the filter is starting with a fairly well-fitting

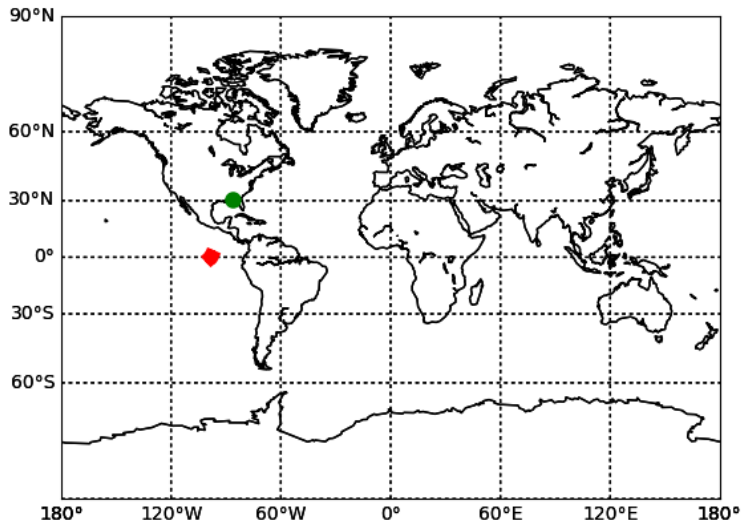

Figure 5.12: Observation geometry for the similarity measure test case. The satellite position is marked in red, and the ground station is the green circle.

extent matrix.

The initial ν parameter (which indicates the confidence in the initial extent estimate) ranges from 15 to 25 in all of these cases, while the bounds are $\nu_{\min} = 8$ and $\nu_{\max} = 30$. A small process noise matrix is used in several of the cases. The matrix used is pQ, where p is a small multiplier on the order of 1×10^{-9} and $Q = \text{diag}(10^{-9}, 10^{-7}, 10^{-7}, 10^{-7}, 10^{-7}, 10^{-3})$ is defined in equinoctial element space. Table 5.2 lists the values of these tuning parameters used in each case. All parameters were manually adjusted by running the filter several times and changing the values.

Because the range of magnitudes for each similarity measure varies significantly, and since, in the particle filter formulation, only relative magnitude matters, the results are plotted as normalized trend lines. Additionally, the measures are

Likelihood	Covariance multiplier k	Process noise multiplier p	ν_0
Inverse Frobenius	15	1×10^{-7}	20
Compound similarity	15	1×10^{-7}	17
Bhattacharyya	15	0	20
Förstner	15	1×10^{-9}	20
Kullback-Leibler	20	1×10^{-7}	25
Hellinger	10	0	20

Table 5.2: Table of parameters used for each test case.

inverted such that the results represent error between the estimated extent matrix and the true extent matrix. Thus, a value that converges to zero (an arbitrary value on these plots) is expected for a well-performing filter. Because the similarity comparison is performed in measurement space, and we are concerned with extent similarity in state space, all similarity measures are plotted for each likelihood option. Figures 5.13–5.24 show the extent projections and trend lines from each case. The extent estimates are plotted as green lines, the true extent is plotted in red, and the cluster components are blue dots.

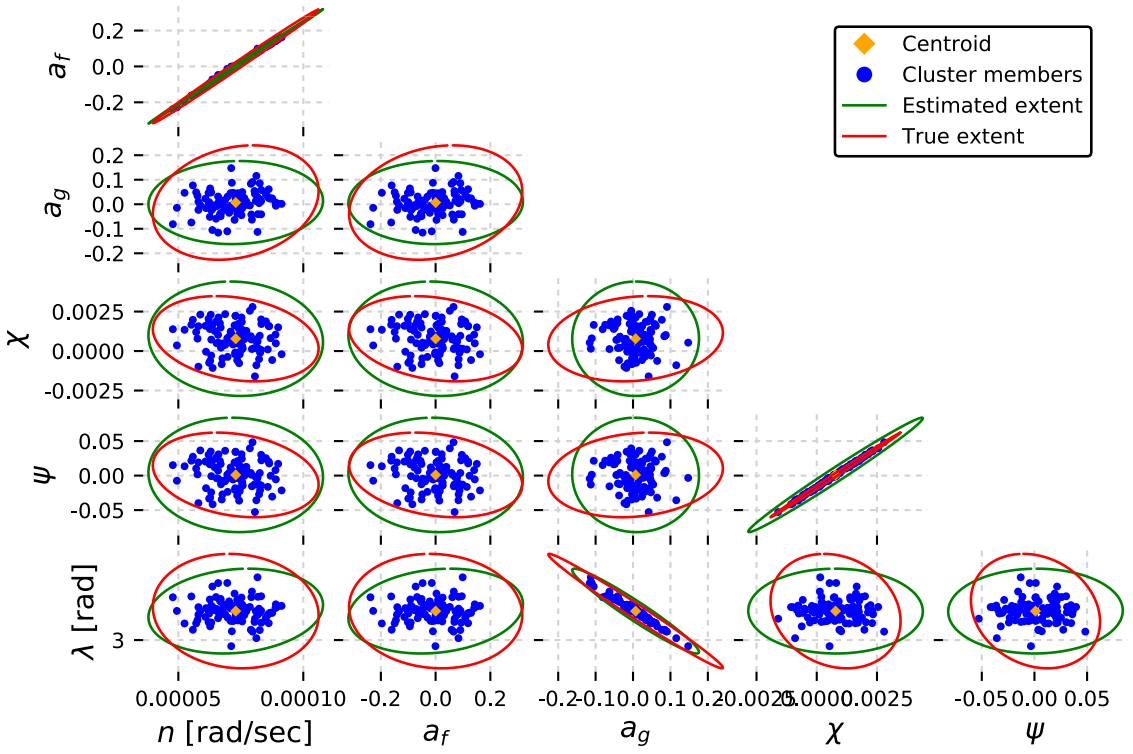

Figure 5.13: Extent estimate projections using the inverse Frobenius norm as the likelihood.

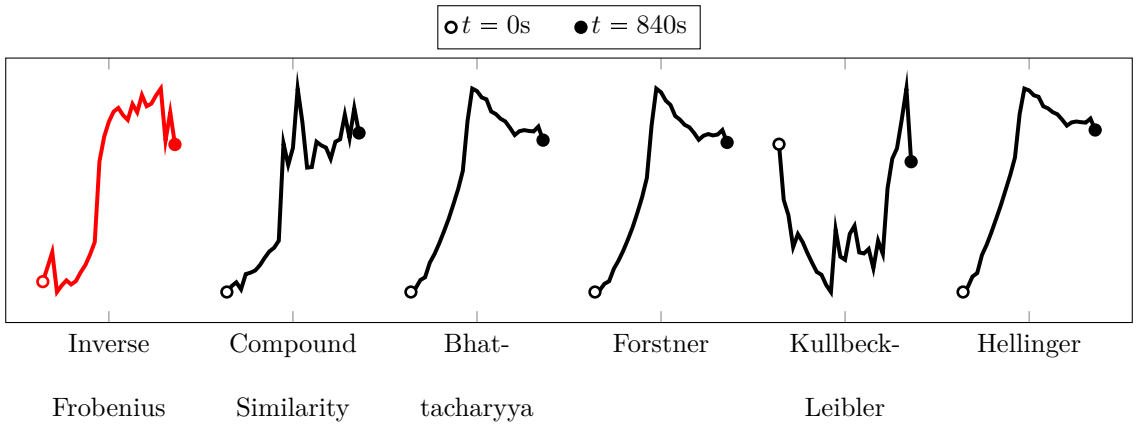

Figure 5.14: Error between true extent and extent estimated using each of the six similarities, when using the inverse Frobenius norm as the likelihood.

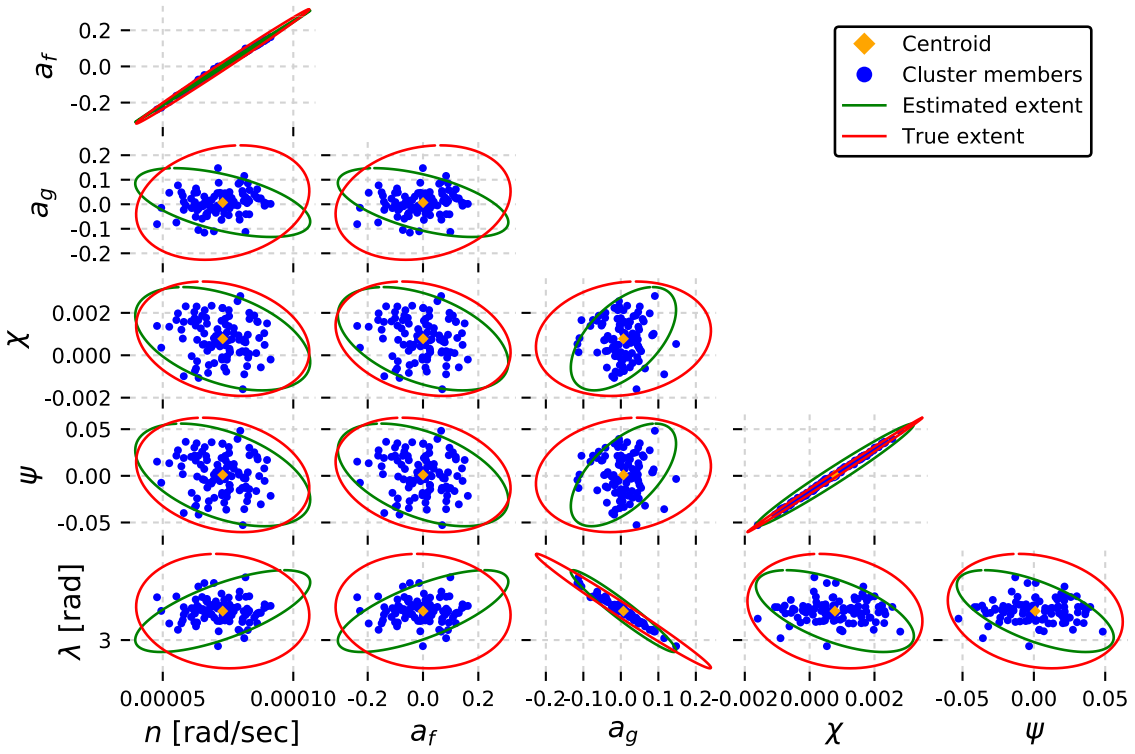

Figure 5.15: Extent estimate projections using compound similarity as the likelihood.

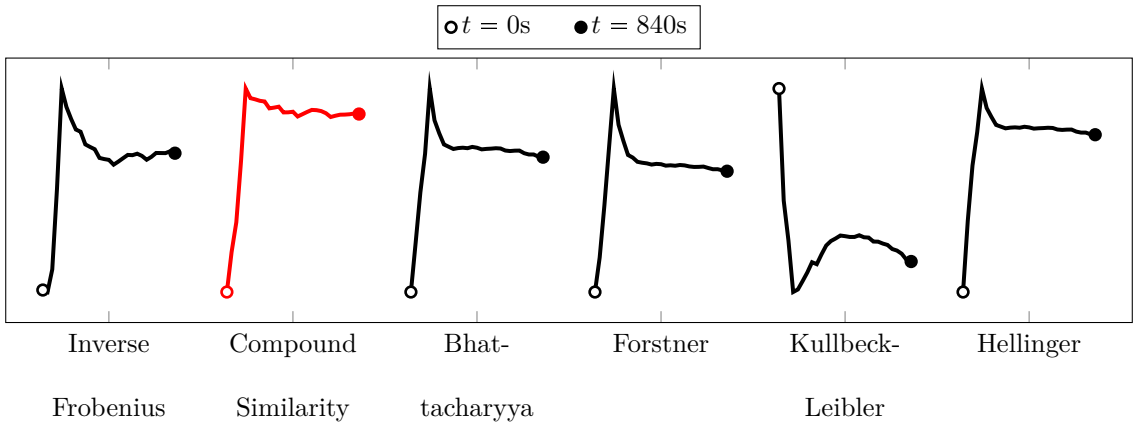

Figure 5.16: Error between true extent and extent estimated using each of the six similarities, when using compound similarity as the likelihood.

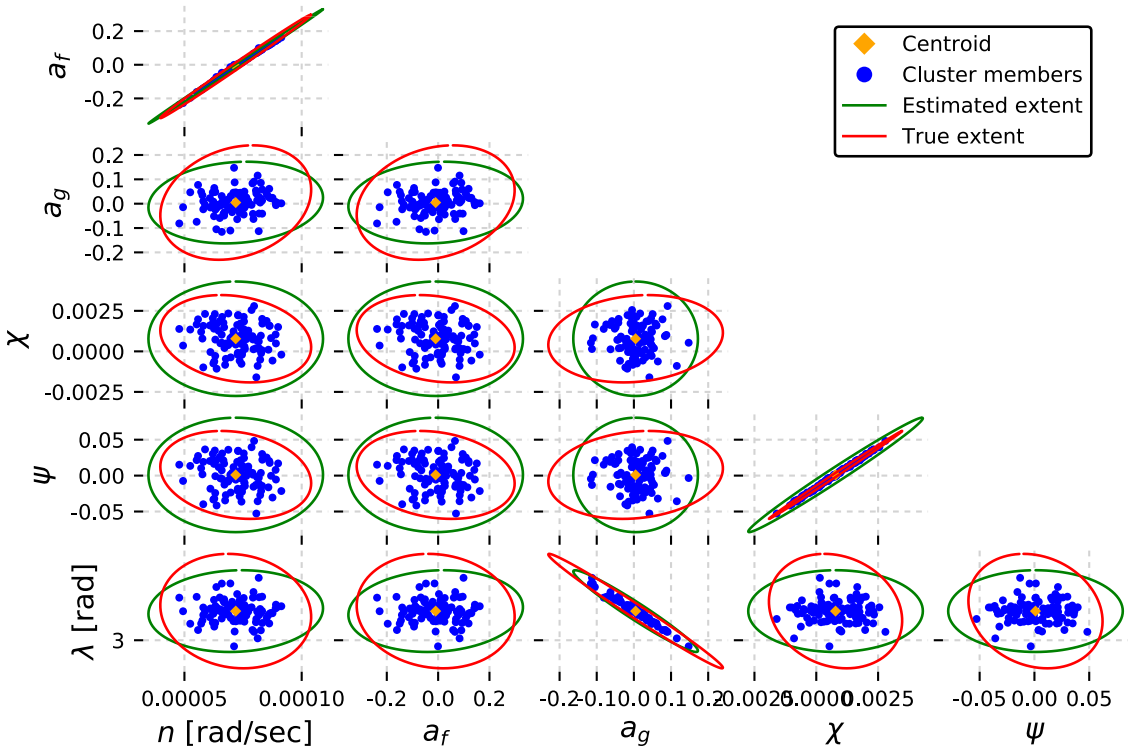

Figure 5.17: Extent estimate projections using Bhattacharyya distance as the likelihood.

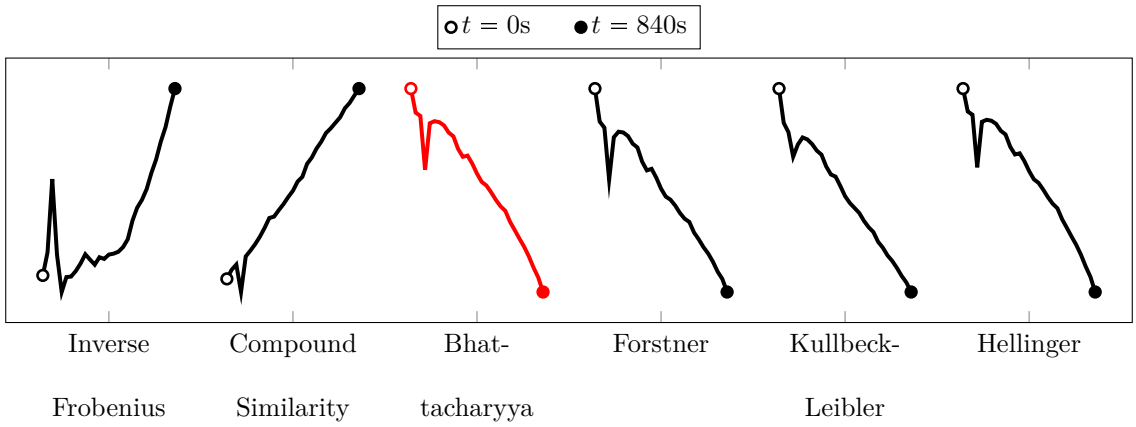

Figure 5.18: Error between true extent and extent estimated using each of the six similarities, when using Bhattacharyya distance as the likelihood.

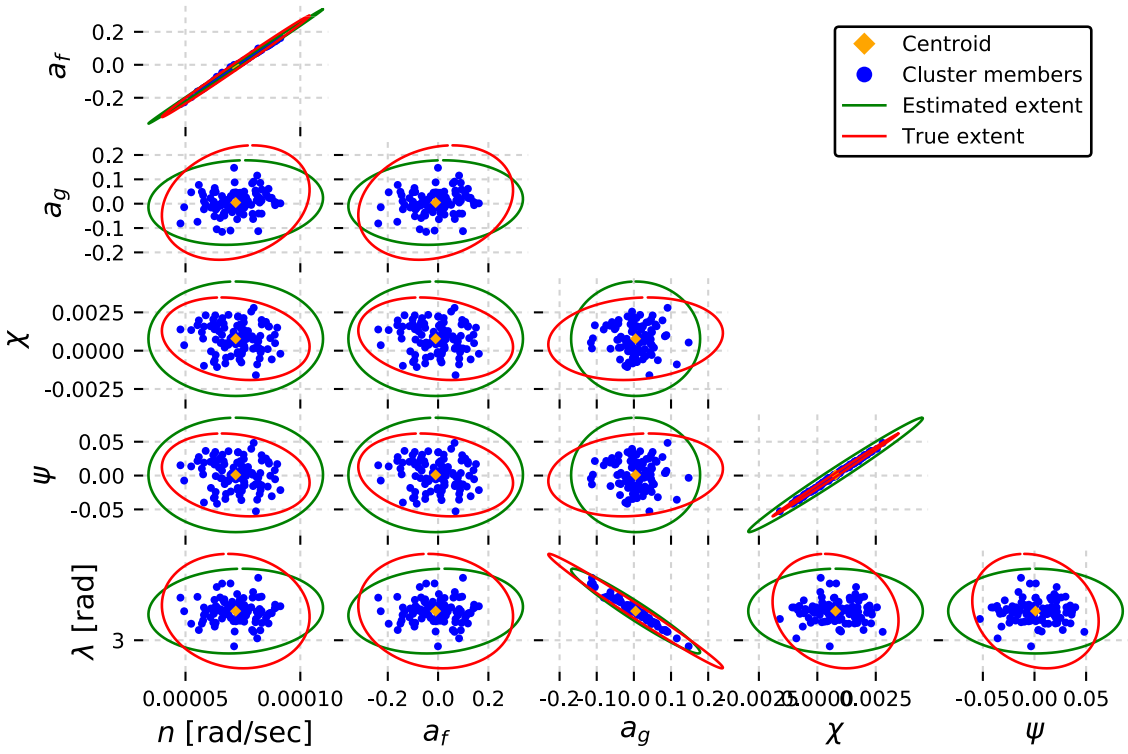

Figure 5.19: Extent estimate projections using Förstner distance as the likelihood.

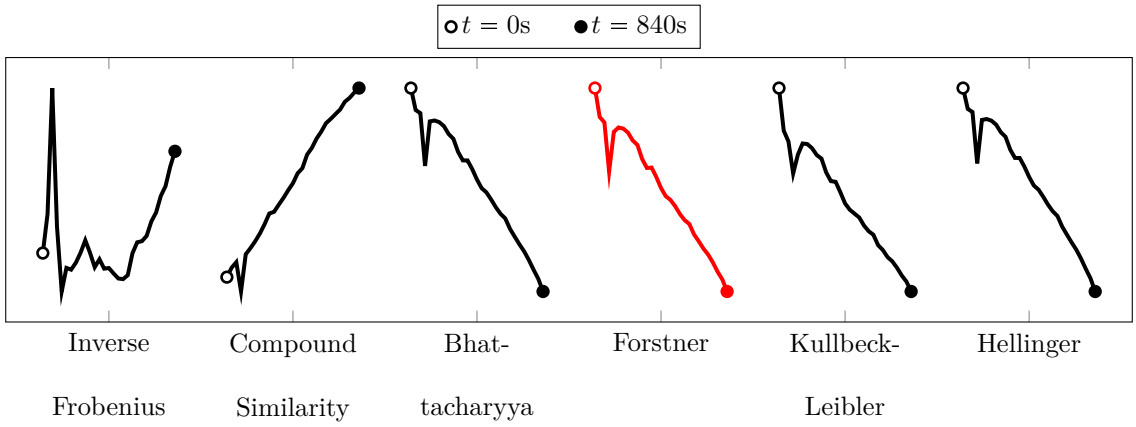

Figure 5.20: Error between true extent and extent estimated using each of the six similarities, when using Förstner distance as the likelihood.

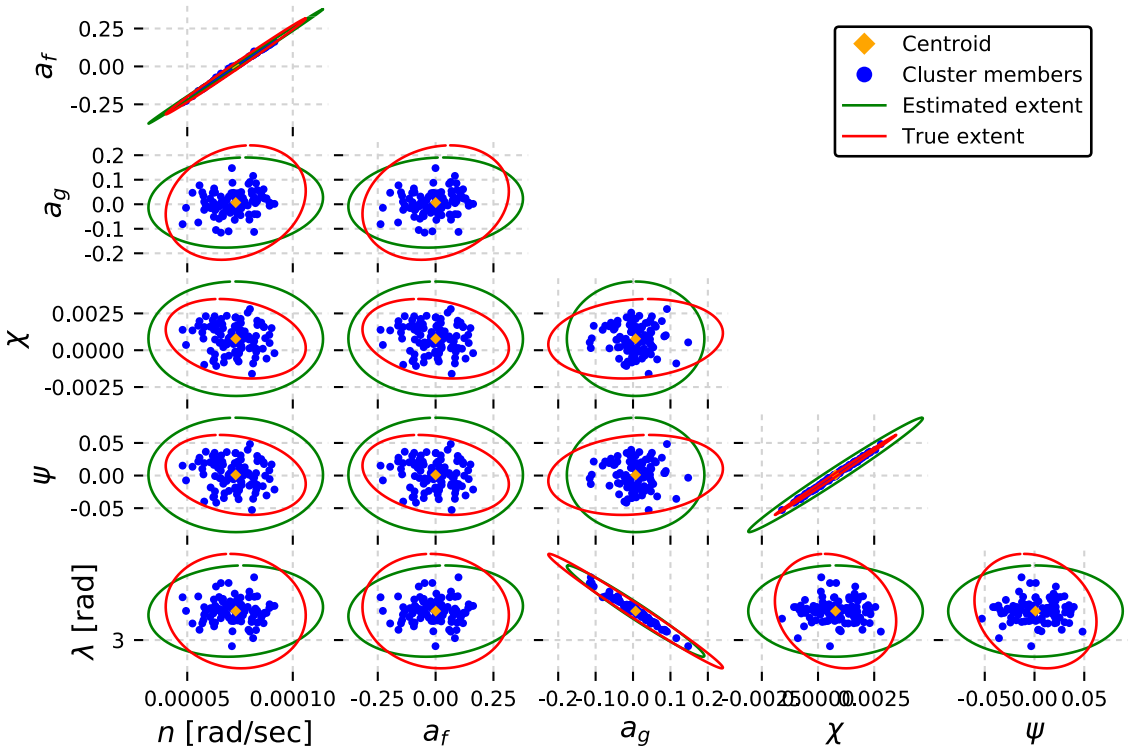

Figure 5.21: Extent estimate projections using Kullback-Leibler distance as the likelihood.

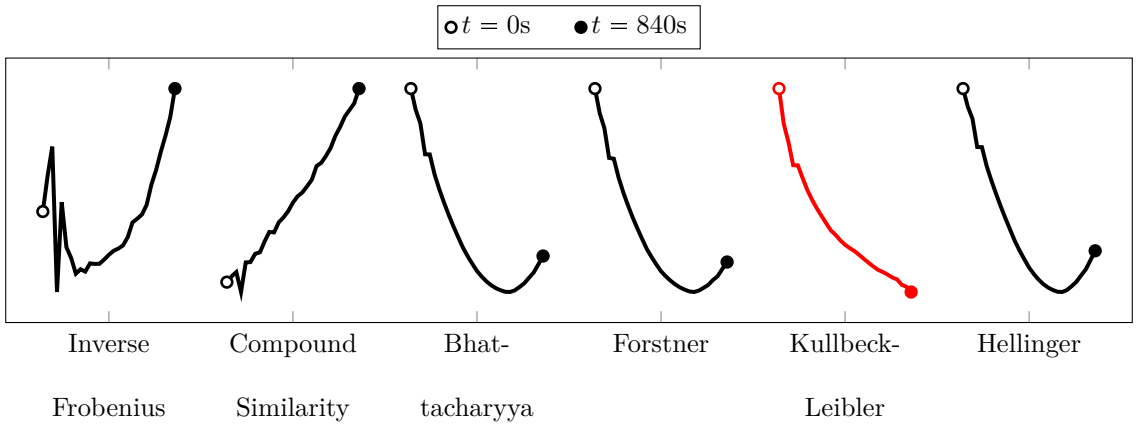

Figure 5.22: Error between true extent and extent estimated using each of the six similarities, when using Kullback-Leibler distance as the likelihood.

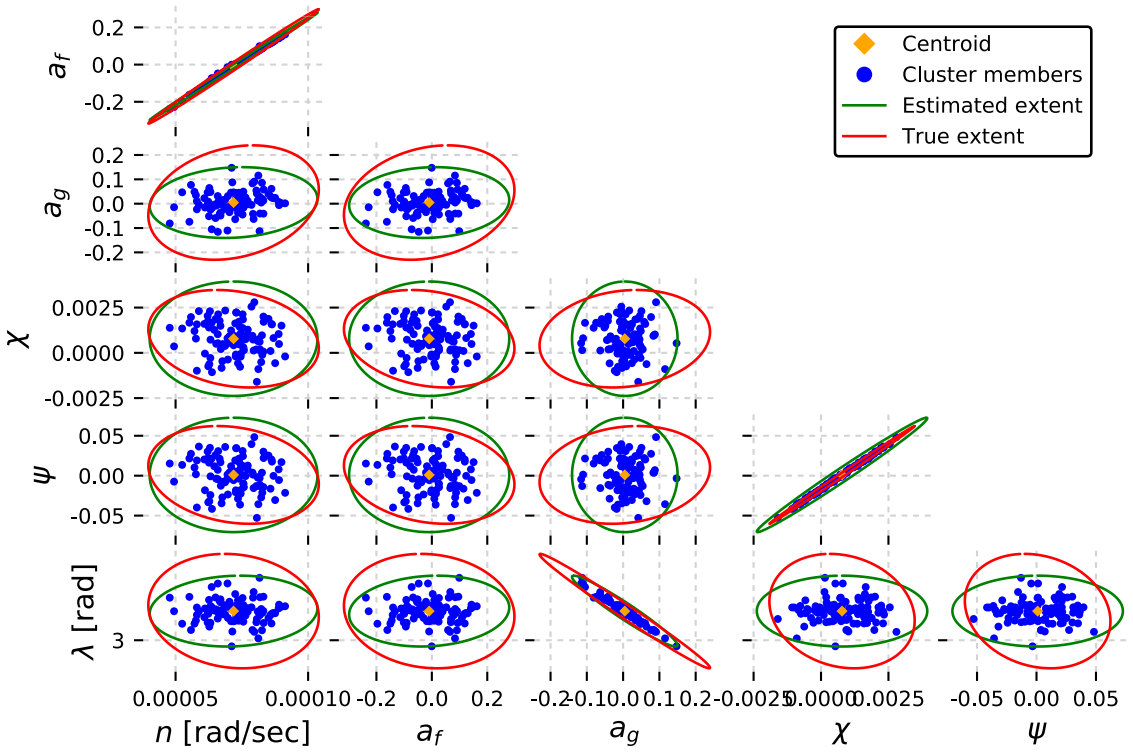

Figure 5.23: Extent estimate projections using Hellinger distance as the likelihood.

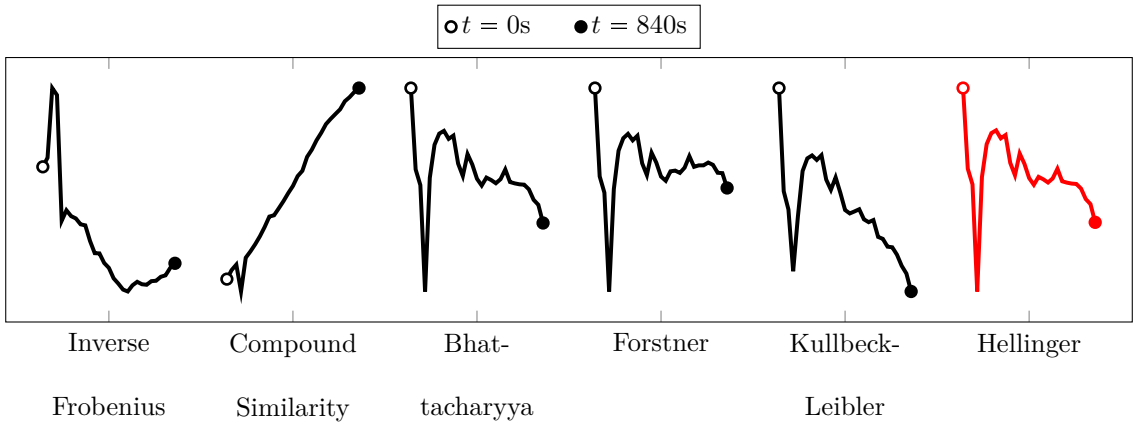

Figure 5.24: Error between true extent and extent estimated using each of the six similarities, when using Hellinger distance as the likelihood.

The final cluster estimate in most cases (the compound similarity case being a slight outlier) fits the cluster quite well when compared with the true extent,

computed by finding the minimum value enclosing ellipsoid for the set of cluster members. Note that in all cases, some of the projections, such as the (a_g, χ) plane, are consistently off from the true extent. This is likely due to the observability of the system, as this is a relatively short track of observations compared to the orbit period. Another factor is the sensitivity of the true extent to the convex hull of state-space points from which it is computed. This sensitivity is most apparent in the (n, a_f) projection, which is the most degenerate plane in this case.[3] A very small perturbation of a single cluster state in either of these elements will affect the true extent calculation noticeably, while the extent estimation process will not reflect such sensitivity. In fact, as discussed in Section 5.2.2, this sensitivity of the minimum volume enclosing ellipse is the main reason it is not chosen to model the measurement extent.

The similarity measures differ between cases significantly, however. The inverse Frobenius norm case shows little improvement over the estimation period by any measure. Compound similarity seems to change the estimate little, as all error measures level out to some biased value. Moreover, the compound similarity metric was found to be rather unstable, yielding values varying by many orders of magnitude.

The remaining four cases (Bhattacharyya, Förstner, Hellinger, Kullback-Leibler, and Hellinger) show some similarities. Over the estimation period, the quality of the extent estimate as measured by these same four metrics improves, while the quality as measured by the inverse Frobenius and compound similarity metrics seemingly

[3]See Section B.1 for a discussion of degeneracy in this context.

diverges. Importantly, those two metrics are some function of an element-wise matrix norm, while the four converging metrics are all functions of the determinant of the extent matrices. This suggests that element-wise matrix norms are not a good choice for a likelihood function in this application.

A likely reason for this could simply be the scale of the different components—for example, when comparing measurement extent matrices in range, azimuth, and elevation space, the likelihood will be dominated by the differences in the range terms of the matrices (as they are larger than the angle terms). This de-emphasis on the angles components could cause a diverging estimate, particularly in the state-space elements most directly connected with observations in those coordinates. Similarly, the extent comparisons in state space over-emphasize larger components, such as mean longitude, over smaller ones such as mean motion. The measurement space comparisons could potentially be improved via scaling of the range component of the measurement extent matrices, such that all components are on the same order of magnitude.

In order to illustrate that this is indeed the problem with the metrics based on element-wise matrix norms, a test case consisting of a single frame of measurement data from the GEO test case above is constructed. The measurement extent is computed both in native measurement space ($\boldsymbol{Y}$, units of kilometers and radians) and scaled measurement space ($\boldsymbol{Y}'$, Earth radii and radians) (GEO altitude is 35.786 km or about 5.61 Earth radii). 10,000 samples are drawn from the inverse Wishart distributions $\mathcal{IW}(\nu, \boldsymbol{Y})$ and $\mathcal{IW}(\nu, \boldsymbol{Y}')$, and ν is varied from 5 to 20. The value of each distance metric is recorded for every sample (relative to the base matrix), and

the values obtained using $\boldsymbol{Y}$ are differenced with those using $\boldsymbol{Y}'$. Importantly, the two sets of metrics should not differ significantly–the units of range should not affect the measure of similarity between matrices. Figures 5.25 and 5.26 show the average value of the difference between metrics for each ν value. Note that in Figure 5.25, the values all oscillate near zero, regardless of the value of ν. Meanwhile, there is obvious structure in Figure 5.26, indicating dependence on the scale of the measurement extent matrix. In other words, these metrics tend to be dominated by the range component of the measurement extent simply because it has a larger magnitude.

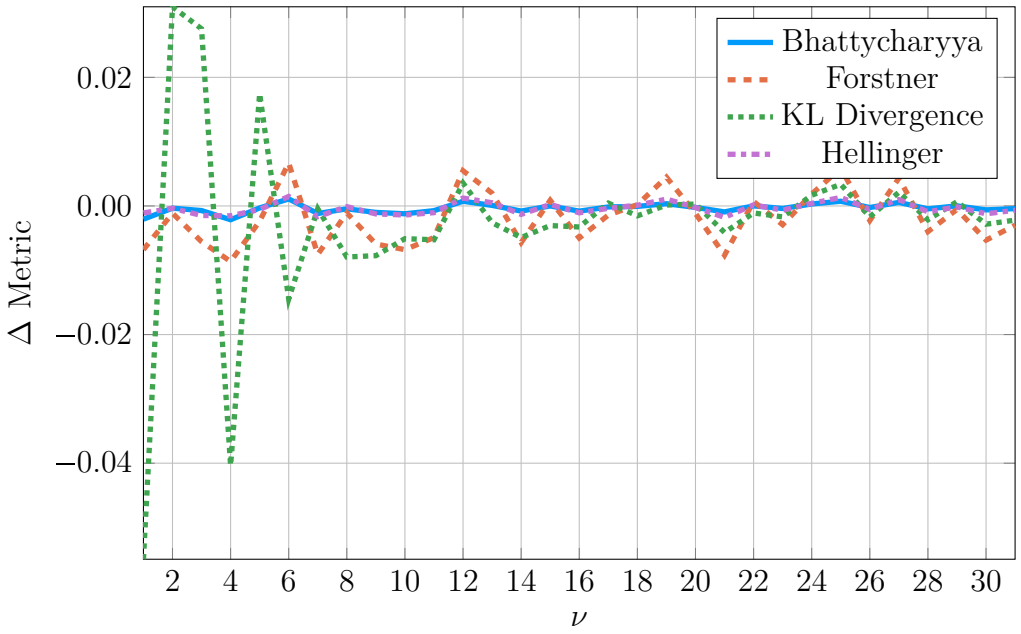

Figure 5.25: Sample mean of the differenced Bhattacharyya, Förstner, Kullback-Leibler, and Hellinger metrics using a scaled matrix versus an unscaled one.

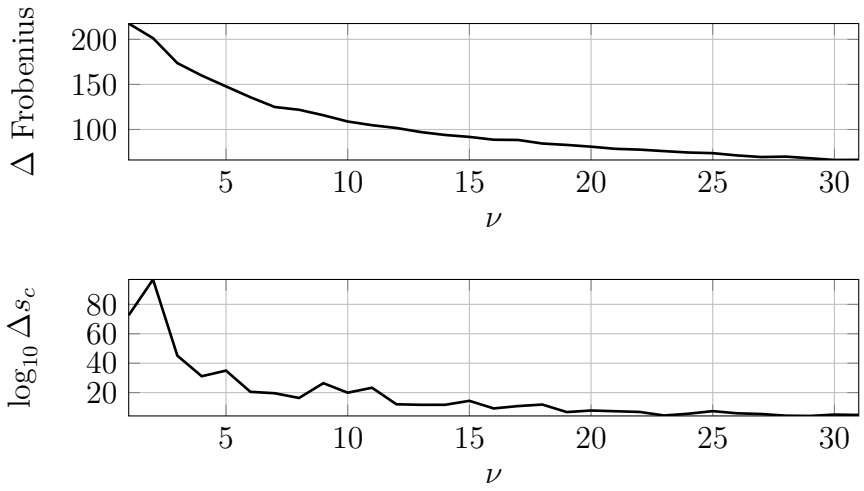

Figure 5.26: Sample mean of the differenced Frobenius and compound similarity metrics using a scaled matrix versus an unscaled one.

5.4.4 Consistency

It is also constructive to investigate the behavior of each similarity measure in the context of the matrix particle filter presented in Section 5.3. This is accomplished in the following way. First, $10,000$ samples are generated from the six-dimensional Inverse Wishart distribution $\mathcal{IW}(\nu, (\nu - p - 1)\mathbb{1}_6)$ (see Equation (5.17)) These samples may be treated as full-state observations of the base matrix $\mathbb{1}_6$. A weighted, moving-window average is then computed for increasing subsets of samples, using the aforementioned likelihood function candidates as weighting functions. The difference between the weighted means and the base matrix (using the same metric as the weighting function) is plotted for increasing sample size. This is shown for several discrete values of ν.

In the filtering scheme presented in Section 5.3, ν increases with successive observations. Therefore, there are several desirable features in each plot. First, for

all values of ν, the estimate should asymptotically approach the true value (or, the difference should asymptotically approach zero), with minimal bias. Second, these trends should improve for increasing values of ν, indicating improvement as the filter increases its measure of certainty in the estimate.

Generally, more samples reduces the error in all cases. The element-wise measures are biased from zero, as is the Hellinger distance. The Kullback-Leibler divergence and Bhattacharyya distance show the fastest convergence across all values of ν.

5.4.5 Summary

The main motivation for this section was to seek a consistent metric for evaluating the quality of a cluster estimate–that is, the metric should perform well as a likelihood function during extent estimation, and that quality should be reflected in the *a posteriori* evaluation. These results show that determinant-based metrics provide greater consistency than element-wise matrix norms. Most of the four determinant-based metrics proposed yield similar performance, as seen in the GEO-based demonstration case presented. Investigating the consistency of the proposed measures also proved to be informative. Based on these results, both the Bhattacharyya distance and the Kullback-Leibler divergence appear to be acceptable approximations for the measurement likelihood function in the extent filter (perhaps unsurprising, given the similarities the two share). In the next chapter, the Bhattacharyya distance is used in the extent filter.

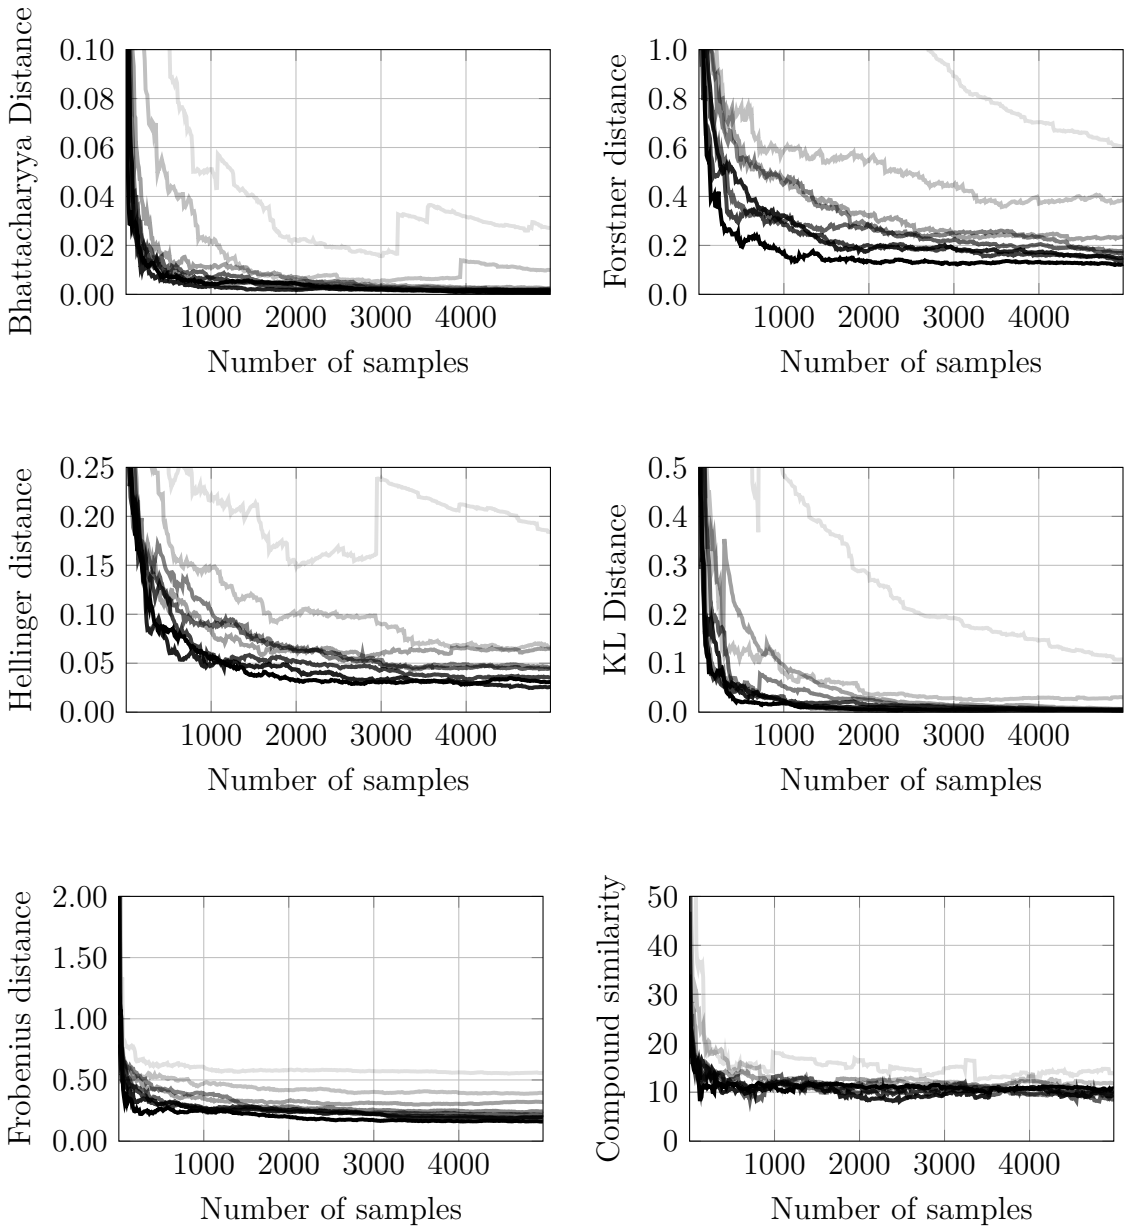

Figure 5.27: Moving-window average of each matrix difference metric. Darker lines correspond to higher values of ν; the lightest line corresponds to $\nu = 8$, and the darkest line corresponds to $\nu = 15$.

Chapter 6: Results

In order to validate the extent estimation method, three simulated test cases are analyzed. All three involve an on-orbit breakup (generated using the NASA standard breakup model, described below), with observations collected some short time after the event. All cases use realistic perturbing force models for the given orbit class.

6.1 NASA Standard Breakup Model

NASA's Standard Breakup Model (SBM), a component of the broader EVOLVE 4.0 and LEGEND [188] debris models, is a set of distributions driven by empirical data that describe the population of debris resulting from on-orbit breakups [189]. Using the minimum characteristic length (L_c) of the debris population as the independent variable, SBM gives distributions for the total quantity, as well as the area to mass ratio and ΔV of each piece. This dissertation uses the SBM models associated with explosive fragments of upper stages with an initial mass of 600-1000 kg. A brief summary of the relevant parts of SBM as presented in Johnson *et al.* [189] follows.

The number of fragments N is determined by the power law

$$N(L_c) = 6L_c^{-1.6}. \tag{6.1}$$

The area to mass ratio distribution $D_{A/m}$ is a weighted combination of normal distributions dependent on $\lambda_c = \log_{10} L_c$, described in full in Johnson *et al.* [189]. The distribution of $\log_{10}(\Delta V)$ is

$$D_{\Delta V} = \mathcal{N}(\mu, \sigma), \tag{6.2}$$

where

$$\mu = 0.2\chi + 1.85 \tag{6.3}$$

$$\chi = \log_{10}(A/m) \tag{6.4}$$

$$\sigma = 0.4. \tag{6.5}$$

A debris population from a breakup is generated in the following way. Given the desired minimum characteristic length, the number of pieces N is computed from Equation (6.1). The area to mass ratio of each piece is determined by taking N draws from $D_{A/m}$. Finally, the ΔV magnitude of each piece is determined by a random draw from Equation (6.2), using the area to mass ratio of each object to determine χ. The model assumes fragments are dispersed isotropically. A uniformly-distributed random direction vector is computed by

$$\hat{\boldsymbol{v}} = \frac{\boldsymbol{q}}{\|\boldsymbol{q}\|}, \tag{6.6}$$

where $\boldsymbol{q}$ is a vector of three draws from the unit normal distribution $\mathcal{N}(0, 1)$ [190]. Figures 6.1 and 6.2 show 1000 draws from the area to mass ratio and ΔV

distribution functions, respectively, along with the distribution function itself.

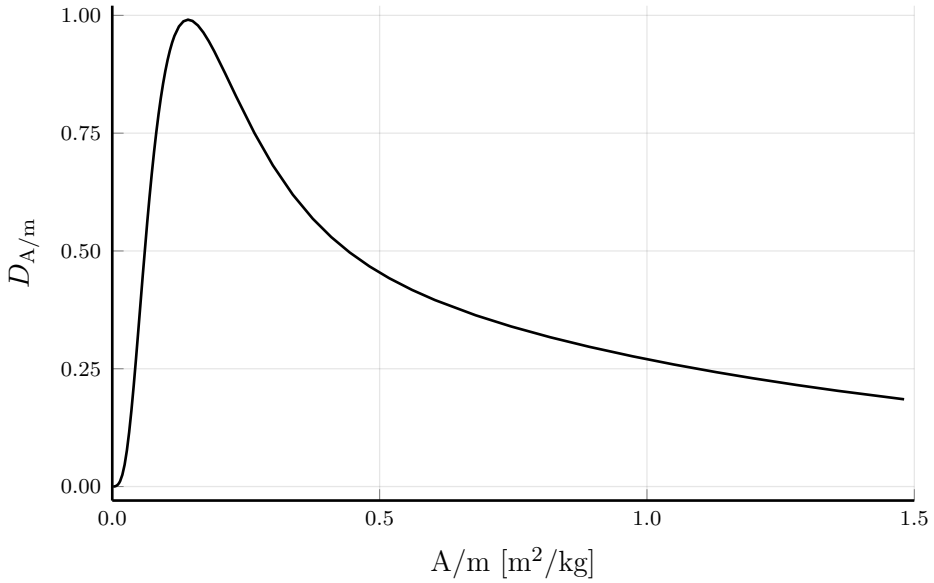

Figure 6.1: Plot of the probability density function $D_{\mathrm{A/m}}$ for $\lambda_c = 10$ cm.

6.2 LEO Test Case

The first test case explored is a breakup in a 500 km circular orbit. The pre-breakup state is listed in Table 6.1. The breakup occurs over the Pacific, and the cluster is observed by a fictitious ground station in New Mexico about four minutes later. The threshold characteristic length used for the NASA breakup model is 10 cm (said to be the approximate lower sensitivity of the U.S. Space Surveillance Network [189]), resulting in an initial population of 227 objects. Equinoctial element projections of the initial cloud distribution are plotted in Figure 6.4. Propagation includes the effects of an 8×8 spherical harmonic Earth gravity model and drag using an exponential atmosphere model (see Chapter 3 for details), and the estimator uses this same dynamics model.

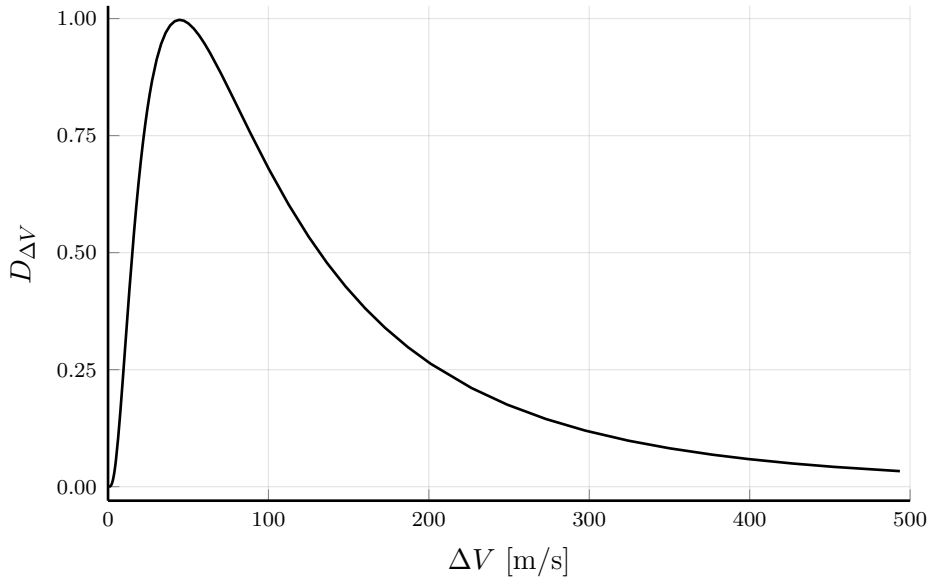

Figure 6.2: Plot of the probability density function $D_{\Delta V}$ for $\lambda_c = 10$ cm.

Table 6.1: Pre-breakup orbit parameters for the LEO test case.

a	6875.7 km	n	0.00110737 rad s^{-1}
e	0.000596618	a_f	0.000170519
i	0.610618 rad	a_g	0.000571731
Ω	−1.39651 rad	χ	−0.310388
ω	2.67746 rad	ψ	0.0546494
ν	4.33258 rad	λ	5.61464 rad

Area/mass ratio: 0.004 m^2/kg
Epoch: 1 Jan 2016 00:07:00.00 UTC

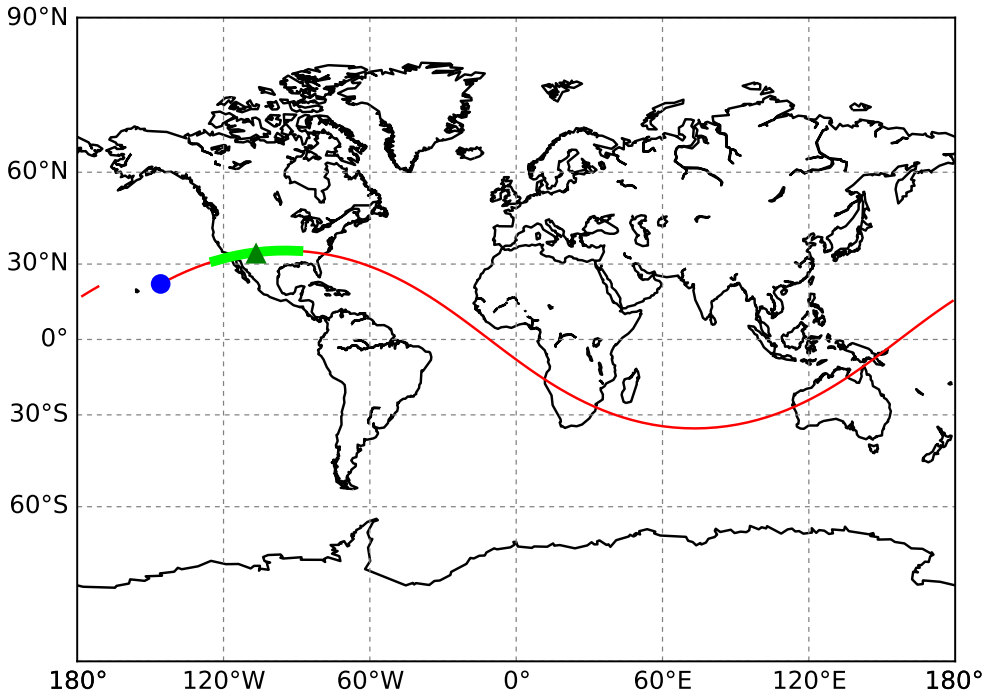

Figure 6.3: Ground track map for the LEO test case scenario. A dark green triangle marks the ground station, the visible portion of the ground track is in light green, and the breakup point is in blue.

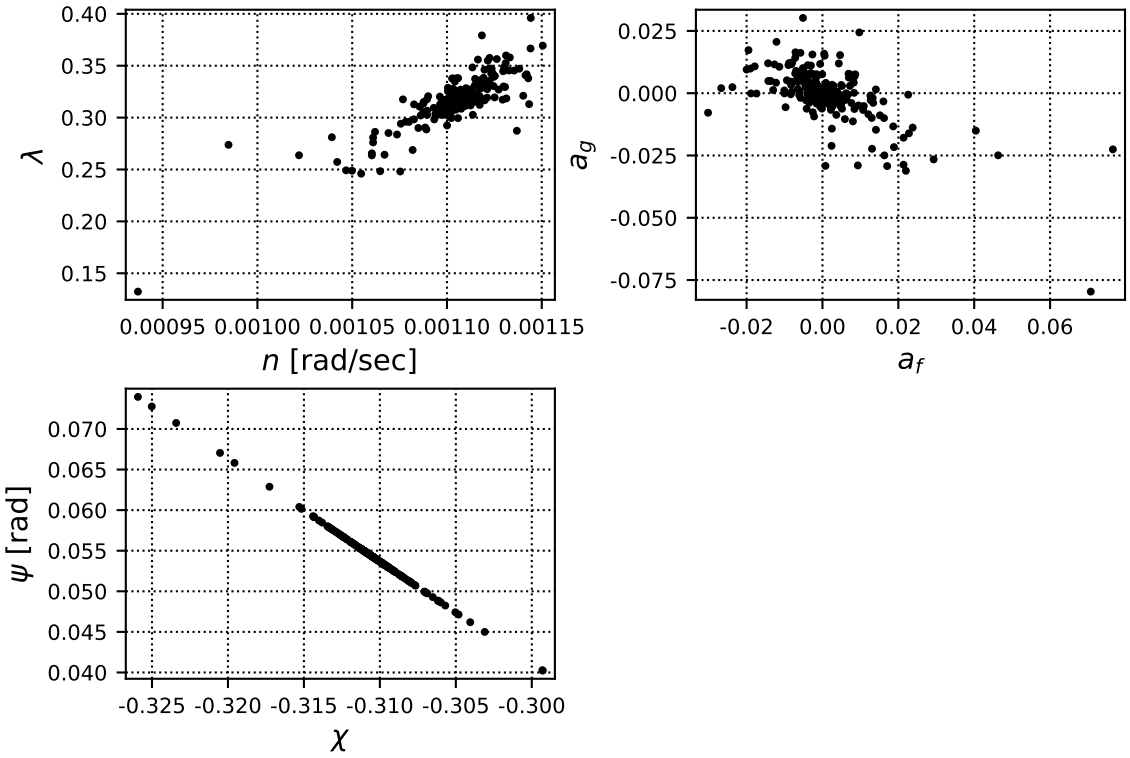

Figure 6.4: Initial distribution of the debris cloud for the LEO test case.

Range and angles observations, such as from a radar, are corrupted with zero-mean Gaussian noise:

$$\boldsymbol{y} = \boldsymbol{y}_{\text{true}} + \boldsymbol{\nu} \tag{6.7}$$

$$\boldsymbol{\nu} \sim \mathcal{N}(\mathbf{0}, R), \tag{6.8}$$

where $R = \text{diag}(\sigma^2_{\text{Range}}, \sigma^2_{\text{Azimuth}}, \sigma^2_{\text{Elevation}})$. The noise variances used in the simulation are listed in Table 6.2. The pass over the New Mexico ground station lasts approximately 10 min and reaches a maximum elevation of about 80°.

The extent estimate is initialized as follows: the centroid is initialized with the

Table 6.2: Observation noise parameters.

σ_{Range}	0.015 km
σ_{Azimuth}	0.015 rad
$\sigma_{\text{Elevation}}$	0.015 rad

Table 6.3: Summary description of filter parameters.

Parameter	Description
$N_{\text{particles}}$	Number of extent matrix particles
$\nu_{\min}$	Minimum value of extent estimate "confidence" ν
$\nu_{\max}$	Maximum value of extent estimate "confidence" ν
τ	With β, forms the time constant for the evolution of ν
β	With τ, forms the time constant for the evolution of ν
g	Measurement gate size

pre-breakup state, the extent is initialized with the spherical Cartesian covariance matrix $P_{\text{cart},0} = \text{diag}(50\ \text{km}^2, 50\ \text{km}^2, 50\ \text{km}^2, 0.05\ \text{km/s}^2, 0.05\ \text{km/s}^2, 0.05\ \text{km/s}^2)$ (converted to equinoctial element space via the unscented transform), and $\nu_0 = 10$. See Section 5.3.5 for a discussion of extent estimate initialization. In this scenario, zero process noise is used. The other free parameters in the extent particle filter (described in Table 6.3) are set to the values listed in Table 6.4.

Three equinoctial element projections of the estimated extent and centroid at the end of the pass are plotted in Figure 6.5. The gating process eliminated 14 outliers. The evolution of the fit metric for the extent matrix versus the true

Table 6.4: Filter tuning parameters used for the LEO test case.

Parameter	Value
$N_{\text{particles}}$	10,000
$\nu_{\min}$	8
$\nu_{\max}$	30
τ	5400
β	600
g	1.2

extent (Bhattacharyya distance, d_B) is plotted in Figure 6.6.[1] Each point on the plot represents a measurement frame, and filled points are times at which the entire estimated cluster is in view, thus there is no imputation of occluded parts of the cluster (see Section 5.3.6). As both of these figures show, the estimated extent is very close to the true extent. Sudden jumps in these plots, such as the one at 280 seconds elapsed time in Figure 6.6, correspond to the gating process including or excluding observations from a cluster member at that time step, resulting in a temporary lag between the estimated extent and the true extent (which changes instantaneously), until the filter "catches up" by processing the new observations.

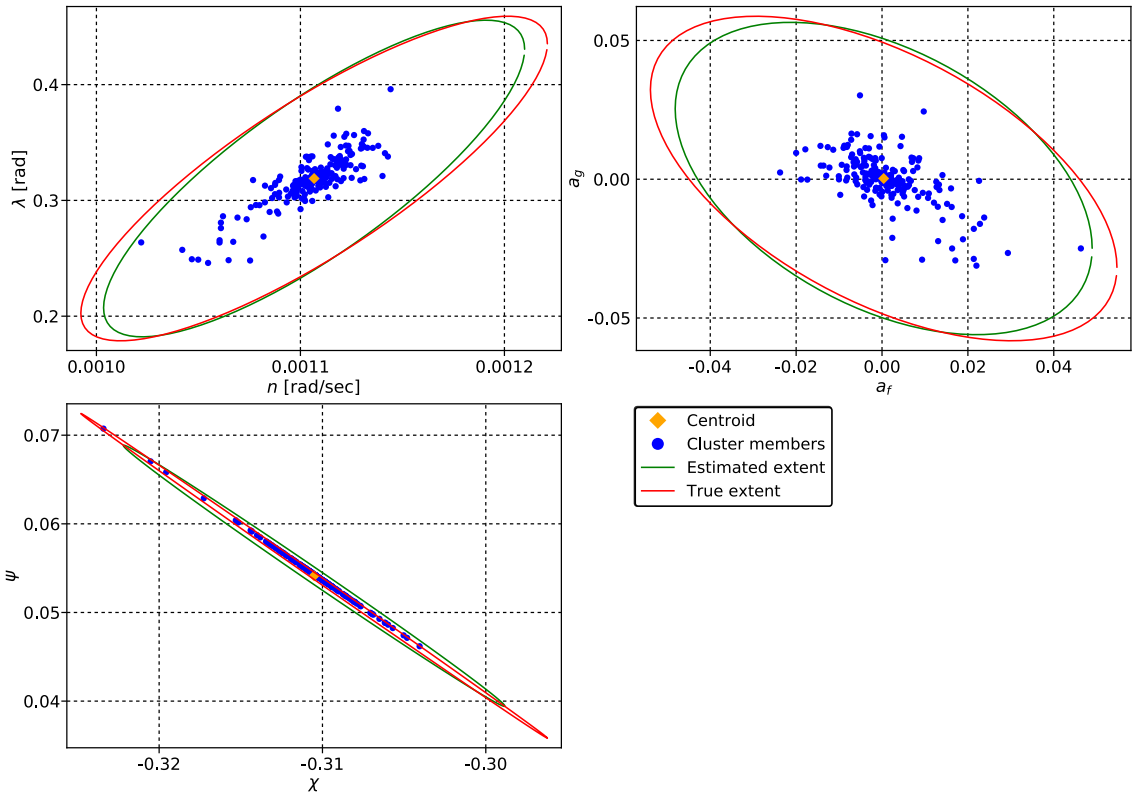

Figure 6.5: Extent estimate at the end of the pass for the LEO test case.

The error and uncertainty in the centroid estimate are shown in Figure 6.7.

[1]Note that while the inverse Bhattacharyya distance is used as the likelihood function, the actual distance is shown for plotting purposes.

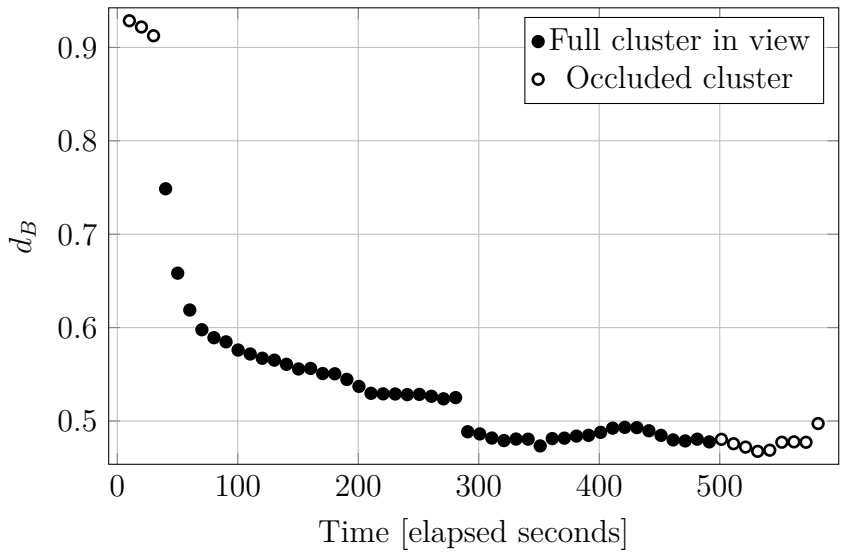

Figure 6.6: Fit metric evolution for the extent estimate for the LEO test case.

As soon as the extent is fully in-view of the ground station (about 30 seconds into the simulation), the uncertainty in the centroid estimate is driven down, reaching a steady-state value of about 8 km. The error remains below 5 km over the entire pass. While this is not particularly precise when compared to a single object orbit determination, it is approximately 0.5% of the largest extent dimension, and in line with the results shown in Chapter 4. As discussed above, the sudden changes in the error correspond to changes in cluster membership as determined by the gating process.

In addition to the similarity metric between the two extent matrices, an informative way to visualize the behavior of the extent estimate is to plot the Cartesian volume over time. This is meaningful because of the way the extent estimates are likely to be used–for defining volumes in space which should be avoided (if flying a spacecraft) or sought out (if tasking a sensor network). Both of these tasks require a description of the cluster in Cartesian space, instead of the native equinoctial

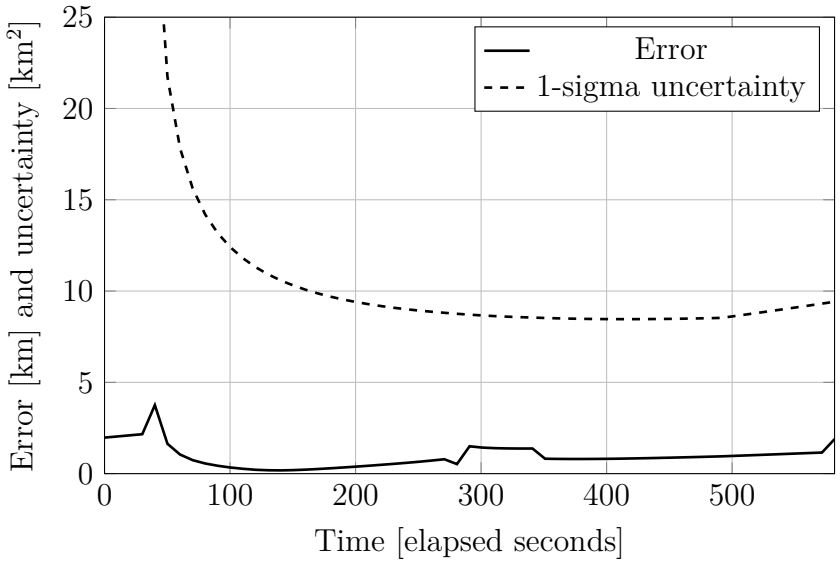

Figure 6.7: Cartesian position error and uncertainty for the LEO test case.

element space of the estimator.

The extent estimate is propagated to each point using the similarity transform (see Section 3.3), then transformed to Cartesian space, where the volume is computed. This is plotted with the true extent, which is calculated by propagating each cluster member to each point in time, computing the minimum volume enclosing ellipse in equinoctial space, then transforming into Cartesian space to calculate the volume. Figure 6.8 shows the evolution of the volume of the cluster, predicted over the next two periods[2]. The estimated volume tracks the true value very closely. Note that at the pinch point and anti-pinch line (as described by Hujsak [18] and Chobotov *et al.* [19], i.e. the breakup point and a half period later), the volume is minimized, while maximum volume occurs 90° away from these points. The volume difference between predicted and estimated is also minimized at the pinch and anti-pinch points.

[2]The times are relative to the breakup time, not the end of the ground station pass.

The estimated volume is systematically lower than the true volume. A likely explanation for this is that the expansion rate of the cluster is not strongly observable over the ground station pass. Velocity information is limited when using only position (range and angle) measurements over a relatively short span, even for single object orbit determination. Additionally, since the true volume calculation does not account for objects that might typically be gated out of the estimation process, a small number of individual objects might cause the volume to grow substantially by drifting away from the original reference orbit faster than the other objects in the cluster. During the observation span, these objects may be near the middle of the cluster, masking their contributions to the extent size and shape.

As another way of visualizing the extent estimate, the cluster members (black dots) and the extent projection (green area) are plotted in the perifocal orbit frame referenced to the centroid in Figures 6.9 and 6.10[3]. The growth in the in-track direction is clearly evident after two periods, and the extent contains all cluster members. This shows that the estimated extent is able to predict the gross size, shape and behavior of the debris cloud, which is a primary goal of this work, and implies the potential ability to accurately estimate the breakup time by propagating backwards and searching for minima. It is also notable that these estimates were formed over only a single pass, using about 10 minutes of data.

[3]Due to the nonlinear transformation between the native extent coordinate system (equinoctial elements) and the Cartesian system, the extent projections are visualized by sampling uniformly over the entire volume of the extent in equinoctial elements, then transforming the points individually into the Cartesian perifocal frame.

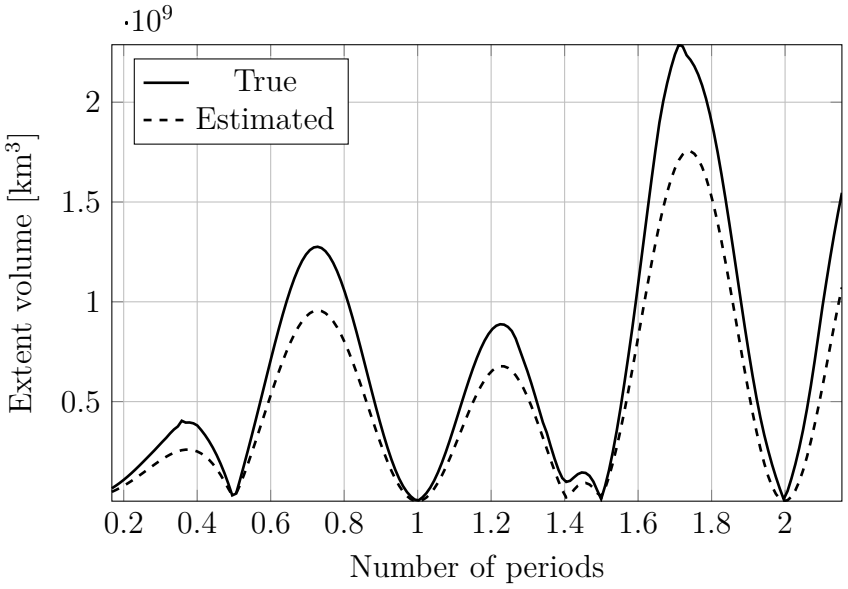

Figure 6.8: Predicted Cartesian volume of the estimated extent, with the volume of the true extent for the LEO test case.

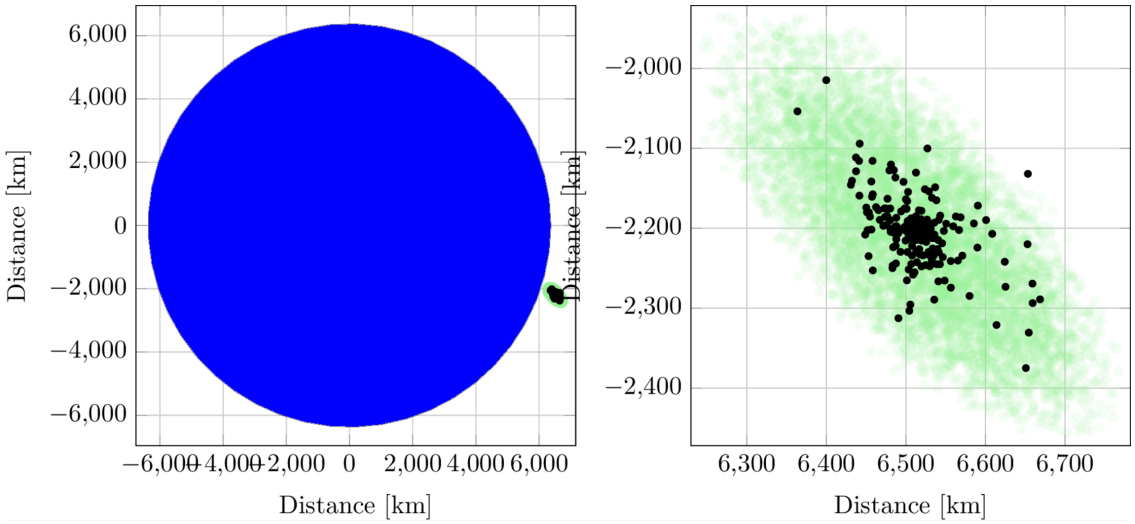

Figure 6.9: Cluster and extent estimates at the end of the fit span for the LEO test case, plotted in the centroid's perifocal frame.

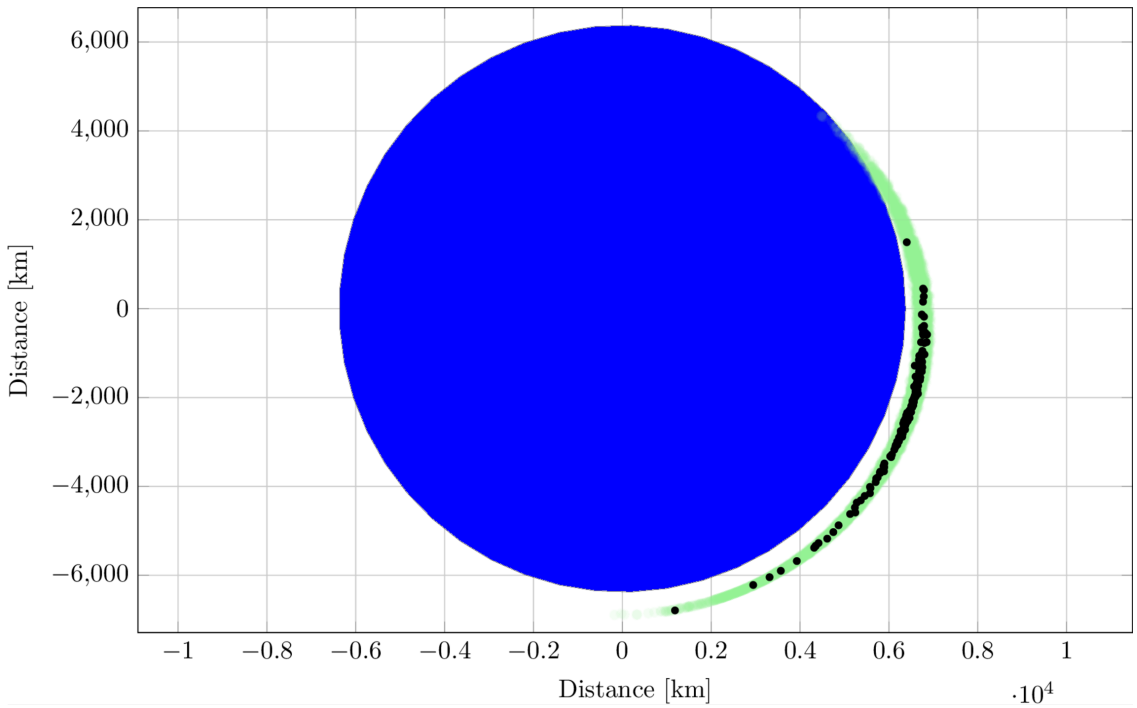

Figure 6.10: Cluster and extent estimates for the LEO test case propagated forward two periods, plotted in the centroid's perifocal frame.

6.3 GEO Test Case

Next, a breakup in geosynchronous orbit (GEO) is simulated. The pre-breakup orbit is shown in Table 6.5, and mapped in Figure 6.11. The NASA breakup model is again used to generate the debris cloud, this time with a threshold characteristic length of 15 cm, to account for the much farther detection distance. The initial distribution of 125 cluster members is shown in Figure 6.12. Range and angle measurements are simulated from a single ground station in Florida using the noise characteristics in Table 6.2 at a cadence of 30 s over 30 min. The simulation includes the effects of a 4×4 spherical harmonic Earth gravity model and solar radiation pressure.

The extent estimate is initialized using the Cartesian covariance matrix $P_{\text{cart},0} =$

Table 6.5: Pre-breakup orbit parameters for the GEO test case.

a	42,166.2 km	n	7.2916×10^{-5} rad s^{-1}
e	0.0	a_f	0.0
i	0.00147678 rad	a_g	0.0
Ω	1.53829 rad	χ	0.000738
ω	0.0 rad	ψ	2.4×10^{-5}
ν	0.0 rad	λ	3.185 rad

Area/mass ratio: 0.004 m^2/kg
Epoch: 1 Jan 2016 12:00:00.00 UTC

diag(200 km^2, 200 km^2, 200 km^2, 0.20 km/s^2, 0.20 km/s^2, 0.20 km/s^2), converted to equinoctial element space and conditioned using the method in Section 5.3.5. A heuristic tuning process was used for these simulations–several filter runs were performed while varying the tuning parameters one at a time, and filter performance (namely convergence speed and error characteristics) was noted. While this resulted in good performance for these simulations, one area of proposed future work is to develop systematic methods for selecting these parameters. The Tapley process noise covariance (see Appendix D) is used in the filter with the tuning parameter $q = 5 \times 10^{-5}$. The rest of the filter parameters used are listed in Table 6.6.

The primary projections of the estimated and true extent, along with the propagated fragments, are plotted in Figure 6.13. The quality of the estimate is evident here and Figure 6.14 shows how the fit metric improved over the pass. 12 outliers were gated out of the cluster. Unlike the LEO test case, there is no occlusion this soon after the event, and the estimate does not quite reach steady state over the 30 minute pass.

The position error and uncertainty in the centroid estimate are plotted in Figure 6.15. The error converges to less than 5 km. After dropping from the initialized

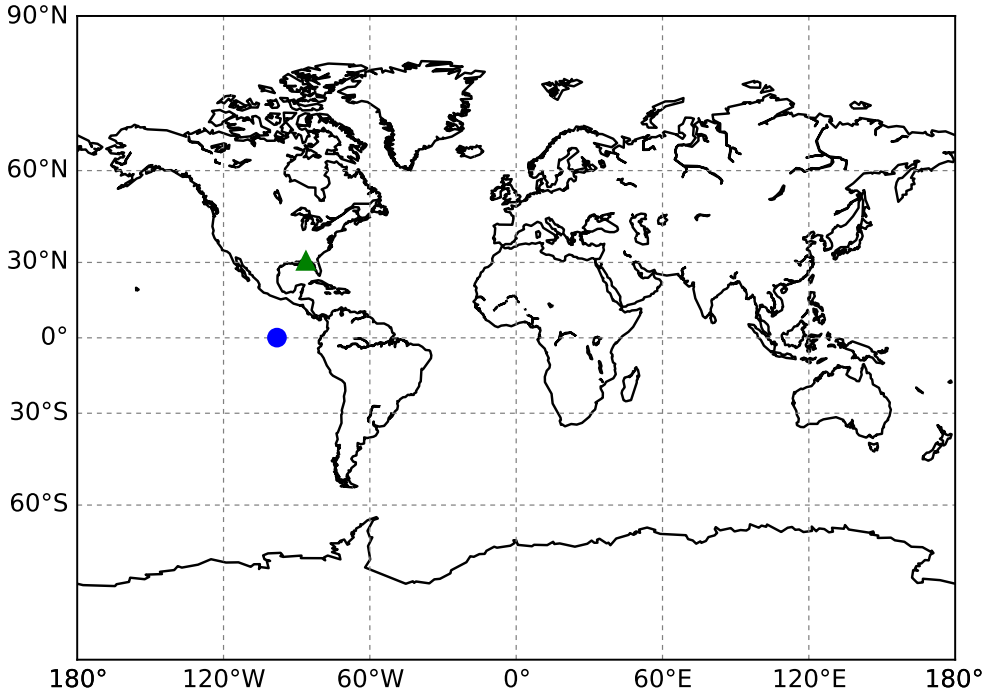

Figure 6.11: Ground track map for the GEO test case scenario. Stations are in green, and the breakup point is in blue.

Table 6.6: Filter tuning parameters used for the GEO test case.

Parameter	Value
$N_{\text{particles}}$	10,000
ν_{min}	8
ν_{max}	30
τ	86,400
β	600
g	1.25

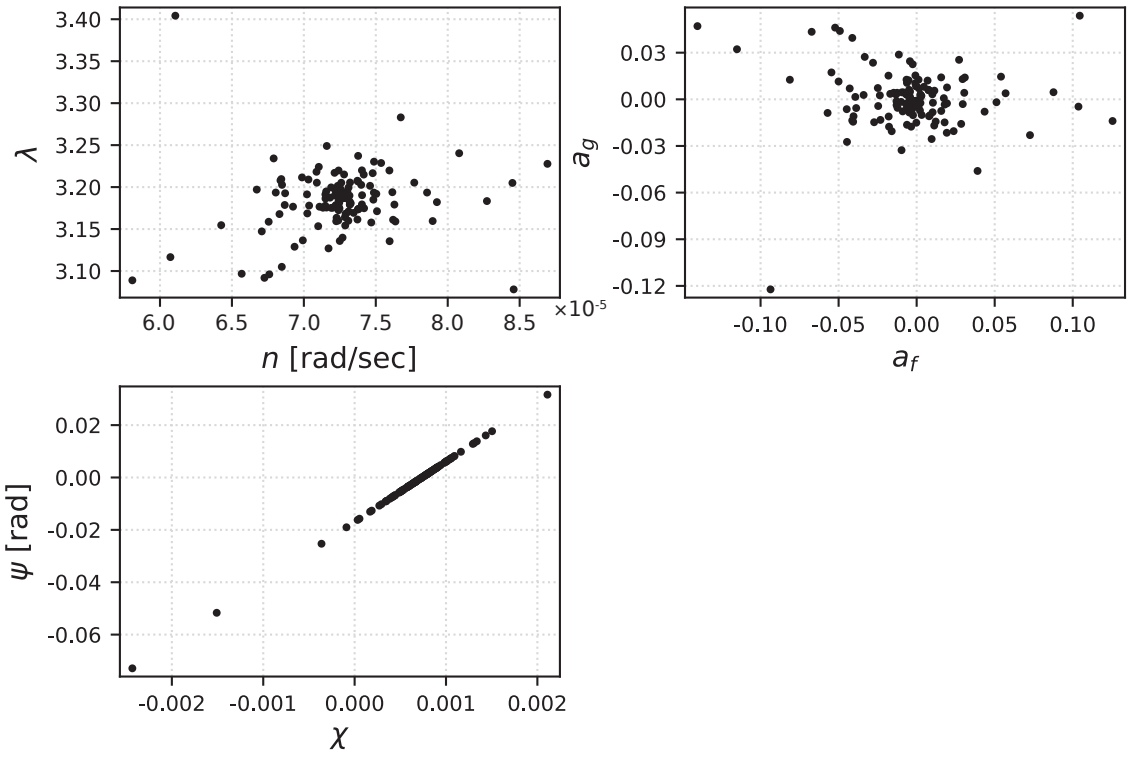

Figure 6.12: Initial distribution of cluster members for the GEO test case.

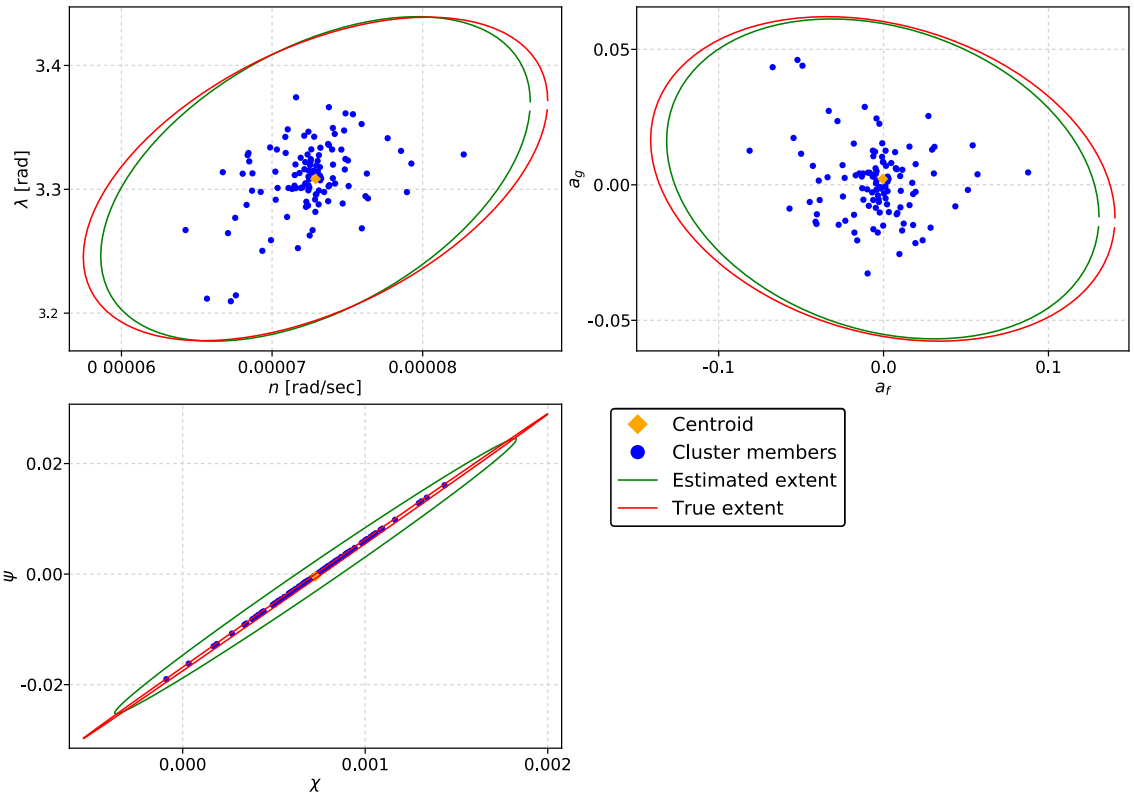

Figure 6.13: Extent estimate at the end of the pass.

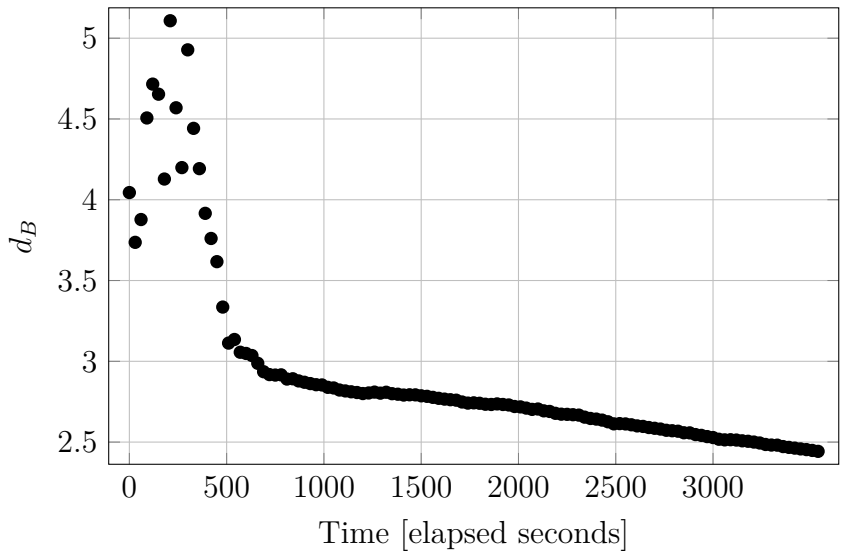

Figure 6.14: Evolution of the fit metric for the GEO test case.

value to about 2 km (1σ), the position uncertainty begins growing, a result of the process noise that is added to the model, and the growing adaptive measurement noise resulting from the expanding cluster (Section 5.3.7).

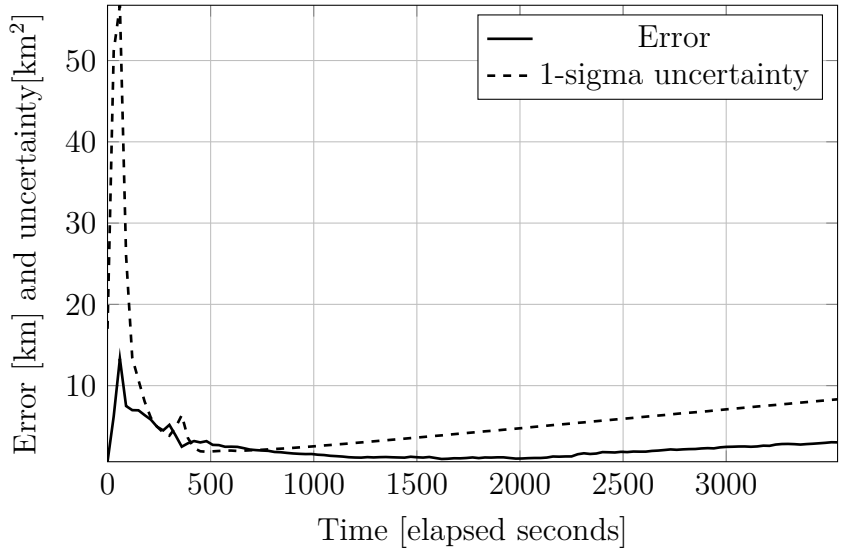

Figure 6.15: Centroid position error and uncertainty for the GEO test case over the pass.

As described above for the LEO scenario, the cluster volume predictions over two centroid orbit periods are plotted in Figure 6.16. The evolution of the cluster is

very similar to that of the LEO test case, with minima coinciding with the pinch and anti-pinch points, and maxima approximately a quarter period after the minima. These trends match the expected behavior quite well. The perifocal projections in Figures 6.17 and 6.18 also reflect the quality of the fit. After only two periods, some of the cluster members are more than 180° out of phase with one another, but the propagated extent estimate reflects this. While this scenario had no further passes to update the cluster, the more separated cluster members would subsequently be gated out of the larger cluster, as they would be easily tracked as single objects.

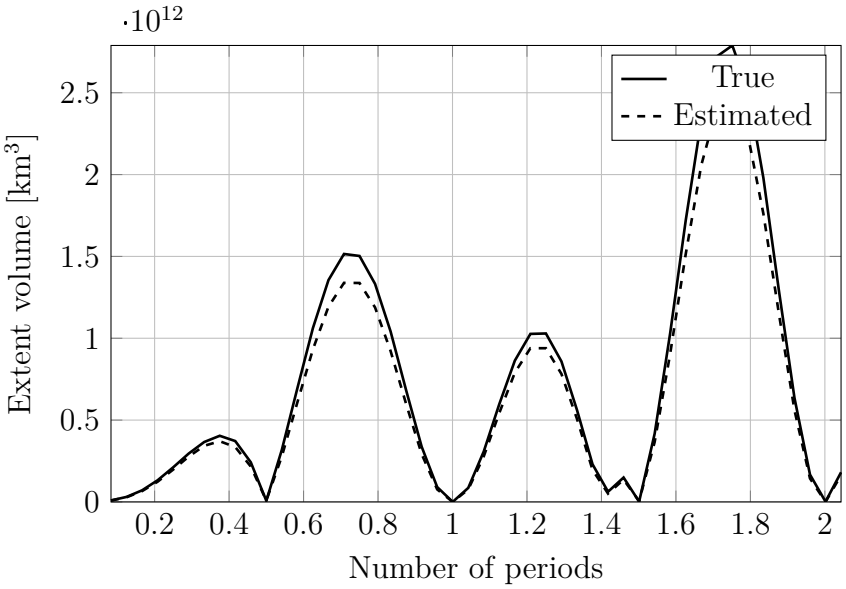

Figure 6.16: Predicted and actual volume of the satellite cluster in the GEO test case over two periods.

6.4 HEO Test Case

Finally, a breakup in a highly-elliptical orbit is simulated. The pre-breakup reference orbit is listed in Table 6.7. This is a typical geosynchronous transfer orbit, which is an intermediate orbit between the LEO parking orbit attained by the launch vehicle

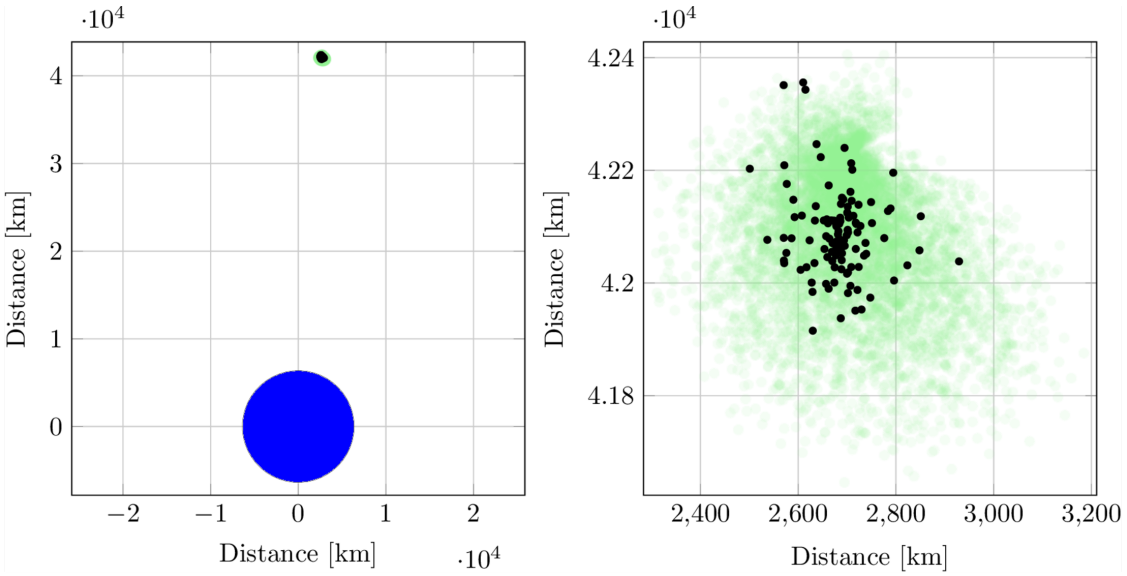

Figure 6.17: Cluster and extent estimates at the end of the fit span for the GEO test case, plotted in the centroid's perifocal frame.

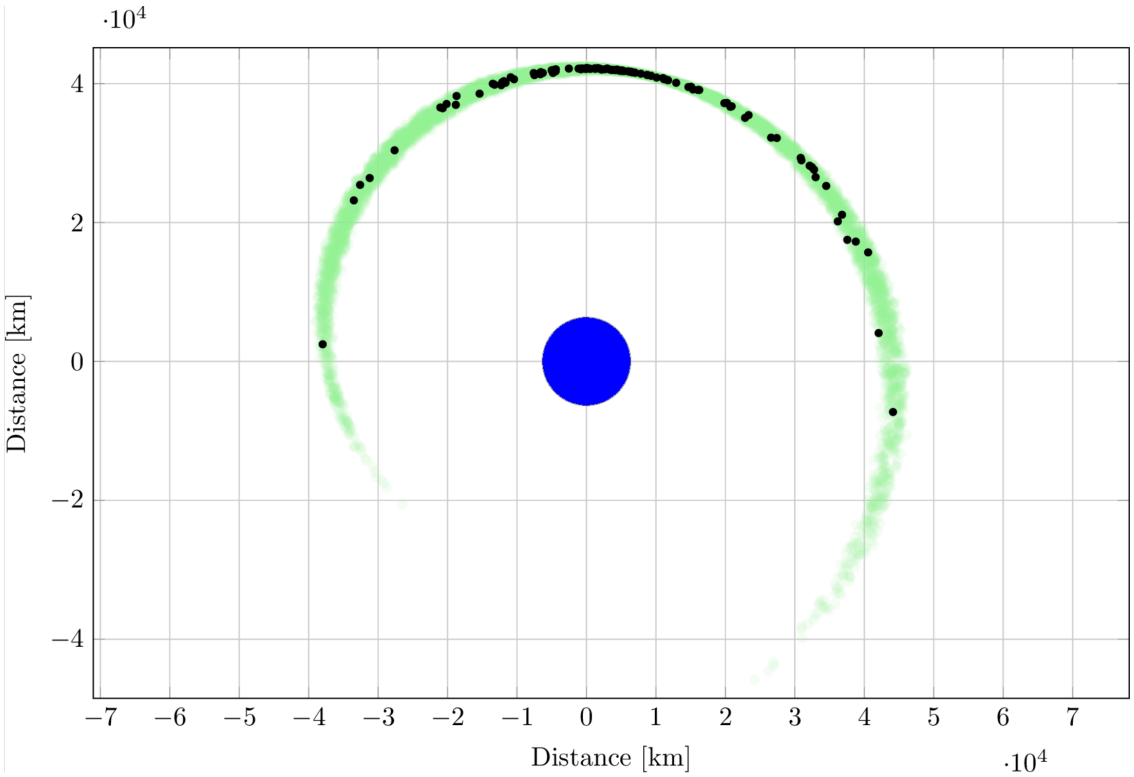

Figure 6.18: Cluster and extent estimate for the GEO test case propagated forward two periods, plotted in the centroid's perifocal frame.

Table 6.7: Pre-breakup orbit parameters for the HEO test case.

a	24,420.9 km	n	0.000165434 rad s^{-1}
e	0.72654	a_f	0.72654
i	0.174534 rad	a_g	-2.81092×10^{-6}
Ω	-7.15649×10^{-6} rad	χ	-6.26118×10^{-7}
ω	3.28759×10^{-6} rad	ψ	0.0874894
ν	2.47479 rad	λ	0.989781 rad

Area/mass ratio: 0.004 m^2/kg
Epoch: 1 Jan 2016 07:00:00.00 UTC

immediately after launch and the final GEO mission orbit. This is a uniquely difficult regime to track, due to the highly variable distances involved and the different perturbations affecting the orbit at different times–drag and higher-order Earth gravity terms near perigee, and solar radiation pressure near apogee [66].

As with previous test cases, the NASA breakup model is used to generate the initial debris cloud, using a threshold characteristic length of 25 cm, resulting in a cluster size of 43 objects. The dynamics model includes the effects of an 8×8 spherical harmonic Earth gravity model, drag (using an exponential atmosphere model [136]), and solar radiation pressure. A ground station on Ascension Island in the Atlantic Ocean records range and angles observations for all objects in the cluster every 30 s for 45 min while the cluster is about 3.5 h from apogee. Observation noise parameters are identical to previous test cases, and the other filter parameters are listed in Table 6.8. A single revolution of the pre-breakup orbit, the ground station location, and the cluster position at the beginning of the pass are plotted in Figure 6.19, and the initial distribution of cluster objects is plotted in Figure 6.20.

The resulting primary extent projections are plotted in Figure 6.21. The one obvious outlier in Figure 6.20 is gated out of the cluster, but the rest of the cluster

Table 6.8: Filter tuning parameters used for the HEO test case.

Parameter	Value
$N_{\text{particles}}$	10,000
$\nu_{\min}$	8
$\nu_{\max}$	30
τ	73,200
β	600
g	1.5

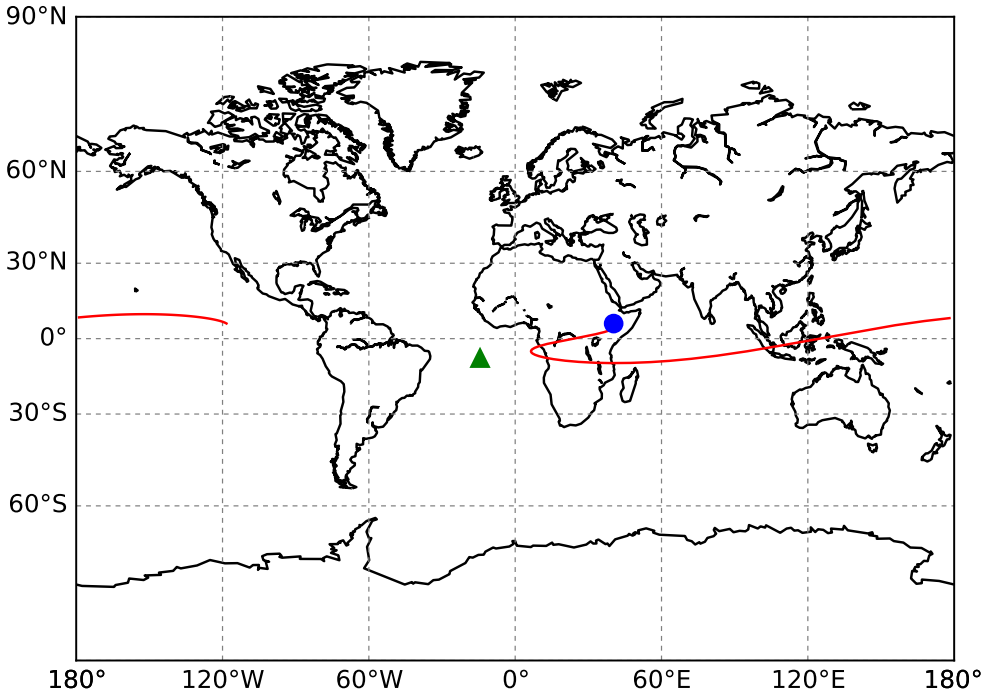

Figure 6.19: Ground track map for the HEO test case. The ground station is in green, the breakup point is in blue.

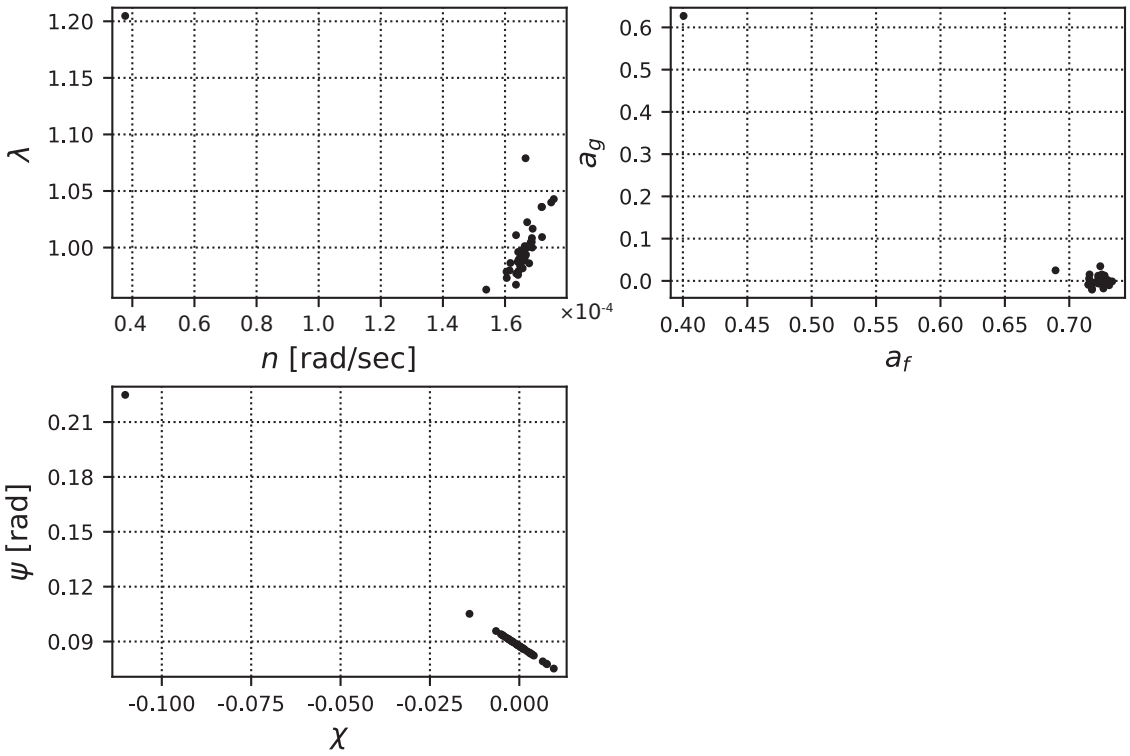

Figure 6.20: Initial distribution of cluster members in the HEO test case.

members remain. The quality of the extent estimate is evident both here and in Figure 6.22, which shows the evolution of the fit metric over the pass. There is a noticeable amount of noise in this plot early in the pass, but it settles about 12 min in, and appears to reach a steady-state value of about 1.9 approximately 25 min into the pass.

The centroid estimate error and uncertainty are plotted in Figure 6.23. The error estimate is low but biased near 1 km, while the 1σ position uncertainty collapses to nearly 1 km early in the pass before slowly increasing.

Figure 6.24 shows the true and predicted volume of the satellite cluster extent over two periods. The volume evolution is less regular than in the LEO and GEO cases due to the eccentricity of the orbit, but similar features, such as the pinch

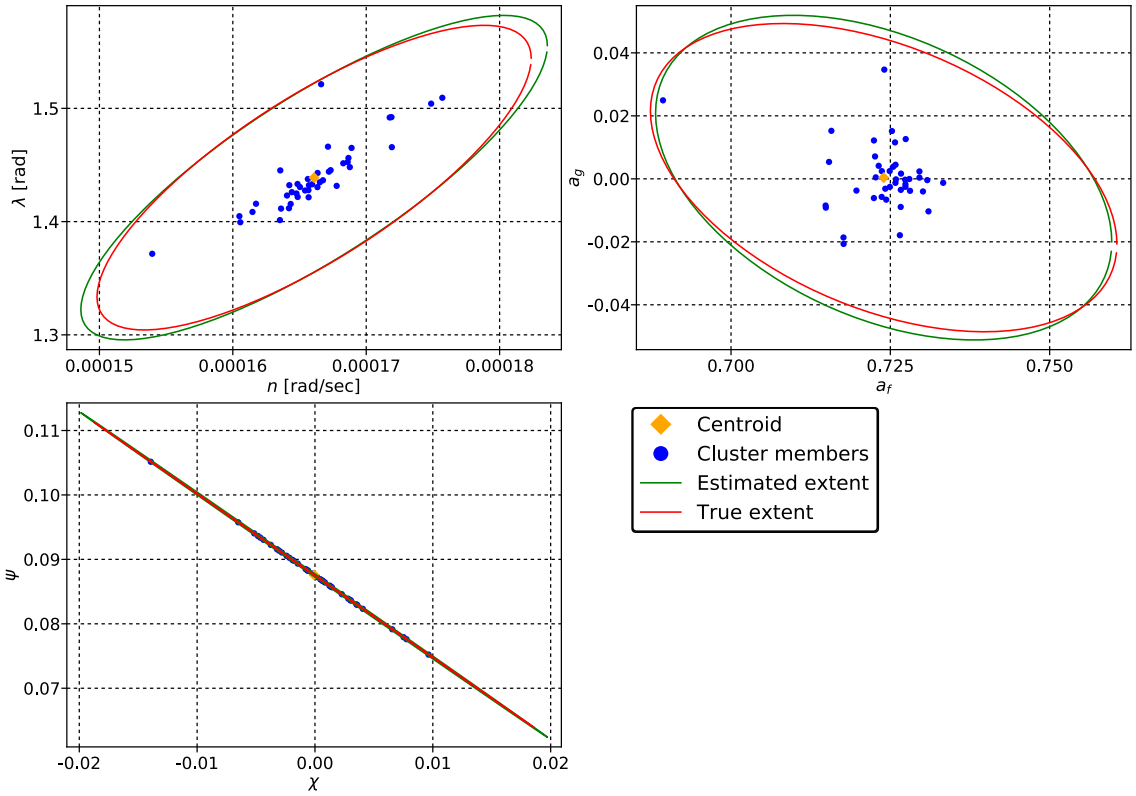

Figure 6.21: Extent estimate at the end of the pass for the HEO test case.

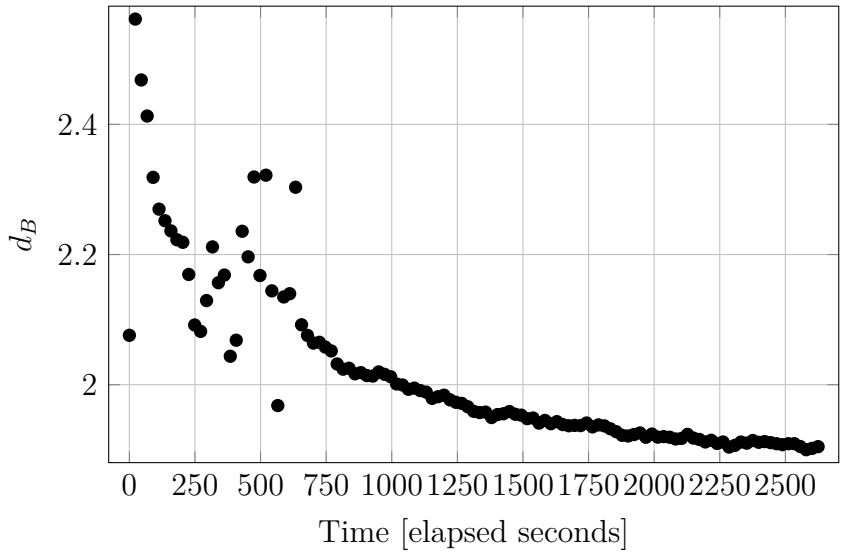

Figure 6.22: Evolution of the fit metric for the HEO test case.

point, are apparent. The maximum volume points are marked on this figure as well as Figure 6.25—they occur at the same angle away from the apogee as the pinch point. Figures 6.26 and 6.27 show the extent projections in the perifocal frame. With the exception of the single cluster member that "escapes" (and would subsequently be tracked as an individual object), the extent estimate holds over the propagation period.

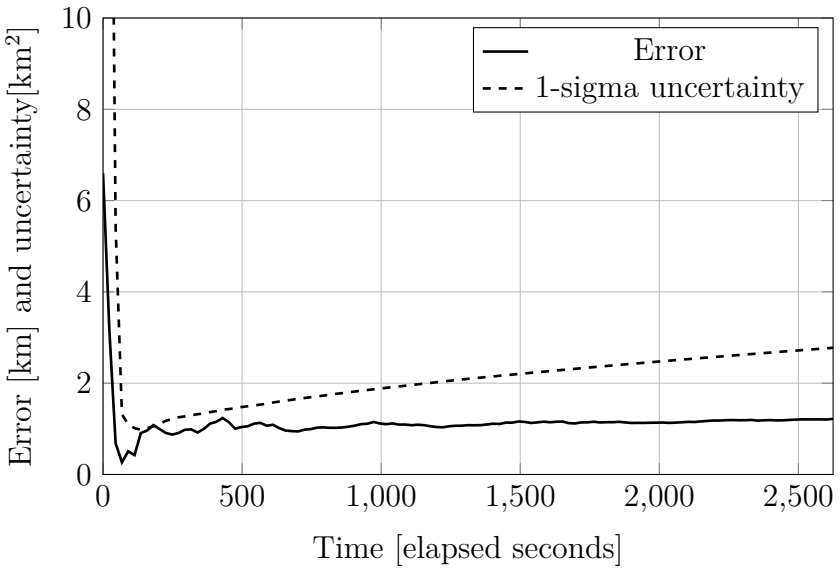

Figure 6.23: Centroid position error and uncertainty for the HEO test case over the pass.

One thing to keep in mind when considering the predicted extent plots is that the true extent is often driven by a small number of outliers, perhaps only one or two. As these plots show pure predictions, with no extent estimate updates, this manifests as large differences between the prediction and the true extents. However, subsequent passes would allow these outliers to be gated out of the cluster, allowing for a more accurate extent estimate.

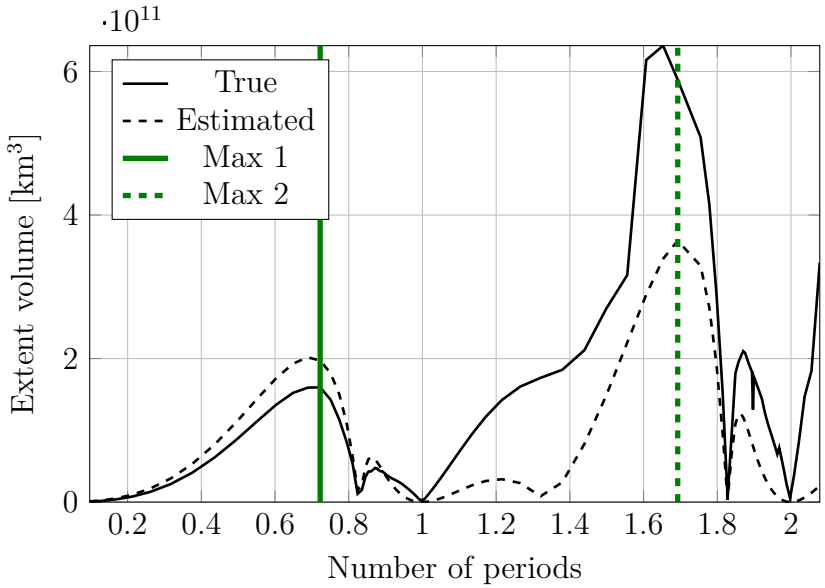

Figure 6.24: Predicted and actual volume of the satellite cluster in the HEO test case over two periods.

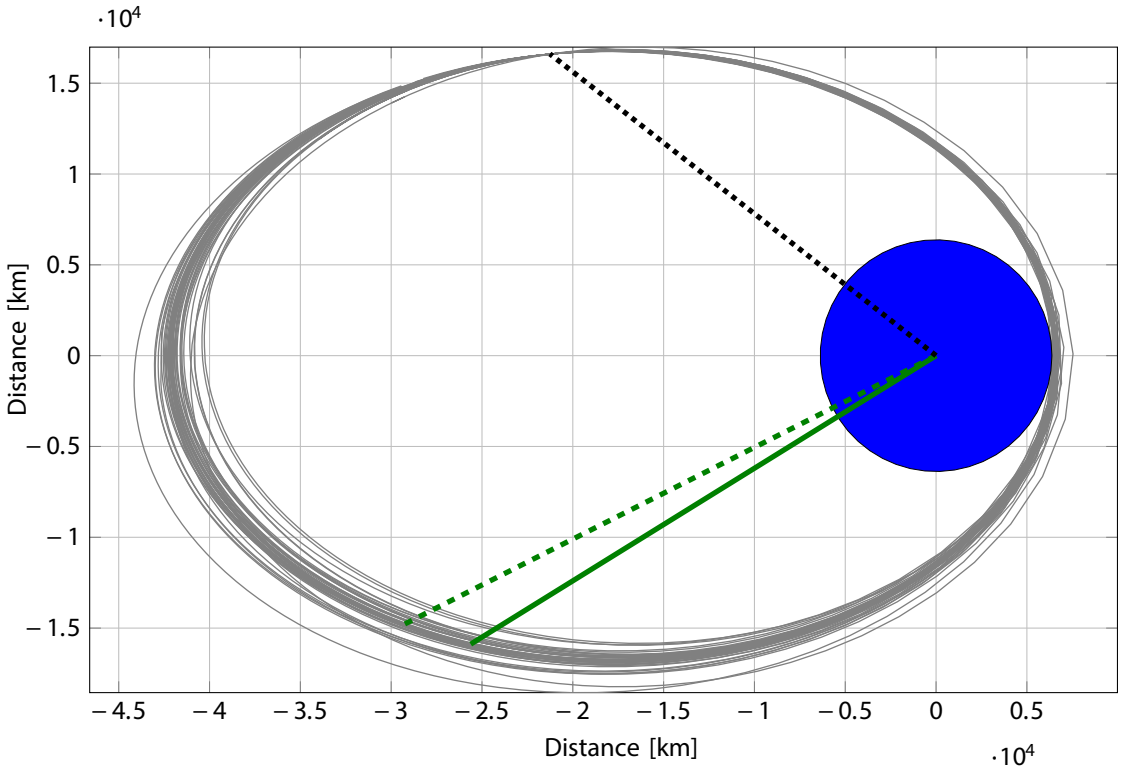

Figure 6.25: Orbit traces of all cluster members in the centroid's perifocal frame, annotated with features from Figure 6.24. Additionally, the black dotted line shows the location of the initial breakup point.

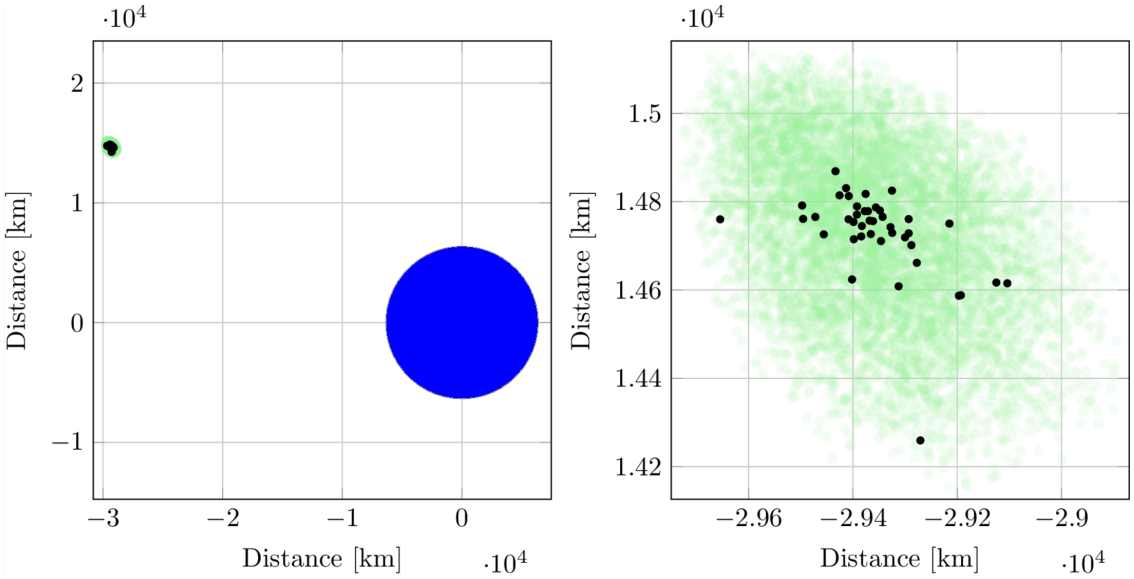

Figure 6.26: Cluster and extent estimates at the end of the fit span for the HEO test case, plotted in the centroid's perifocal frame.

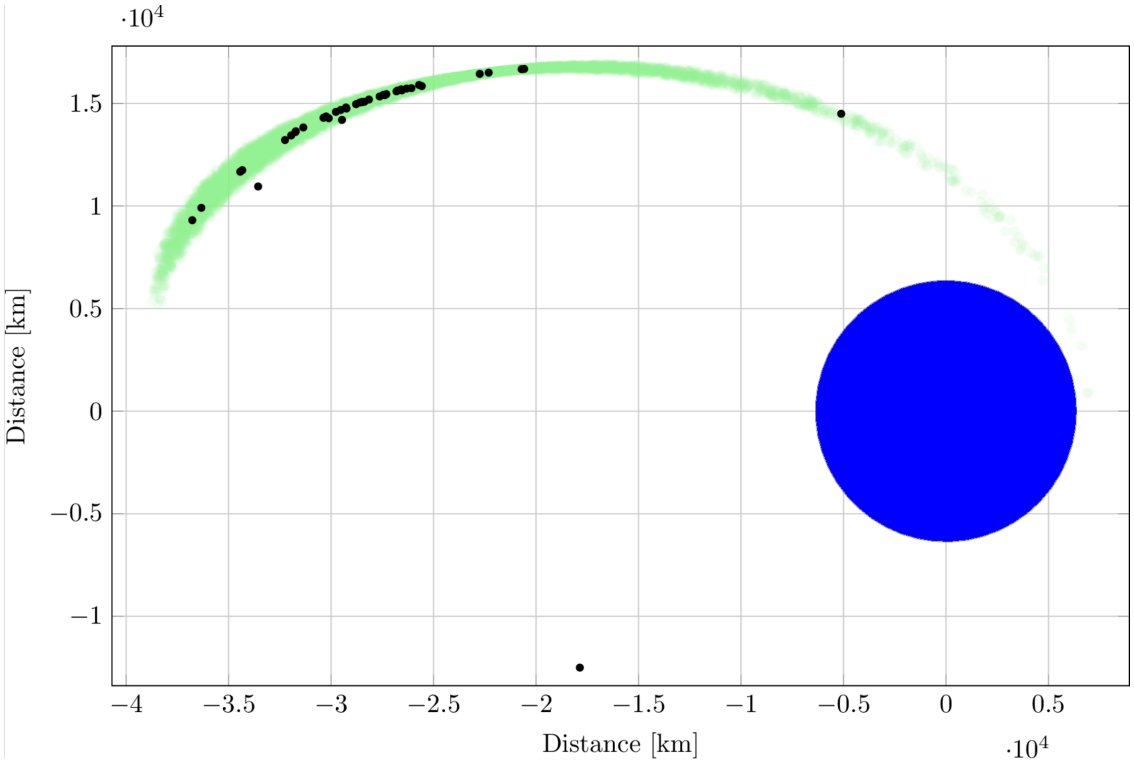

Figure 6.27: Cluster and extent estimates after being propagated two periods, plotted in the centroid's perifocal frame.

6.5 Test Case Summary

As demonstrated here, the extent particle filter method described in this work is robust enough to provide accurate extent and centroid estimates over a variety of orbits, including high-eccentricity cases, which can be particularly stressing. Centroid position error around 1 km is typical over a single pass, which is quite small when compared to the size of the cluster extent itself.

A moderate amount of manual tuning was required for the filter in these cases. Several tuning parameters, including τ, β, N, $\nu_{\min}$, and $\nu_{\max}$, had marginal effect on the performance of the filter. Meanwhile, the gate size g was critical not only in fine-tuning, but in convergence of the filter. It can easily be set too high (allowing all cluster objects to be accepted), or too low (resulting in a diverging extent estimate due to too many objects being incorrectly rejected). Notably, zero process noise was used in the LEO and HEO test cases, while a small process noise covariance was required for convergence of the GEO filter. At first, this seems at odds with the conclusions of Chapter 4, but the centroid measurement covariance in these simulations was sufficiently high to avoid divergence of the centroid filter (except in the GEO case).

The observability of the 6×6 extent matrix with three-dimensional observations is a potential concern here, particularly when given relatively short measurement arcs. For the test cases presented here, the method of initializing the extent from a spherical covariance yields good results, although this is not an indication that the system is generally observable. Using observations besides range and

bearing (angles-only telescope observations, for instance) will likely result in poorer estimates, as range information becomes observable only given sufficient geometric diversity in the measurements, requiring longer measurement arcs. In general, the observability of the cluster extent will follow single-object observability trends [191], [192].

Quantification of the quality of the extent estimate is nuanced, but the technique used here involving predicting the estimate and comparing it with the computed true extent allows for several key insights. First, the extent estimate is closest to the true extent at the pinch point, or the point along the reference orbit at which the breakup occurred. The extent estimate quality remains good at subsequent pinch point transits as well. This trend is also evident 180° away from the pinch point as well, at the anti-pinch line. The largest differences between the true and estimated extent occur 90° away from the pinch point in the circular orbit cases, which also correspond to the largest cluster extent. The extent estimates track this pattern rather well, even when using only a single pass of observations.

In the immediate aftermath of a breakup, these extent estimates may be the only quantitative description of the resulting cluster available. Precision state estimates of individual objects are not yet possible due to the limited amount of data and high probability of mis-tagged observations. The extent estimates demonstrated here enable conjunction assessments between a spacecraft of interest and a defined volume of space, rather than the latest pre-event state of the broken-up spacecraft. The estimates also allow for sensor tasking against a volume of space (or, more precisely, the projection of 3D volume into observation space), increasing the likelihood

that all pieces of a debris cloud are captured. While an accurate description of the cluster is desirable, the nature of the resulting product does not require the level of precision required for single target tracking.

Chapter 7: Conclusions

The rapidly changing space environment is necessitating changes to the way objects in space are tracked. This dissertation presented a novel method for tracking clusters of closely-spaced objects in orbit, as in the immediate aftermath of a fragmentation event such as a breakup or collision. Using established group tracking methods as a starting point, the problem was decoupled into a centroid tracking filter and an extent particle filter, allowing a cluster to be tracked by its orbit as well as its size, shape and orientation.

The problem motivation and background were provided in Chapter 1, a review of the relevant group tracking and space surveillance literature was presented in Chapter 2, and the necessary tools for modeling the motion of satellites and representing the extent as an ellipsoid are presented in Chapter 3.

Chapter 4 explored the concept of an average orbit, and the implications of computing the centroid using different orbit state representations. It was found that averaging using the equinoctial element set presented in Equation (3.1) results in a more realistic representation of an average orbit than Cartesian or classical orbital element representations, as measured by the difference between averaged observables (range and angles) and averaged states. Although some error is inherent in the fact

that the centroid does not generally follow the same equations of motion as its constituent pieces, this error is found to be small in comparison to the extent itself, and thus acceptable for our purposes (as demonstrated in Chapter 6).

Chapter 5 describes the method for sequentially estimating the extent matrix using a matrix-variate particle filter. Because the actual measurement likelihood is unknown, a suitable approximation is investigated by evaluating several potential matrix comparison metrics. By considering the qualitative behavior of the candidate metrics, then the performance of each in a representative filtering scenario, the inverse Bhattacharyya distance is chosen for the simulations in Chapter 6. More generally, it is found that metrics based on element-wise differences perform worse overall than determinant-based metrics, largely because of the difference in scale of one coordinate over the others. The consistency of using the weighted average of Inverse Wishart-distributed samples is also demonstrated in this chapter, suggesting that the particle filter will converge to the correct value given a sufficient number of samples.

Results for three test cases are presented in Chapter 6. The test cases consist of a spontaneous satellite breakup (modeled using the NASA standard breakup model) in LEO, HEO, and GEO. The filtering method developed here is shown to be robust to the choice of orbit, including high-eccentricity cases. Extent estimates sufficiently capture all members of the clusters and compare favorably to truth, both during the estimation span and when propagated forward for two orbit periods. This allows for a specific "region of concern" to be defined, predicted, and refined given more observational data. The centroid estimate, similarly, compares favorably with

the true centroid. The volume of the cluster extent is found to follow the familiar periodic behavior described by past analytical studies of debris clouds. Careful selection of the gate size and process noise covariance was found to be critical to the filter performance, in some cases making the difference between a diverged estimate and a converged one. Further, no significant issues relating to observability of the extent matrix were encountered.

This method can significantly improve spaceflight safety over the current practices. As discussed in Chapter 1, it is difficult for a ground-based space surveillance system to quickly characterize and catalog large-scale on-orbit debris events. The presented method gives operators a way to characterize such events very quickly after they occur, enabling warnings or other mitigation efforts on a much faster timeline. This will become increasingly important as the number of active payloads in Earth orbit grows in the coming years.

In addition to providing a way to characterize the potential threat to spacecraft from a fragmentation soon after the event, extent tracking could also enable a more systematic way of tracking and cataloging individual objects resulting from such events. Sensors could be tasked based on the extended target tracking information, and perhaps cued to stare at the region of space the cluster is expected to pass through, increasing the probability of detection of the entire cluster. As individual pieces drift away from the main group, estimates can be bootstrapped with information from the group tracking estimate for more rapid orbit determination and cataloging.

7.1 Limitations

While the extent filter provides a way to rapidly characterize on-orbit breakups and fragmentations without waiting for the pieces to separate into individually-tracked objects, there are several limitations to this method. The estimates are valid for a relatively short period after the event itself–primarily during the pulsating ellipsoid phase of the debris cloud [69]. This typically lasts for only a few orbit periods. It should be noted, however, that as the primary cluster spreads out, smaller clusters could be broken out and tracked in the same way (see below).

Another clear drawback of modeling a cluster in this way is that the extent matrix contains no information about the internal distribution of cluster constituents. Instead, it is assumed that the spatial density of objects inside the extent envelope is high enough to be dangerous for any other objects passing through. As is seen by the results in Chapter 6, there is a potential for the estimated extent to be overly conservative. This could be mitigated somewhat by more careful tuning of the extent filter.

Finally, the estimates resulting from this cluster tracking method are less precise than those resulting from a more typical single object orbit determination. While the latter might result in estimates with tens or hundreds of meters of uncertainty, the former is often one or two orders of magnitude less precise. This is inherent to the large cluster dimensions and the way extent is estimated, as well as the issues with centroid tracking discussed in Chapter 4. Nevertheless, the intended use cases for these tracking products is significantly different than those for

traditional orbit determination products, and thus there is still significant expected utility of these cluster estimates, as discussed above.

7.2 Future Work

Considering these limitations, and the fact that this is a new paradigm for tracking clusters of objects in space, there are myriad ways this research could be improved and extended.

Validation using real-world events, such as breakups or cubesat deployments, is obviously desirable, although obtaining the requisite observational data is often a challenge. However, the increasing availability of this data outside of the typical government channels in recent years (e.g. from commercial providers) makes it more likely that this type of study can be pursued as future work.

A multi-cluster tracking scheme could be useful as the cluster continues to spread, and a single cluster ceases to accurately represent the hazardous areas of space. The single parent cluster would split into several smaller clusters, each tracked using the technique described in this work. Parent cluster information could be used to initialize child clusters, potentially allowing for faster convergence versus initializing each child cluster with the technique described in Section 5.3.5.

Higher-order extent propagation, as in the second-order EKF[193, Section 3.7], may improve performance by more accurately predicting the evolution of the extent over time. Instead of truncating the Taylor series expansion at first order (including only the Jacobian of the dynamics function), a second-order formulation also

includes the Hessian term. This would improve the accuracy of the transformations involved in the filter, however because the second-order partial derivatives term needs to be found numerically, there is an associated computational cost in using it. This increased cost would have to be weighed against the marginal improvement to accuracy.

A more detailed study of the detection limits, noise characteristics, and pre-processing algorithms of common sensors used for space surveillance would allow for a more accurate cluster measurement model. The current model simply assumes all targets produce a point measurement, but being able to accurately model, for example, merged or missing targets would allow for more realistic simulations. A more accurate measurement simulation would better reflect the expected filter performance using real-world data.

Optimal tasking for clusters is another interesting way this research could be extended. The optimal tasking schedule for a cluster could yield some efficiencies: currently, breakups and other debris events receive very high tasking priority in order to observe all associated pieces. If the cluster is quantified early on, this could relieve some of the tasking pressure, allowing for a more focused search. The extent volume trends can also help to identify the time of the event, and thus the pinch points; observing at these pinch points has been identified as an efficient way to increase the chances of observing the entire cluster [194].

Eventually, the constituents of all clusters will drift apart enough to be trackable as point targets (unless the cluster is being actively maintained). Instead of treating each of these newly-separated targets as a new object, using the cluster

estimate to initialize the new track could reduce the number of observations needed for a sufficiently converged solution. In the group tracking spectrum defined by Drummond, Blackman, and Pretrisor [75] (see Chapter 2), this would be moving from Class 1 (cluster tracking only) to Class 4 (individual target tracks only), using cluster information to inform individual target track formation (Classes 2 and 3).

A new approach to modeling density evolution directly has recently been explored by Healy, Binz, and Kindl [195] [17], [196], [197]. This approach uses the transformation of variables technique to exactly propagate arbitrary distributions, rather than relying on linearization of the dynamics models. The estimation techniques developed in this dissertation could be coupled with this density model (termed *Eulerian orbit density*) and used to update the distribution using observational data.

Appendices

Appendix A: Spherical Harmonic Gravity Model

The partial derivatives of the spherical harmonic gravity potential function in Equation 3.13 are[126, p. 548]:

$$\frac{\partial U}{\partial r} = -\frac{\mu}{r^2}\sum_{l=2}^{\infty}\sum_{m=0}^{l}\left(\frac{R_\oplus}{r}\right)^l (l+1)P_{l,m}[\sin\phi_{gc}]\times \tag{A.1}$$

$$\{C_{l,m}\cos(m\lambda) + S_{l,m}\sin(m\lambda)\} \tag{A.2}$$

$$\frac{\partial U}{\partial \phi_{gc}} = \frac{\mu}{r}\sum_{l=2}^{\infty}\sum_{m=0}^{l}\left(\frac{R_\oplus}{r}\right)^l \{P_{l,m+1}[\sin\phi_{gc}] - m\tan\phi_{gc}P_{l,m}[\sin\phi_{gc}]\}\times \tag{A.3}$$

$$\{C_{l,m}\cos(m\lambda) + S_{l,m}\sin(m\lambda)\} \tag{A.4}$$

$$\frac{\partial U}{\partial \lambda} = \frac{\mu}{r}\sum_{l=2}^{\infty}\sum_{m=0}^{l}\left(\frac{R_\oplus}{r}\right)^l mP_{l,m}[\sin\phi_{gc}]\{S_{l,m}\cos(m\lambda) - C_{l,m}\sin(m\lambda)\}. \tag{A.5}$$

The inertial acceleration components are then [126, p. 548]

$$a_I = \left\{\frac{1}{r}\frac{\partial U}{\partial r} - \frac{r_K}{r^2\sqrt{r_I^2+r_J^2}}\frac{\partial U}{\partial \phi_{gc}}\right\} r_I - \left\{\frac{1}{r_I^2+r_J^2}\frac{\partial U}{\partial \lambda}\right\} r_J \tag{A.6}$$

$$a_J = \left\{\frac{1}{r}\frac{\partial U}{\partial r} - \frac{r_K}{r^2\sqrt{r_I^2+r_J^2}}\frac{\partial U}{\partial \phi_{gc}}\right\} r_j + \left\{\frac{1}{r_I^2+r_J^2}\frac{\partial U}{\partial \lambda}\right\} r_I \tag{A.7}$$

$$a_K = \frac{1}{r}\frac{\partial U}{\partial r}r_K + \frac{\sqrt{r_I^2+r_J^2}}{r^2}\frac{\partial U}{\partial \phi_{gc}}. \tag{A.8}$$

Appendix B: Ellipsoid properties

B.1 Geometric properties

Define the ellipsoid $\mathcal{E}$ represented by the positive definite matrix S as

$$\mathcal{E}(S) = \{\boldsymbol{y} : \boldsymbol{y}^T S^{-1} \boldsymbol{y} \leqslant 1\}, \tag{B.1}$$

The eigenvalues of S^{-1}, (λ_i), correspond to the semiaxes of the ellipsoid (a_i):

$$a_i = \lambda_i^{-1/2}, \tag{B.2}$$

and the eigenvectors correspond to the orientation.

The term *degeneracy* is used in this dissertation to describe an ellipsoid that has one axis that is very small compared to the others; in other words, $\lambda_i/\lambda_j \to 0$. In two dimensions, this has the appearance of the ellipse being stretched to nearly a straight line.

We may compute the volume of $\mathcal{E}$ as [147]

$$\mathrm{Vol}(\mathcal{E}) = \rho_n \sqrt{\det S}, \tag{B.3}$$

where ρ_n is the volume of the n-dimensional unit sphere, or the unit n-ball [198][199, Section 5.19]:

$$\rho_n = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)}. \tag{B.4}$$

B.2 Uniform sampling over the volume of an ellipsoid

Uniform sampling over the interior of an n-dimensional ellipsoid is discussed at length in [200]. The method presented here involves transforming a sample from an n-dimensional hypersphere into the ellipsoid. Although Li [201] states that this does not result in a set of uniformly-distributed points over the ellipsoid, Sun and Farooq [200] states that this assertion is incorrect, and while their proof is "omitted …for brevity", Gammell and Barfoot [202] fills in the final gap in the literature by presenting a formal proof.

Let $\mathcal{E}(S)$ represent an ellipsoid as defined in Equation B.1. Let $\boldsymbol{x}_i \sim \mathcal{N}(\mathbf{0}, \mathbb{1}_n)$ be a random vector drawn from the multivariate unit normal distribution. Then, $\boldsymbol{x}_i/||\boldsymbol{x}_i||$ is uniform on the surface of the unit hypersphere. If $u \sim \mathcal{U}(0,1)$ is uniform on $[0,1]$, then $\boldsymbol{x}_{\text{sphere}} = u^{1/n}\boldsymbol{x}_i/||\boldsymbol{x}_i||$ is uniform over the interior of the unit n-dimensional hypersphere.

Finally, if L is the lower-triangular Cholesky decomposition of S^{-1}, i.e.

$$LL^T = S^{-1}, \tag{B.5}$$

the uniform point in the hypersphere $\boldsymbol{x}_{\text{sphere}}$ can be transformed to the ellipsoid $\mathcal{E}$ via $L\boldsymbol{x}_{\text{sphere}}$. Further, if the center of the ellipsoid is somewhere other than the

origin, the point may be translated simply by adding the center vector.

Appendix C: Measurements

The measurements considered throughout this dissertation are topocentric range, azimuth, and elevation, such as from a radar [126, Section 4.3]. Given the satellite and ground station position in an Earth-fixed reference frame (ECEF), the relative position vector is $\boldsymbol{\rho}_{\text{ECEF}} = \boldsymbol{r}_{\text{sat}} - \boldsymbol{r}_{\text{site}}$. Then $\boldsymbol{\rho}_{\text{ECEF}}$ must be rotated to the topocentric SEZ frame, which is centered on the ground station, and defined by the $\hat{Z}$ unit vector along the line from the center of the Earth to the ground station, and the $\hat{E}$ unit vector defined as $\frac{\hat{K}\times\hat{Z}}{|\hat{K}\times\hat{Z}|}$, where $\hat{K}$ is the third axis in the Earth-centered inertial frame. $\hat{S}$ completes the triad. The rotation is $\boldsymbol{\rho}_{\text{SEZ}} = R_2(90^\circ - \phi_{\text{geodetic}})R_3(\lambda)\boldsymbol{\rho}_{\text{ECEF}}$, where ϕ_{geodetic} is the geodetic latitude, λ is the longitude of the ground site, and R_2 and R_3 are rotation matrices about the second and third axes, respectively.

Range is simply the magnitude of $\boldsymbol{\rho}_{\text{SEZ}}$, and azimuth and elevation may be calculated:

$$el = \arcsin\frac{\rho_Z}{\rho} \tag{C.1}$$

$$az = \text{atan2}(\rho_E, -\rho_S). \tag{C.2}$$

C.1 Partial derivatives

In order to relate these measurements to a satellite's state, the partial derivatives matrix $\frac{\partial \boldsymbol{y}}{\partial \boldsymbol{x}_{\text{eq}}}$ must be calculated, where $\boldsymbol{y}$ is a vector of range, azimuth, and elevation measurements, and $\boldsymbol{x}_{\text{eq}}$ is the equinoctial state vector. This can be decomposed into the following chain:

$$\frac{\partial \boldsymbol{y}}{\partial \boldsymbol{x}_{\text{eq}}} = \frac{\partial \boldsymbol{y}}{\partial \boldsymbol{x}_{\text{ECEF}}} \frac{\partial \boldsymbol{x}_{\text{ECEF}}}{\partial \boldsymbol{x}_{\text{ECI}}} \frac{\partial \boldsymbol{x}_{\text{ECI}}}{\partial \boldsymbol{x}_{\text{RSW}}} \frac{\partial \boldsymbol{x}_{\text{RSW}}}{\partial \boldsymbol{x}_{\text{eq}}}. \tag{C.3}$$

While these are complicated matrices, they may be found analytically, avoiding the need for finite-differencing methods. $\frac{\partial \boldsymbol{y}}{\partial \boldsymbol{x}_{\text{ECEF}}}$ is given in Vallado [126, p. 792]. An approximation of $\frac{\partial \boldsymbol{x}_{\text{ECEF}}}{\partial \boldsymbol{x}_{\text{ECI}}}$ is

$$\begin{bmatrix} T & \mathbf{0} \\ \dot{T} & T \end{bmatrix}, \tag{C.4}$$

where $T = R_3(GAST)$ and

$$\dot{T} = \omega_{\oplus} \begin{bmatrix} -\sin(GAST) & \cos(GAST) & 0 \\ -\cos(GAST) & -\sin(GAST) & 0 \\ 0 & 0 & 0 \end{bmatrix}. \tag{C.5}$$

$GAST$ is the Greenwich apparent sidereal time[203, p. 116]. $\frac{\partial \boldsymbol{x}_{\text{ECI}}}{\partial \boldsymbol{x}_{\text{RSW}}}$ is

$$\left[\begin{array}{c|c|c} \hat{R} & \hat{S} & \hat{W} \end{array} \right]^T, \tag{C.6}$$

where $\hat{R}, \hat{S}, \hat{W}$ are the radial, in-track, and cross-track unit vectors[126, p.163].

$\frac{\partial \boldsymbol{x}_{\mathrm{RSW}}}{\partial \boldsymbol{x}_{\mathrm{eq}}}$ is given in Long *et al.* [28, Section 3.3.9].

Appendix D: Process noise

This Appendix describes the process noise covariance matrix used in this dissertation.

Tapley, Schutz, and Born derive the process noise covariance matrix Q in Cartesian coordinates for perturbing accelerations following a Gaussian distribution [204, Appendix F]:

$$Q(t) = \sigma_u^2 \begin{bmatrix} \frac{1}{3}(t-t_0)^3 \mathbb{1}_{3\times 3} & \frac{1}{2}(t-t_0)^2 \mathbb{1}_{3\times 3} \\ \frac{1}{2}(t-t_0)^2 \mathbb{1}_{3\times 3} & t-t_0 \end{bmatrix}. \tag{D.1}$$

The parameter σ_u^2 is the variance of the distribution governing the perturbing acceleration, and can be used as a tuning parameter. $(t-t_0)$ is the time between measurement updates. This matrix is transformed into equinoctial element space by use of the similarity transform.

Bibliography

[1] D. King-Hele, *A Tapestry of Orbits*. Cambridge; New York, NY, USA: Cambridge University Press, 1992.

[2] National Research Council, *Navy's Needs in Space for Providing Future Capabilities*. Washington, DC: The National Academies Press, 2005. DOI: 10.17226/11299.

[3] R. Sridharan and A. F. Pensa, Eds., *Perspecitves in Space Surveillance*, ser. MIT Lincoln Laboratory Series. Cambridge, MA: The MIT Press, 2017.

[4] D. L. Pechkis, N. S. Pacheco, and T. W. Botting, "Statistical approach to the operational testing of space fence," *IEEE Aerospace and Electronic Systems Magazine*, vol. 31, no. 11, pp. 30–39, Nov. 2016. DOI: 10.1109/MAES.2016.150176.

[5] J. C. Zingarelli, E. Pearce, R. Lambour, T. Blake, C. J. R. Peterson, and S. Cain, "Improving the Space Surveillance Telescope's Performance Using Multi-Hypothesis Testing," *The Astronomical Journal*, vol. 147, no. 5, p. 111, May 2014. DOI: 10.1088/0004-6256/147/5/111.

[6] J. McDowell, *The Cubesat Explosion*, Slides, Boston, MA, Jun. 2014.

[7] J. A. Atchison and A. Q. Rogers, "Operational Methodology for Large-Scale Deployment of Nanosatellites into Low Earth Orbit," *Journal of Spacecraft and Rockets*, vol. 53, no. 5, pp. 799–810, 2016. DOI: 10.2514/1.A33361.

[8] Z. Manchester, M. Peck, and A. Filo, "KickSat: A Crowd-Funded Mission to Demonstrate the World's Smallest Spacecraft," *Proceedings of the 2013 Small Satellite Conference*, 2013.

[9] J. Radtke, C. Kebschull, and E. Stoll, "Interactions of the space debris environment with mega constellations—Using the example of the OneWeb constellation," *Acta Astronautica*, vol. 131, pp. 55–68, Feb. 2017. DOI: 10.1016/j.actaastro.2016.11.021.

[10] F. R. Hoots, R. L. Roehrich, and T. Kelso, "Spacetrack Report No. 3," Office of Astrodynamics, Aerospace Defense Center, ADC/DO6, Peterson AFB, CO, Tech. Rep., 1980.

[11] T. Kelso, "Analysis of the 2007 Chinese ASAT Test and the Impact of its Debris on the Space Environment," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2007.

[12] S. Clark, *Rocket explosion raises worries over space debris*, https://spaceflightnow.com/news/n1210/23breezem/, Oct. 2012.

[13] T. Kelso, *Breeze-M Rocket Body (38746) Breakup*, Sep. 2012. [Online]. Available: http://celestrak.com/events/38746-Breakup/.

[14] J. McDowell, *Jonathan's Space Report No. 687*, planet4589.org, Sep. 2013. [Online]. Available: http://planet4589.org/space/jsr/back/news.687.

[15] F. Letizia, C. Colombo, and H. G. Lewis, "Analytical Model for the Propagation of Small-Debris-Object Clouds After Fragmentations," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 8, pp. 1478–1491, 2015. DOI: 10.2514/1.G000695.

[16] C. McInnes, "An analytical model for the catastrophic production of orbital debris," *ESA Journal*, vol. 17, no. 4, pp. 293–305, 1993.

[17] L. Healy, S. Kindl, E. Rolfe, and C. Binz, "Structure and evolution of a debris cloud in the early phases," in *Spaceflight Mechanics 2016*, ser. Advances in the Astronautical Sciences, AAS 16-430, vol. 158, San Diego, CA: Univelt, Inc., 2016, pp. 480–504. DOI: 10.5281/zenodo.1406548.

[18] R. K. Hujsak, "Nonlinear dynamical model of relative motion for the orbiting debris problem," *Journal of Guidance, Control, and Dynamics*, vol. 14, no. 2, pp. 460–465, 1991. DOI: 10.2514/3.20660.

[19] V. A. Chobotov, D. B. Spencer, D. L. Schmitt, R. P. Gupta, R. G. Hopkins, and D. T. Knapp, "Dynamics of Debris Motion and the Collision Hazard to Spacecraft Resulting from an Orbital Breakup," Aerospace Corporation, Tech. Rep. TOR-0086A(2430-02)-1, Jan. 1988.

[20] M. J. Waxman and O. E. Drummond, "A bibliography of cluster (group) tracking," in *Signal and Data Processing of Small Targets 2004*, ser. Proceedings of SPIE, vol. 5428, 2004, pp. 551–560. DOI: 10.1117/12.548357.

[21] R. Jehn, "Dispersion of debris clouds from on-orbit fragmentation events," *ESA Journal*, vol. 15, no. 1, pp. 63–77, 1991.

[22] D. Brouwer, "Solution of the problem of artificial satellite theory without drag," *The Astronomical Journal*, vol. 64, p. 378, Nov. 1959. DOI: 10.1086/107958.

[23] Y. Kozai, "The motion of a close earth satellite," *The Astronomical Journal*, vol. 64, p. 367, Nov. 1959. DOI: 10.1086/107957.

[24] F. R. Hoots, P. W. Schumacher Jr., and R. A. Glover, "History of Analytical Orbit Modeling in the U. S. Space Surveillance System," *Journal of Guidance, Control, and Dynamics*, vol. 27, no. 2, pp. 174–185, Mar. 2004. DOI: 10.2514/1.9161.

[25] H. L. Neal, S. L. Coffey, and S. Knowles, "Maintaining the space object catalog with special perturbations," in *Astrodynamics 1997*, ser. Advances in the Astronautical Sciences, vol. 97, San Diego, CA: Univelt, Inc., 1997, pp. 1349–1360.

[26] S. L. Coffey, H. L. Neal, C. L. Visel, and P. Conolly, "Demonstration of a Special-Perturbations-Based Catalog in the Naval Space Command System," in *Space Flight Mechanics 1998*, vol. 99, San Diego, CA: Univelt, Inc., 1998, pp. 227–248.

[27] P. W. Schumacher Jr. and R. A. Glover, "Analytic Orbit Model for U.S. Naval Space Surveillance: An Overview," in *Proceedings of the U.S-Russian Second Space Surveillance Workshop*, AAS 95-427, Poznan, Poland, Jul. 1996.

[28] A. C. Long, J. O. Capellari Jr., C. E. Velez, and A. J. Fuchs, "Goddard trajectory determination system (GTDS) mathematical theory, Revision 1," NASA Goddard Space Flight Center, Tech. Rep., 1989.

[29] J. B. Dunham, A. C. Long, H. M. Sielski, and K. A. Preiss, "Onboard orbit estimation with tracking and data relay satellite system data," *Journal of Guidance, Control, and Dynamics*, vol. 6, no. 4, pp. 292–301, 1983. DOI: 10.2514/3.19832.

[30] S. Julier, J. Uhlmann, and H. F. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 477–482, Mar. 2000. DOI: 10.1109/9.847726.

[31] S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte, "A new approach for filtering nonlinear systems," in *American Control Conference, Proceedings of the 1995*, vol. 3, Jun. 1995, 1628–1632 vol.3. DOI: 10.1109/ACC.1995.529783.

[32] D.-J. Lee and K. T. Alfriend, "Sigma Point Filtering for Sequential Orbit Estimation and Prediction," *Journal of Spacecraft and Rockets*, vol. 44, no. 2, pp. 388–398, 2007. DOI: 10.2514/1.20702.

[33] T. Kelecy, M. Jah, and K. DeMars, "Application of a Multiple Hypothesis Filter to near GEO high area-to-mass ratio space objects state estimation," *Acta Astronautica*, vol. 81, no. 2, pp. 435–444, Dec. 2012. DOI: 10.1016/j.actaastro.2012.08.006.

[34] N. Singh, J. Horwood, J. Aristoff, and A. Poore, "Multiple hypothesis tracking (MHT) for space surveillance: Results and simulation studies," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, 2013.

[35] E.-S. Park, S.-Y. Park, K.-M. Roh, and K.-H. Choi, "Satellite orbit determination using a batch filter based on the unscented transformation," *Aerospace Science and Technology*, vol. 14, no. 6, pp. 387–396, Sep. 2010. DOI: 10.1016/j.ast.2010.03.007.

[36] C. Sabol, K. Hill, K. Alfriend, and T. Sukut, "Nonlinear effects in the correlation of tracks and covariance propagation," *Acta Astronautica*, vol. 84, pp. 69–80, Mar. 2013. DOI: 10.1016/j.actaastro.2012.08.023.

[37] K. Hill, K. T. Alfriend, and C. Sabol, "Covariance-based uncorrelated track association," in *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, AIAA, 2008. DOI: 10.2514/6.2008-7211.

[38] J. G. Miller, "Special perturbations uncorrelated track processing," in *Astrodynamics 2007*, ser. Advances in the Astronautical Sciences, vol. 129, San Diego, CA: Univelt, Inc., 2007. DOI: AAS07-288.

[39] D. Cappellucci, "Uncorrelated Track Processing and Identification in SPADOC 4C," in *Proceedings of the 1994 Space Surveillance Workshop*, MIT Lincoln Laboratory, Lexington, MA: 15-24, Apr. 1994.

[40] P. Schumacher Jr., "Tracking Satellite Break-ups," in *Proceedings of the U.S-Russian Second Space Surveillance Workshop*, Poznan, Poland, Jul. 1996, pp. 174–182.

[41] J. L. Junkins, "Von Karman Lecture: Adventures on the Interface of Dynamics and Control," *Journal of Guidance, Control, and Dynamics*, vol. 20, no. 6, pp. 1058–1071, 1997. DOI: 10.2514/2.4176.

[42] J. L. Junkins, M. R. Akella, and K. T. Alfriend, "Non-Gaussian error propagation in orbital mechanics," *Journal of the Astronautical Sciences*, vol. 44, no. 4, pp. 541–563, 1996.

[43] J. L. Junkins and P. Singla, "How nonlinear is it? A tutorial on nonlinearity of orbit and attitude dynamics," *Journal of the Astronautical Sciences*, vol. 52, no. 1, pp. 7–60, 2004.

[44] C. Sabol, T. Sukut, K. Hill, K. T. Alfriend, B. Wright, Y. Li, and P. Schumacher, "Linearized orbit covariance generation and propagation analysis via simple Monte Carlo simulations," in *Spaceflight Mechanics 2010*, ser. Advances in the Astronautical Sciences, AAS 10-134, vol. 136, San Diego, CA: Univelt, Inc., 2010, pp. 509–526.

[45] R. S. Park and D. J. Scheeres, "Nonlinear mapping of Gaussian statistics: Theory and applications to spacecraft trajectory design," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 6, pp. 1367–1375, 2006. DOI: 10.2514/1.20177.

[46] K. Fujimoto, D. J. Scheeres, and K. T. Alfriend, "Analytical nonlinear propagation of uncertainty in the two-body problem," *Journal of Guidance, Control, and Dynamics*, vol. 35, no. 2, pp. 497–509, 2012. DOI: 10.2514/1.54385.

[47] M. Majji, J. L. Junkins, and J. D. Turner, "A high order method for estimation of dynamic systems," *The Journal of the Astronautical Sciences*, vol. 56, no. 3, pp. 401–440, 2008. DOI: 10.1007/BF03256560.

[48] M. Majji, R. Weisman, and K. T. Alfriend, "Solution of the Liouville's equation for Keplerian motion: Application to uncertainty calculations," in *Spaceflight Mechanics 2012*, ser. Advances in the Astronautical Sciences, AAS 12-262, vol. 143, San Diego, CA: Univelt, Inc., 2012.

[49] R. M. Weisman and M. K. Jah, "Uncertainty quantification for angles-only initial orbit determination," in *Spaceflight Mechanics 2014*, ser. Advances in the Astronautical Sciences, AAS 14-434, San Diego, CA: Univelt, Inc., 2014.

[50] R. Weisman, M. Majji, and K. Alfriend, "Analytic characterization of measurement uncertainty and initial orbit determination on orbital element representations," *Celestial Mechanics and Dynamical Astronomy*, vol. 118, no. 2, pp. 165–195, 2014. DOI: 10.1007/s10569-013-9529-0.

[51] G. Terejanu, P. Singla, T. Singh, and P. D. Scott, "Uncertainty propagation for nonlinear dynamic systems using Gaussian mixture models," *Journal of Guidance, Control and Dynamics*, vol. 31, no. 6, pp. 1623–1633, 2008.

[52] D. Alspach and H. Sorenson, "Nonlinear Bayesian estimation using Gaussian sum approximations," *IEEE Transactions on Automatic Control*, vol. 17, no. 4, pp. 439–448, Aug. 1972. DOI: 10.1109/TAC.1972.1100034.

[53] K. Ito and K. Xiong, "Gaussian filters for nonlinear filtering problems," *IEEE Transactions on Automatic Control*, vol. 45, no. 5, pp. 910–927, May 2000. DOI: 10.1109/9.855552.

[54] D. Giza, P. Singla, J. Crassidis, R. Linares, P. Cefola, and K. Hill, "Entropy-based space object data association using an adaptive Gaussian sum filter," in *AIAA Guidance, Navigation, and Control Conference*, AIAA Paper # 2010-7526, American Institute of Aeronautics and Astronautics, 2010.

[55] J. T. Horwood, N. D. Aragon, and A. B. Poore, "Gaussian Sum Filters for Space Surveillance: Theory and Simulations," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1839–1851, November–December 2011. DOI: 10.2514/1.53793.

[56] K. J. DeMars, R. H. Bishop, and M. K. Jah, "Entropy-Based Approach for Uncertainty Propagation of Nonlinear Dynamical Systems," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 4, pp. 1047–1057, 2013. DOI: 10.2514/1.58987.

[57] K. J. DeMars and M. K. Jah, "Probabilistic Initial Orbit Determination Using Gaussian Mixture Models," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 5, pp. 1324–1335, 2013. DOI: 10.2514/1.59844.

[58] K. Vishwajeet, P. Singla, and M. Jah, "Nonlinear uncertainty propagation for perturbed two-body orbits," *Journal of Guidance, Control, and Dynamics*, vol. 37, no. 5, pp. 1415–1425, Apr. 2014. DOI: 10.2514/1.G000472.

[59] B. A. Jones, A. Doostan, and G. H. Born, "Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 2, pp. 430–444, 2013. DOI: 10.2514/1.57599.

[60] B. A. Jones, N. Parrish, and A. Doostan, "Postmaneuver Collision Probability Estimation Using Sparse Polynomial Chaos Expansions," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 8, pp. 1425–1437, 2014. DOI: 10.2514/1.G000595.

[61] B. A. Jones and A. Doostan, "Satellite collision probability estimation using polynomial chaos expansions," *Advances in Space Research*, vol. 52, no. 11, pp. 1860–1875, Dec. 2013. DOI: 10.1016/j.asr.2013.08.027.

[62] J. T. Horwood, J. M. Aristoff, N. Singh, and A. B. Poore, "A comparative study of new non-linear uncertainty propagation methods for space surveillance," in *Signal and Data Processing of Small Targets 2014*, ser. Proceedings of SPIE, vol. 9092, 2014, pp. 1–24. DOI: 10.1117/12.2051353.

[63] J. Horwood and A. Poore, "Gauss von Mises Distribution for Improved Uncertainty Realism in Space Situational Awareness," *SIAM/ASA Journal on Uncertainty Quantification*, pp. 276–304, Jan. 2014. DOI: 10.1137/130917296.

[64] C. R. Binz and L. M. Healy, "Association of satellite observations using Bayesian inference," in *Spaceflight Mechanics 2013*, ser. Advances in the Astronautical Sciences, AAS 13-245, vol. 148, San Diego, CA: Univelt, Inc., 2013.

[65] J. Aristoff, J. Horwood, N. Singh, C. Poore A.B. Sheaff, and M. Jah, "Multiple hypothesis tracking (MHT) for space surveillance: Theoretical framework," in *Astrodynamics 2013*, ser. Advances in the Astronautical Sciences, AAS 13-705, vol. 149, San Diego, CA: Univelt, Inc., 2013.

[66] N. L. Johnson, E. Stansbery, D. O. Whitlock, K. J. Abercromby, and D. Shoots, "History of On-orbit Satellite Fragmentations (14th Edition)," NASA Orbital Debris Program Office, Tech. Rep. TM-2008-214779, May 2008.

[67] S. Barrows, G. Swinerd, and R. Crowther, "Review of debris-cloud modeling techniques," *Journal of Spacecraft and Rockets*, vol. 33, no. 4, 1996. DOI: 10.2514/3.26798.

[68] V. Chobotov and D. Spencer, "Debris evolution and lifetime following an orbital breakup," *Journal of Spacecraft and Rockets*, vol. 28, no. 6, pp. 670–676, 1991. DOI: 10.2514/3.26298.

[69] D. McKnight, "A phased approach to collision hazard analysis," *Advances in Space Research*, vol. 10, no. 3–4, pp. 385–388, 1990. DOI: 10.1016/0273-1177(90)90374-9.

[70] J. Ashenberg, "Formulas for the phase characteristics in the problem of low-Earth-orbital debris," *Journal of Spacecraft and Rockets*, vol. 31, no. 6, pp. 1044–1049, 1994. DOI: 10.2514/3.26556.

[71] S. Blackman, *Multiple-Target Tracking with Radar Applications*. Dedham, MA: Artech House, Inc., 1986.

[72] E. Taenzer, "Tracking Multiple Targets Simultaneously with a Phased Array Radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-16, no. 5, pp. 604–614, Sep. 1980. DOI: 10.1109/TAES.1980.308926.

[73] ——, "Tracking multiple targets simultaneously with a phased array radar," in *EASCON-77*, Arlington, VA: IEEE, 1977, 10–6A–10–6R.

[74] E. H. Flad, "Tracking of formation flying aircraft," in *Proceedings of the IEEE International Radar Conference*, vol. 10, London, 1977, pp. 160–163.

[75] O. E. Drummond, S. S. Blackman, and G. C. Pretrisor, "Tracking clusters and extended objects with multiple sensors," in *Signal and Data Processing of Small Targets 1990*, ser. Proceedings of SPIE, vol. 1305, 1990, pp. 362–375. DOI: 10.1117/12.21605.

[76] M. A. Zimmer and M.-J. Tsai, "Tracking of a single cluster of closely spaced objects using one and two passive optical sensors," in *Signal and Data Processing of Small Targets 1992*, ser. Proceedings of SPIE, vol. 1698, 1992, pp. 268–280. DOI: 10.1117/12.139378.

[77] F. E. Daum and R. J. Fitzgerald, "Importance of resolution in multiple-target tracking," in *Signal and Data Processing of Small Targets 1994*, ser. Proceedings of SPIE, vol. 2235, 1994, pp. 329–338. DOI: 10.1117/12.179063.

[78] C.-Y. Yang, J.-M. Qu, S.-Y. Mao, and S. Li, "An initialization method for group tracking," in *Proceedings of the IEEE 1995 National Aerospace and Electronics Conference, 1995. NAECON 1995.*, vol. 1, May 1995, 303–308 vol.1. DOI: 10.1109/NAECON.1995.521955.

[79] P. J. Shea, "Tracking groups of detections for use in surveillance tracking problems," in *Signal Processing, Sensor Fusion, and Target Recognition XIV*, ser. Proceedings of SPIE, vol. 5809, 2005, pp. 23–32. DOI: 10.1117/12.604230.

[80] N. T. Gordon, D. J. Salmond, and D. J. Fisher, "Bayesian target tracking after group pattern distortion," in *Signal and Data Processing of Small Targets 1997*, ser. Proceedings of SPIE, vol. 3163, 1997, pp. 238–248. DOI: 10.1117/12.283961.

[81] D. Salmond and N. Gordon, "Group and extended object tracking," in *Target Tracking: Algorithms and Applications (Ref. No. 1999/090, 1999/215), IEE Colloquium On*, 1999, pp. 16/1–16/4. DOI: 10.1049/ic:19990517.

[82] D. J. Salmond and N. J. Gordon, "Group tracking with limited sensor resolution and finite field of view," in *Signal and Data Processing of Small Targets 2000*, ser. Proceedings of SPIE, vol. 4048, 2000, pp. 532–540. DOI: 10.1117/12.392015.

[83] T. Connare, E. Blasch, J. Schmitz, F. Salvatore, and F. Scarpino, "Group IMM tracking utilizing track and identification fusion," in *Proc. of the Workshop on Estimation, Tracking, and Fusion; A Tribute to Yaakov Bar Shalom*, 2001, pp. 205–220.

[84] Y. Bar-Shalom, F. Daum, and J. Huang, "The probabilistic data association filter," *IEEE Control Systems*, vol. 29, no. 6, pp. 82 –100, Dec. 2009. DOI: 10.1109/MCS.2009.934469.

[85] I. R. Goodman, R. P. Mahler, and H. T. Nguyen, *Mathematics of Data Fusion*. Norwell, MA, USA: Kluwer Academic Publishers, 1997.

[86] R. P. S. Mahler, "Detecting, tracking, and classifying group targets: A unified approach," in *Signal Processing, Sensor Fusion, and Target Recognition X*, ser. Proceedings of SPIE, vol. 4380, 2001, pp. 217–228. DOI: 10.1117/12.436950.

[87] R. Mahler, ""Statistics 101" for multisensor, multitarget data fusion," *IEEE Aerospace and Electronic Systems Magazine*, vol. 19, no. 1, pp. 53–64, Jan. 2004. DOI: 10.1109/MAES.2004.1263231.

[88] K. Gilholm, S. Godsill, S. Maskell, and D. Salmond, "Poisson models for extended target and group tracking," in *Signal and Data Processing of Small Targets 2005*, ser. Proceedings of SPIE, vol. 5913, 2005, 59130R. DOI: 10.1117/12.618730.

[89] R. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003. DOI: 10.1109/TAES.2003.1261119.

[90] S. Challa, M. R. Morelande, D. Musicki, and R. J. Evans, *Fundamentals of Object Tracking*. Cambridge, UK: Cambridge University Press, 2011.

[91] B.-N. Vo and W.-K. Ma, "The Gaussian Mixture Probability Hypothesis Density Filter," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006. DOI: 10.1109/TSP.2006.881190.

[92] K. Granström, C. Lundquist, and U. Orguner, "A Gaussian mixture PHD filter for extended target tracking," in *2010 13th Conference on Information Fusion (FUSION)*, Jul. 2010, pp. 1–8.

[93] K. Granström, C. Lundquist, and O. Orguner, "Extended target tracking using a Gaussian-mixture PHD filter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 4, pp. 3268 –3286, Oct. 2012. DOI: 10.1109/TAES.2012.6324703.

[94] D. Clark and S. Godsill, "Group Target Tracking with the Gaussian Mixture Probability Hypothesis Density Filter," in *3rd International Conference on Intelligent Sensors, Sensor Networks and Information, 2007.*, Dec. 2007, pp. 149–154. DOI: 10.1109/ISSNIP.2007.4496835.

[95] D. Li and D.-Y. Yi, "Tracking of Midcourse Ballistic Target Group on Space-based Infrared Focal Plane using GM-CPHD Filter," *Defence Science Journal*, vol. 62, no. 6, pp. 431–436, Nov. 2012. DOI: 10.14429/dsj.62.1199.

[96] F. Lian, J. Liu, M. Mallick, and C. Han, "Joint detection and estimation error bounds for an unresolved target-group using single or multiple sensors," *Digital Signal Processing*, vol. 36, pp. 26–38, Jan. 2015. DOI: 10.1016/j.dsp.2014.09.016.

[97] J. McCabe and K. DeMars, "Integrated detection and tracking for multiple space objects," in *Spaceflight Mechanics 2015*, ser. Advances in the Astronautical Sciences, vol. 155, San Diego, CA: Univelt, Inc., 2015.

[98] K. LeGrand and K. DeMars, "Space-based relative multitarget tracking," in *Spaceflight Mechanics 2015*, ser. Advances in the Astronautical Sciences, vol. 155, San Diego, CA: Univelt, Inc., 2015.

[99] K. J. DeMars, I. I. Hussein, C. Frueh, M. K. Jah, and R. S. Erwin, "Multiple-Object Space Surveillance Tracking Using Finite-Set Statistics," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 9, pp. 1741–1756, 2015. DOI: 10.2514/1.G000987.

[100] I. Hussein, K. DeMars, C. Fruh, M. Jah, and R. Erwin, "An AEGIS-FISST Algorithm for Multiple Object Tracking in Space Situational Awareness," in *AIAA/AAS Astrodynamics Specialist Conference*, American Institute of Aeronautics and Astronautics, 2012.

[101] I. I. Hussein, K. J. DeMars, C. Fruh, R. S. Erwin, and M. K. Jah, "An AEGIS-FISST integrated detection and tracking approach to Space Situational Awareness," in *2012 15th International Conference on Information Fusion (FUSION)*, Jul. 2012, pp. 2065–2072.

[102] I. I. Hussein, C. Fruh, R. S. Erwin, and M. K. Jah, "An AEGIS-FISST algorithm for joint detection, classification and tracking," in *Spaceflight Mechanics 2013*, ser. Advances in the Astronautical Sciences, AAS 13-432, vol. 146, San Diego, CA: Univelt, Inc., 2013.

[103] S. Gehly, "Estimation of Geosynchronous Space Objects Using Finite Set Statistics Filtering Methods," Ph.D. University of Colorado at Boulder, Boulder, CO, Jan. 2016.

[104] W. Koch, "On Bayesian Tracking of Extended Objects," in *2006 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Sep. 2006, pp. 209–216. DOI: 10.1109/MFI.2006.265585.

[105] J. Koch, "Bayesian approach to extended object and cluster tracking using random matrices," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 44, no. 3, pp. 1042–1059, Jul. 2008. DOI: 10.1109/TAES.2008.4655362.

[106] M. Feldmann and D. Franken, "Advances on tracking of extended objects and group targets using random matrices," in *12th International Conference on Information Fusion, 2009. FUSION '09*, Jul. 2009, pp. 1029–1036.

[107] M. Feldmann, D. Franken, and W. Koch, "Tracking of Extended Objects and Group Targets Using Random Matrices," *IEEE Transactions on Signal Processing*, vol. 59, no. 4, pp. 1409–1420, Apr. 2011. DOI: 10.1109/TSP.2010.2101064.

[108] J. Lan and X. Li, "Tracking of extended object or target group using random matrix - Part I: New model and approach," in *2012 15th International Conference on Information Fusion (FUSION)*, Jul. 2012, pp. 2177–2184.

[109] ——, "Tracking of extended object or target group using random matrix - Part II: Irregular object," in *2012 15th International Conference on Information Fusion (FUSION)*, Jul. 2012, pp. 2185–2192.

[110] K. Granström and U. Orguner, "New prediction for extended targets with random matrices," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 2, pp. 1577–1589, Apr. 2014. DOI: 10.1109/TAES.2014.120211.

[111] S. Kullback and R. A. Leibler, "On Information and Sufficiency," *The Annals of Mathematical Statistics*, vol. 22, no. 1, pp. 79–86, Mar. 1951.

[112] M. Baum, B. Noack, and U. Hanebeck, "Extended object and group tracking with Elliptic Random Hypersurface Models," in *2010 13th Conference on Information Fusion (FUSION)*, Jul. 2010, pp. 1–8. DOI: 10.1109/ICIF.2010.5711854.

[113] M. Baum and U. Hanebeck, "Shape tracking of extended objects and group targets with star-convex RHMs," in *2011 Proceedings of the 14th International Conference on Information Fusion (FUSION)*, Jul. 2011, pp. 1–8.

[114] M. Baum and U. D. Hanebeck, "Extended Object Tracking with Random Hypersurface Models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 1, pp. 149–159, Jan. 2014. DOI: 10.1109/TAES.2013.120107.

[115] L. Sun, X. Li, and J. Lan, "Modeling of extended objects based on support functions and extended Gaussian images for target tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 50, no. 4, pp. 3021–3035, Oct. 2014. DOI: 10.1109/TAES.2014.130547.

[116] M. Wieneke and W. Koch, "A PMHT Approach for Extended Objects and Object Groups," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 3, pp. 2349–2370, Jul. 2012. DOI: 10.1109/TAES.2012.6237596.

[117] S. K. Pang, J. Li, and S. Godsill, "Detection and Tracking of Coordinated Groups," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 47, no. 1, pp. 472–502, Jan. 2011. DOI: 10.1109/TAES.2011.5705687.

[118] ——, "Models and Algorithms for Detection and Tracking of Coordinated Groups," in *2008 IEEE Aerospace Conference*, Mar. 2008, pp. 1–17. DOI: 10.1109/AERO.2008.4526445.

[119] J. Huang and W. Hu, "MCMC-Particle-based group tracking of space objects within Bayesian framework," *Advances in Space Research*, vol. 53, no. 2, pp. 280–294, Jan. 2014. DOI: 10.1016/j.asr.2013.10.032.

[120] C. R. Binz and L. M. Healy, "Centroid dynamics for group object tracking," in *Spaceflight Mechanics 2015*, ser. Advances in the Astronautical Sciences, AAS 15-365, San Diego, CA: Univelt, Inc., 2015.

[121] L. Mihaylova, A. Carmi, F. Septier, A. Gning, S. Pang, and S. Godsill, "Overview of Bayesian sequential Monte Carlo methods for group and extended object tracking," *Digital Signal Processing: A Review Journal*, vol. 25, no. 1, pp. 1–16, 2014. DOI: 10.1016/j.dsp.2013.11.006.

[122] G. Gennari and G. Hager, "Probabilistic data association methods in visual tracking of groups," in *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004*, vol. 2, Jun. 2004, pp. 876–881. DOI: 10.1109/CVPR.2004.1315257.

[123] G. Duan, H. Ai, S. Cao, and S. Lao, "Group Tracking: Exploring Mutual Relations for Multiple Object Tracking," in *Computer Vision – ECCV 2012*, ser. Lecture Notes in Computer Science 7574, A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, Eds., Springer Berlin Heidelberg, 2012, pp. 129–143.

[124] L. Bazzani, M. Cristani, and V. Murino, "Decentralized particle filter for joint individual-group tracking," in *2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2012, pp. 1886–1893. DOI: 10.1109/CVPR.2012.6247888.

[125] B. Cancela, M. Ortega, A. Fernández, and M. G. Penedo, "Hierarchical framework for robust and fast multiple-target tracking in surveillance scenarios," *Expert Systems with Applications*, vol. 40, no. 4, pp. 1116–1131, Mar. 2013. DOI: 10.1016/j.eswa.2012.08.025.

[126] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*, 3rd. Hawthorne, CA: Microcosm Press, 2007.

[127] J. L. Arsenault, K. C. Ford, and P. E. Koskela, "Orbit determination using analytic partial derivatives of perturbed motion," *AIAA Journal*, vol. 8, no. 1, pp. 4–12, 1970. DOI: 10.2514/3.5597.

[128] J. R. Wright, "LS GEO IOD," *The Journal of the Astronautical Sciences*, vol. 59, no. 1-2, pp. 352–369, Jun. 2012. DOI: 10.1007/s40295-013-0022-5.

[129] O. Montenbruck and E. Gill, *Satellite Orbits: Models, Methods, and Applications*. Springer Science & Business Media, 2000.

[130] M. Walker, B. Ireland, and J. Owens, "A set of modified equinoctial orbit elements," *Celestial Mechanics*, vol. 36, no. 4, pp. 409–419, 1985. DOI: 10.1007/BF01227493.

[131] J. R. Dormand and P. J. Prince, "A family of embedded Runge-Kutta formulae," *Journal of Computational and Applied Mathematics*, vol. 6, no. 1, pp. 19–26, Mar. 1980. DOI: 10.1016/0771-050X(80)90013-3.

[132] C. Rackauckas and Q. Nie, "DifferentialEquations.jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia," *Journal of Open Research Software*, vol. 5, no. 1, May 2017. DOI: 10.5334/jors.151.

[133] D. G. King-Hele and G. E. Cook, "Refining the Earth's Pear Shape," *Nature*, vol. 246, no. 5428, pp. 86–88, Nov. 1973. DOI: 10.1038/246086a0.

[134] W. H. Press, *Numerical Recipes in FORTRAN 90: The Art of Parallel Scientific Computing*. Cambridge; New York: Cambridge University Press, 1996.

[135] NGA, *Earth Gravitational Model 2008 (EGM2008)*, earth-info.nga.mil, 2008. [Online]. Available: http://earth-info.nga.mil/GandG/wgs84/gravitymod/egm2008/.

[136] NASA, NOAA, and USAF, "U.S. Standard Atmosphere, 1976," Committee on Extension to the Standard Atmosphere, Washington, D.C., Tech. Rep., 1976.

[137] T. Kubo-oka and A. Sengoku, "Solar radiation pressure model for the relay satellite of SELENE," *Earth, Planets and Space*, vol. 51, no. 9, pp. 979–986, Jun. 2014. DOI: 10.1186/BF03351568.

[138] K. J. DeMars, R. H. Bishop, and M. K. Jah, "Space object tracking in the presence of attitude-dependent solar radiation pressure effects," in *Astrodynamics 2011*, ser. Advances in the Astronautical Sciences, vol. 142, San Diego, CA: Univelt, Inc., 2011. DOI: AAS11-582.

[139] A. Gelb, Ed., *Applied Optimal Estimation*. Cambridge, MA: MIT press, 1974.

[140] D. A. Vallado, "Covariance trasnformations for satellite flight dynamics operations," in *Astrodynamics 2003*, ser. Advances in the Astronautical Sciences, AAS 03-526, vol. 116, San Diego, CA: Univelt, Inc., 2003, pp. 373–407.

[141] G. Strang, *Linear Algebra and Its Applications*, 2nd. New York: Academic Press, 1980.

[142] S. J. Julier and J. K. Uhlmann, "Unscented filtering and nonlinear estimation," *Proceedings of the IEEE*, vol. 92, no. 3, pp. 401–422, Mar. 2004. DOI: 10.1109/JPROC.2003.823141.

[143] C. R. Binz and L. M. Healy, "Development of a satellite group tracking method," in *Spaceflight Mechanics 2016*, ser. Advances in the Astronautical Sciences, AAS 16-399, vol. 158, San Diego, CA: Univelt, Inc., 2016, pp. 1059–1076.

[144] J. L. Crassidis and J. L. Junkins, *Optimal Estimation of Dynamic Systems*, Second, ser. Chapman & Hall/CRC Applied Mathematics & Nonlinear Science. Boca Raton, FL: CRC Press, 2011.

[145] L. D. Stone, R. L. Streit, T. L. Corwin, and K. L. Bell, *Bayesian Multiple Target Tracking*, 2nd. Norwood, MA: Artech House, Dec. 2013.

[146] P. C. Mahalanobis, "On the generalized distance in statistics," *Proceedings of the National Institute of Sciences (Calcutta)*, vol. 2, pp. 49–55, 1936.

[147] M. J. Todd and E. A. Yıldırım, "On Khachiyan's algorithm for the computation of minimum-volume enclosing ellipsoids," *Discrete Applied Mathematics*, vol. 155, no. 13, pp. 1731–1744, Aug. 2007. DOI: 10.1016/j.dam.2007.02.013.

[148] M. J. Todd, *Minimum-Volume Ellipsoids: Theory and Applications*, ser. MOS-SIAM Series on Optimization. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2016.

[149] M. E. Muller, "A Note on a Method for Generating Points Uniformly on N-dimensional Spheres," *Commun. ACM*, vol. 2, no. 4, pp. 19–20, Apr. 1959. DOI: 10.1145/377939.377946.

[150] R. Harman and V. Lacko, "On decompositional algorithms for uniform sampling from n-spheres and n-balls," *Journal of Multivariate Analysis*, vol. 101, no. 10, pp. 2297–2304, Nov. 2010. DOI: 10.1016/j.jmva.2010.06.002.

[151] H. Curtis, *Orbital Mechanics for Engineering Students*, 1st. Burlington, MA: Elsevier Butterworth-Heinemann, 2004.

[152] K. V. Mardia and P. E. Jupp, *Directional Statistics*. John Wiley & Sons, Inc., 1999.

[153] K. Alfriend, S. R. Vadali, P. Gurfil, J. How, and L. Breger, *Spacecraft Formation Flying: Dynamics, Control and Navigation*. Butterworth-Heinemann, Nov. 2009.

[154] I. Sobol, "Uniformly distributed sequences with an additional uniform property," *USSR Computational Mathematics and Mathematical Physics*, vol. 16, no. 5, pp. 236–242, 1976. DOI: 10.1016/0041-5553(76)90154-3.

[155] J. S. Brauchart, J. Dick, and L. Fang, "Spatial low-discrepancy sequences, spherical cone discrepancy, and applications in financial modeling," *Journal of Computational and Applied Mathematics*, vol. 286, pp. 28–53, Oct. 2015. DOI: 10.1016/j.cam.2015.02.023.

[156] C. R. Binz and L. M. Healy, "Direct Uncertainty Estimates from Angles-Only Initial Orbit Determination," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 1, pp. 34–46, 2018. DOI: 10.2514/1.G002227.

[157] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. Storrs, CT: YBS Publishing, 2011.

[158] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, Feb. 2002. DOI: 10.1109/78.978374.

[159] S. Julier and J. Uhlmann, "New extension of the Kalman filter to nonlinear systems," in *Signal Processing, Sensor Fusion, and Target Recognition VI*, ser. Proceedings of SPIE, vol. 3068, Jul. 1997, pp. 182–193. DOI: 10.1117/12.280797.

[160] S. Särkkä, "Recursive Bayesian Inference on Stochastic Differential Equations," Doctoral Dissertation, Helsinki University of Technology, Helsinki, Finland, 2006.

[161] ——, *Bayesian Filtering and Smoothing*. New York, NY, USA: Cambridge University Press, 2013.

[162] J. Hartikainen, A. Solin, and S. Särkkä, *Optimal Filtering with Kalman Filters and Smoothers*. Espoo, Finland: Aalto University, 2011. [Online]. Available: http://becs.aalto.fi/en/research/bayes/ekfukf/documentation.pdf.

[163] R. Van Der Merwe and E. A. Wan, "The square-root unscented Kalman filter for state and parameter-estimation," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings*, IEEE, 2001, pp. 3461–3464. DOI: 10.1109/ICASSP.2001.940586.

[164] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*, ser. Chapman & Hall/CRC Monographs and Surveys in Pure and Applied Mathematics 104. Boca Raton, FL: Chapman & Hall, 2000.

[165] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*, 3rd, ser. Wiley Series in Probability and Statistics. Hoboken, NJ: Wiley, 2003.

[166] J. M. Bernardo, "Reference Analysis," in *Handbook of Statistics*, vol. 25, Amsterdam: Elsevier, 2005, pp. 17–90.

[167] D. T. Hall, M. D. Hejduk, and L. C. Johnson, "Remediating Non-Positive Definite State Covariainces for Collision Probability Estimation," in *Astrodynamics 2017*, ser. Advances in the Astronautical Sciences, vol. 162, San Diego, CA: Univelt, Inc., 2017.

[168] M. Xu, T. Ellis, S. J. Godsill, and G. A. Jones, "Visual tracking of partially observable targets with suboptimal filtering," *IET Computer Vision*, vol. 5, no. 1, pp. 1–13, Jan. 2011. DOI: 10.1049/iet-cvi.2009.0060.

[169] M. Xu and T. Ellis, "Augmented tracking with incomplete observation and probabilistic reasoning," *Image and Vision Computing*, Performance Evaluation of Tracking and Surveillance, vol. 24, no. 11, pp. 1202–1217, Nov. 2006. DOI: 10.1016/j.imavis.2005.06.004.

[170] X. P. Zhang, A. S. Khwaja, J. A. Luo, A. S. Housfater, and A. Anpalagan, "Multiple Imputations Particle Filters: Convergence and Performance Analyses for Nonlinear State Estimation With Missing Data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 8, pp. 1536–1547, Dec. 2015. DOI: 10.1109/JSTSP.2015.2465360.

[171] A. Kong, J. S. Liu, and W. H. Wong, "Sequential Imputations and Bayesian Missing Data Problems," *Journal of the American Statistical Association*, vol. 89, no. 425, pp. 278–288, 1994. DOI: 10.2307/2291224.

[172] P. D. Allison, "Handling missing data by maximum likelihood," in *SAS Global Forum*, vol. 2012, Havenford, PA, 2012.

[173] T. D. Jorgensen, K. J. Forney, J. A. Hall, and S. M. Giles, "Using modern methods for missing data analysis with the social relations model: A bridge to social network analysis," *Social Networks*, vol. 54, pp. 26–40, Jul. 2018. DOI: 10.1016/j.socnet.2017.11.002.

[174] J. P. Hoffbeck and D. A. Landgrebe, "Covariance matrix estimation and classification with limited training data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 7, pp. 763–767, Jul. 1996. DOI: 10.1109/34.506799.

[175] S. Tadjudin and D. A. Landgrebe, "Covariance estimation with limited training samples," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 37, no. 4, pp. 2113–2118, Jul. 1999. DOI: 10.1109/36.774728.

[176] D. Kincaid and E. W Cheney, *Numerical Analysis: Mathematics of Scientific Computing*, 3rd. Pacific Grove, CA: Brooks/Cole, 2002.

[177] M. Moshtaghi, T. C. Havens, J. C. Bezdek, L. Park, C. Leckie, S. Rajasegarar, J. M. Keller, and M. Palaniswami, "Clustering ellipses for anomaly detection," *Pattern Recognition*, vol. 44, no. 1, pp. 55–69, Jan. 2011. DOI: 10.1016/j.patcog.2010.07.024.

[178] J. Goldberger, S. Gordon, and H. Greenspan, "An efficient image similarity measure based on approximations of KL-divergence between two gaussian mixtures," in *Proceedings Ninth IEEE International Conference on Computer Vision*, Oct. 2003, 487–493 vol.1. DOI: 10.1109/ICCV.2003.1238387.

[179] ——, "Unsupervised image-set clustering using an information theoretic framework," *IEEE Transactions on Image Processing*, vol. 15, no. 2, pp. 449–458, Feb. 2006. DOI: 10.1109/TIP.2005.860593.

[180] S. K. Zhou and R. Chellappa, "From sample similarity to ensemble similarity: Probabilistic distance measures in reproducing kernel Hilbert space," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 6, pp. 917–929, Jun. 2006. DOI: 10.1109/TPAMI.2006.120.

[181] K. Misztal and J. Tabor, "Ellipticity and Circularity Measuring via Kullback–Leibler Divergence," *Journal of Mathematical Imaging and Vision*, vol. 55, no. 1, pp. 136–150, May 2016. DOI: 10.1007/s10851-015-0618-4.

[182] W. Förstner and B. Moonen, "A Metric for Covariance Matrices," in *Geodesy-The Challenge of the 3rd Millennium*, P. D. E. W. Grafarend, D. F. W. Krumm, and D. V. S. Schwarze, Eds., Springer Berlin Heidelberg, 2003, pp. 299–309. DOI: 10.1007/978-3-662-05296-9_31.

[183] A. Bhattacharyya, "On a measure of divergence between two statistical populations defined by their probability distributions," *Bulletin of the Calcutta Mathematical Society*, vol. 35, pp. 99–109, 1943.

[184] S. Rajasegarar, J. C. Bezdek, M. Moshtaghi, C. Leckie, T. C. Havens, and M. Palaniswami, "Measures for clustering and anomaly detection in sets of higher dimensional ellipsoids," in *The 2012 International Joint Conference on Neural Networks (IJCNN)*, Jun. 2012, pp. 1–8. DOI: 10.1109/IJCNN.2012.6252703.

[185] J. C. Bezdek, T. C. Havens, J. M. Keller, C. Leckie, L. Park, M. Palaniswami, and S. Rajasegarar, "Clustering elliptical anomalies in sensor networks," in *International Conference on Fuzzy Systems*, Jul. 2010, pp. 1–8. DOI: 10.1109/FUZZY.2010.5584464.

[186] R. N. Tamura and D. D. Boos, "Minimum Hellinger Distance Estimation for Multivariate Location and Covariance," *Journal of the American Statistical Association*, vol. 81, no. 393, pp. 223–229, Mar. 1986. DOI: 10.1080/01621459.1986.10478264.

[187] T. Steerneman, "On the total variation and Hellinger distance between signed measures; an application to product measures," *Proceedings of the American Mathematical Society*, vol. 88, no. 4, pp. 684–688, 1983. DOI: 10.1090/S0002-9939-1983-0702299-0.

[188] J. C. Liou, D. T Hall, P. H Krisko, and J. N Opiela, "LEGEND – a three-dimensional LEO-to-GEO debris evolutionary model," *Advances in Space Research*, Space Debris, vol. 34, no. 5, pp. 981–986, Jan. 2004. DOI: 10.1016/j.asr.2003.02.027.

[189] N. L. Johnson, P. H. Krisko, J. C. Liou, and P. D. Anz-Meador, "NASA's new breakup model of EVOLVE 4.0," *Advances in Space Research*, vol. 28, no. 9, pp. 1377–1384, 2001. DOI: 10.1016/S0273-1177(01)00423-9.

[190] G. Marsaglia, "Choosing a Point from the Surface of a Sphere," *The Annals of Mathematical Statistics*, vol. 43, no. 2, pp. 645–646, Apr. 1972.

[191] R. L. Scott and D. A. Ellery, "An approach to ground Based space surveillance of geostationary on-orbit servicing Operations," *Acta Astronautica*, 2015. DOI: 10.1016/j.actaastro.2015.03.010.

[192] A. M. Friedman and C. Frueh, "Determining characteristics of artificial near-Earth objects using observability analysis," *Acta Astronautica*, vol. 144, pp. 405–421, 2018. DOI: 10.1016/j.actaastro.2017.12.028.

[193] C.-B. Chang and K.-P. Dunn, *Applied State Estimation and Association*, ser. MIT Lincoln Laboratory Series. Cambridge, MA: The MIT Press, 2016.

[194] V. Agapov and A. Lapshin, "Survey and follow-up strategies used in operation of ASPOS OKP to gather observation data on GEO, HEO and MEO objects," in *1st NEO and Debris Detection Conference*, vol. 1, Darmstadt, Germany: ESA Space Safety Programme Office, 2019.

[195] L. M. Healy, C. R. Binz, and S. Kindl, "Orbital Dynamic Admittance and Earth Shadow," *The Journal of the Astronautical Sciences*, Jan. 2019. DOI: 10.1007/s40295-018-00144-1.

[196] L. Healy, S. Kindl, and C. Binz, "Spatial Density Maps from a Debris Cloud," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, Sep. 2016. DOI: 10.5281/zenodo.1409270.

[197] L. M. Healy, S. Kindl, and C. R. Binz, "Spatial density maps from a fragmentation," in *7th European Conference on Space Debris*, Darmstadt, Germany, 2017. DOI: 10.5281/zenodo.1411508.

[198] F. W. J. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, Eds., *NIST Handbook of Mathematical Functions*. New York, NY: Cambridge University Press, 2010, Print companion to the NIST Digital Library of Mathematical Functions.

[199] DLMF, *NIST Digital Library of Mathematical Functions*, http://dlmf.nist.gov/, Release 1.0.11 of 2016-06-08. Online companion to the NIST Handbook of Mathematical Functions., 2016.

[200] H. Sun and M. Farooq, "Note on the generation of random points uniformly distributed in hyper-ellipsoids," in *Proceedings of the Fifth International Conference on Information Fusion, 2002*, vol. 1, Jul. 2002, 489–496 vol.1. DOI: 10.1109/ICIF.2002.1021194.

[201] X. R. Li, "Generation of random points uniformly distributed in hyperellipsoids," in *First IEEE Conference on Control Applications, 1992*, Sep. 1992, 654–658 vol.2. DOI: 10.1109/CCA.1992.269770.

[202] J. D. Gammell and T. D. Barfoot, "The Probability Density Function of a Transformation-based Hyperellipsoid Sampling Technique," *arXiv:1404.1347 [math, stat]*, Apr. 2014. arXiv: 1404.1347 [math, stat].

[203] P. K. Seidelmann, Ed., *Explanatory Supplement to the Astronomical Almanac*. Mill Valley, CA: University Science Books, 1992.

[204] B. D. Tapley, B. E. Schutz, and G. H. Born, *Statistical Orbit Determination*. Burlington, MA: Elsevier Academic Press, 2004.