

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 2/25/08		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 3/01/2007 - 02/29/2008	
4. TITLE AND SUBTITLE Frontiers in Human Information Processing Conference				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-07-1-0311	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER	
6. AUTHOR(S) Dr. Lucy Pao				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Colorado				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME (S) AND ADDRESS (ES) USAF, AFRL/NL AF Office of Scientific Research 875 Randolph Street, Rm. 3112 Arlington, Virginia 22203				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Unlimited distribution				AFRL-OSR-VA-TR-2012- 0026	
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The problem addressed in this project is that of tracking a single satellite in the neighborhood of several other closely spaced, similar satellites where there is measurement mixing between satellites observed at a ground based sensor. This situation has been referred to as the "furball problem."					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code)

DATA ASSOCIATION ALGORITHMS FOR TRACKING SATELLITES

Award No. FA9550-07-1-0311

Project Duration: March 1, 2007 to November 30, 2007

Award amount: \$50,000

Final Report

February 5, 2008

Lucy Y. Pao

University of Colorado at Boulder

Under this 9-month seed grant, we investigated several data association approaches for satellite tracking. Throughout our project, we have interacted with a number of researchers and directors of AFRL's Space Vehicles Directorate. In particular, we have regularly discussed our research progress with Dr. T. Alan Lovell, and we have also met with Drs. R. Scott Erwin L. "Robbie" Robertson, and Frank Chavez to provide updates on our work. A summary of the research results achieved under this grant is detailed in the attached technical report.

PUBLICATIONS

The following papers, funded fully or partially under this award, have been submitted or accepted for publication.

- [A] M. J. Travers, T. D. Murphey, and L. Y. Pao. "Data Association with Ambiguous Measurements," *Proc. American Control Conf.*, Seattle, WA, June 2008, in press.
- [B] M. J. Travers, T. D. Murphey, and L. Y. Pao. "Suboptimal Bayesian Filtering and Data Association for Systems with Measurement Ambiguities," submitted in Jan. 2008 for publication in the *Proc. Robotics: Science and Systems Conf.*, Zurich, Switzerland, June 2008.
- [C] R. M. Powers and L. Y. Pao. "Finite State Machine Analysis of Paired Kolmogorov-Smirnov Tests for Track-Loss Detection," pre-print, to be submitted in Mar. 2008 for publication in the *Proc. IEEE Conf. Decision and Control*, Cancun, Mexico, Dec. 2008.
- [D] R. M. Powers and L. Y. Pao. "Paired Kolmogorov-Smirnov Tests for Track-Loss Detection in the Absence of Truth Data," pre-print, to be submitted in Feb. 2008 for publication in the *IEEE Trans. Aerospace Electronic Systems*.

PROMOTION AND TENURE PROGRESS

Lucy Pao, the P.I. of this project, was promoted to Full Professor in 2004. She has involved Assistant Professor Todd Murphey in this project. Todd Murphey is under Reappointment review this academic year and his review has been progressing smoothly.

20120918133

HONORS AND AWARDS AND SYNERGISTIC ACTIVITIES

During the course of this project, we have received the following awards and have participated in the following activities and positions in various professional societies:

- L. Y. Pao, elected member of the IEEE Control Systems Society, 2005-2007.
- L. Y. Pao, Dean's Scholar, 2005-2008, College of Engineering, University of Colorado at Boulder.
- L. Y. Pao, appointed in April 2007 by the International Federation of Automatic Control (IFAC) to serve on the 2008 IFAC Triennial World Congress Young Author Prize Selection Committee.
- L. Y. Pao, appointed in July 2007 by the American Automatic Control Council to be the General Chair for the 2013 American Control Conference.
- L. Y. Pao, appointed in Aug. 2007 by IFAC to serve on the Organizing Committee for the 2010 IFAC Symposium on Mechatronic Systems.
- L. Y. Pao, appointed in Sept. 2007 by the Colorado Renewable Energy Collaboratory to be the Wind Energy Director at the University of Colorado at Boulder.
- L. Y. Pao, Associate Editor for the International Journal of Control, Automation, and Systems, 2003– .
- L. Y. Pao, selected in Dec. 2007 to receive a Visiting Miller Professorship for research at the University of California at Berkeley during 2008-2009.
- L. Y. Pao, Dean's Faculty Fellowship (2008), awarded by the College of Engineering at the University of Colorado at Boulder.
- L. Y. Pao, IEEE Control Systems Society Distinguished Lecturer, 2008-2011.
- T. D. Murphey, Bruce Holland Excellence in Teaching Award (2007), awarded by the Electrical and Computer Engineering Department at the University of Colorado at Boulder.
- T. D. Murphey, National Science Foundation CAREER Award, "Planning and Control for Overconstrained Mechanisms," 2006-2011.
- L. Y. Pao and T. D. Murphey have given seminars at University of Illinois, California Institute of Technology, University of Southern California, Pennsylvania State University, Georgia Institute of Technology, University of Pennsylvania, and University of California at Berkeley.
- L. Y. Pao and T. D. Murphey have also organized and chaired sessions and workshops at the IEEE Robotics and Automation Conference and the American Control Conference.

Data Association Algorithms for Tracking Satellites

Matthew Travers, Todd Murphey, and Lucy Pao

Award No. FA9550-07-1-0311

Final Report

Project Duration: March 1, 2007 - November 30, 2007

February 5, 2008

1 Introduction

The problem addressed in this project is that of tracking a single satellite in the neighborhood of several other closely spaced, similar satellites where there is measurement mixing between satellites observed at a ground based sensor. This situation has been referred to as the “furball problem.”

Tracking a single object or a group of objects in formation are both well studied and understood processes. What makes the “furball problem” unique is that the problem of tracking a single object in a formation is a generally undeveloped capability, especially with the type of sensor considered for this particular case. The ground based sensor considered takes a single measurement per sampling period. The single measurement can potentially be from any of the satellites in the sensor’s measurement region. This implies that the sensor returns a single stream of measurements that are “mixed” in the sense that the satellite from which the measurement originated is not known to the sensor.

The purpose of this report is to formally address the “furball problem” and present several effective data association methods for finding a viable solution. As will be shown, these data association techniques can be used to effectively sort out measurements and provide reliable tracking.

The report is structured in the following manner. The problem is formally defined and the various data association techniques used to attack this problem are briefly presented in Sections 2 and 3. Section 4 provides a comparison of the computational complexity between existing methods of tracking through batch processes and iterative methods described in Section 3. In Section 5 we provide an in depth description of the types of orbits being considered along with a description of the sensor based frames from which the measurements are taken. Simulation results are presented in Section 6, comparing the performances of the algorithms discussed in Section 3. Finally, some conclusions and possible areas of future work are discussed in Section 6.

2 Problem Definition

The exact problem considered in this report is a system consisting of j satellites in either the same orbit or in orbits very close to each other. For our purposes, two orbits being “close” can be defined as keeping all the orbital parameters the same for each orbit while perturbing the eccentricity of one by a small amount.

A ground based sensor is being used in an attempt to track a single one of the j closely spaced satellites. The model used to describe the dynamics of the satellites written down in an earth centered inertial (ECI) coordinate frame is

$$x_i(k) = f[k, x_i(k-1)] + v_i(k), i = 1, 2, \dots, j \quad (1)$$

where f contains the dynamics of the system, and the additive process noise $v_i(k)$ is zero-mean and white with covariance $Q(k)$. Since maneuvering is not considered in this initial work, the dynamics contain no control vector. The measurement due to each satellite is

$$z_i(k) = h[k, x_i(k-1)] + w_i(k), i = 1, 2, \dots, j \quad (2)$$

where the additive noise $w_i(k)$ is also assumed to be zero-mean and white with covariance $R(k)$, and the function h contains the coordinate transformation between ECI, South/East/Z (SEZ), and Range/Azimuth/Elevation

(R/Az/El) coordinates. The ECI coordinates are Cartesian coordinates such that the x -axis is aligned with the vernal equinox, the z -axis is straight up through the north pole, and the y -axis is 90 degrees away from the x -axis in the equatorial plane such that the coordinates are right-handed. The SEZ coordinates are Cartesian coordinates based at the latitude and longitude position of whatever sensor is being used to take measurements, the x -axis aligned with due South, the y -axis aligned with the East, and the z -axis pointing directly out tangent to the surface of the earth. The range, azimuth, and elevation are measured from the SEZ frame. The range is a radial position measurement from the sensor, the azimuth is the angular measure from true north ($-x$ in SEZ), and the elevation is the angular measure such that the tangent plane at the sensor has measure zero and directly above (z in SEZ) has measure 90 degrees.

The measurements for this system are taken in R/Az/El coordinates and a single measurement is taken at each time step where the origin of that measurement is not known. For example, consider two satellites in closely spaced orbits, where we are trying to track satellite 1. Five measurements are taken, one of which is randomly from satellite 2. One possibility for how the list of measurements recorded by the sensor appears is

$$\{z(k), z(k+1), z(k+2), z(k+3), z(k+4)\} = \{z_1(k), z_1(k+1), z_1(k+2), z_2(k+3), z_1(k+4)\}.$$

Only the set of measurements (with unknown origin) is provided. The actual identity of individual measurements on the right hand side is not known. The data association algorithm must sort out which measurement in the set originated from the satellite of interest.

The above example can be used to introduce uniformity index (UI), a parameter in this study. The uniformity index is the percentage of the total measurements that are assumed to be from the correct satellite of interest. In the example, one of the five measurements is always assumed to be wrong, though its position remains unknown. Thus, the UI for this case is 80%.

Due to the fact that the sensor is considered to be on the surface of the earth, in reality there will clearly be times when the two satellites moving in close proximity will be “out of view.” This problem will be touched upon in the future work section, but for most of this report will be ignored. Thus for this report, it is assumed that there is only a single sensor being used to track satellites through an entire orbit.

3 Algorithms

3.1 The Batch Filter

The Batch Filter is a widely used filtering technique that is often applied to systems where there is no mixing of measurements and where estimation does not need to be done in an online manner. The Batch Filter is used in these instances because it will in general converge to zero RMS error. Section 4 on computational complexity will illustrate a reason the Batch Filter is not always preferable.

The implementation of the Batch Filter is as follows. Measurements from a single satellite are taken. The continuous time dynamics for the system are assumed to be known perfectly and can be written

$$\dot{x}(t) = f(t, x(t)). \quad (3)$$

Then for k measurements taken at discrete time steps dt , the continuous dynamics are used to solve backwards in time k time steps to obtain k estimates of the state at time zero, i.e., $(x_1(0), x_2(0), \dots, x_k(0))$ where $x_1(0)$ is obtained by integrating backwards from the first measurement by dt , $x_2(0)$ by integrating backwards from the second measurement by $2 \cdot dt$, and so on up to k . These k estimates of the initial state are then averaged to obtain a “best estimate,” $\bar{x}(0)$, of the state at $t = 0$. The dynamics (3) can again be used, this time to integrate forward in time, starting from the $\bar{x}(0)$ and integrating forward in time to time k . This value is then the current estimate of state. This process is repeated for time $k + 1$ up to the final measurement. In principle, if the noise associated with the measurements is white with zero mean, the batch estimates should eventually converge to zero RMS error as k goes to infinity.

3.2 The Extended Kalman Filter

The EKF is a well known filtering method for nonlinear systems [1]. The EKF is an estimation algorithm that maps the current estimate of the state, $\hat{x}(k-1|k-1)$, forward one time step. The current estimate is

conditioned on the current state given the current measurement [1]

$$\hat{x}(k-1|k-1) \triangleq E[x(k-1)|Z^{k-1}]$$

where

$$Z^{k-1} \triangleq \{z(j), j = 1, \dots, k-1\}$$

is the cumulative set of measurements up to time $k-1$. The associated state error covariance matrix is

$$P(k-1|k-1) \triangleq E\{[x(k-1) - \hat{x}(k-1|k-1)][x(k-1) - \hat{x}(k-1|k-1)]' | Z^{k-1}\}. \quad (4)$$

We proceed with finding $\hat{x}(k|k-1)$, the prediction, by expanding the nonlinear function (1). The expansion is accomplished by evaluating the Taylor series of (1) about the current estimate: $\hat{x}(k-1|k-1)$

$$\begin{aligned} x(k) = & f[k, \hat{x}(k-1|k-1)] + f_x(k-1)[x(k-1) - \hat{x}(k-1|k-1)] + 1/2 \sum_{i=1}^{n_x} e_i [x(k-1) - \hat{x}(k-1|k-1)]' \\ & \cdot f_{xx}^i(k-1)[x(k-1) - \hat{x}(k-1|k-1)] + \text{higher-order terms} + v(k-1) \end{aligned} \quad (5)$$

where e_i is the i^{th} Cartesian basis vector and $f_x(k-1)$ is the Jacobian of the vector f

$$f_x(k-1) \triangleq [\nabla_x f'(k-1, x)]'_{x=\hat{x}(k-1|k-1)}.$$

Similarly, the Hessian of the i^{th} component of f is

$$f_{xx}^i(k-1) \triangleq [\nabla_x \nabla_x' f^i(k-1, x)]_{x=\hat{x}(k-1|k-1)}.$$

The prediction of the state at time k is then obtained by taking the expectation of the Taylor series expansion (5) conditioned on Z^{k-1} and neglecting higher-order terms:

$$\hat{x}(k|k-1) = f[k, \hat{x}(k-1|k-1)] + 1/2 \sum_{i=1}^{n_x} e_i \text{Tr}[f_{xx}^i(k-1)P(k-1|k-1)].$$

The state prediction error is

$$\begin{aligned} \tilde{x}(k|k-1) = & f_x(k-1)\tilde{x}(k-1|k-1) + 1/2 \sum_{i=1}^{n_x} e_i [\tilde{x}'(k-1|k-1)f_{xx}^i(k-1)\tilde{x}(k-1|k-1) \\ & - \text{Tr}[f_{xx}^i(k-1)P(k-1|k-1)]] + v(k-1) \end{aligned}$$

and its covariance is

$$\begin{aligned} P(k|k-1) \triangleq & E[\tilde{x}(k|k-1)\tilde{x}'(k|k-1)|Z^{k-1}] = f_x(k-1)P(k-1|k-1)f_x' + 1/2 \sum_{i=1}^{n_x} \sum_{j=1}^{n_x} e_i e_j' \text{Tr}[f_{xx}^i(k-1) \\ & \cdot P(k-1|k-1)f_{xx}^j(k-1)P(k-1|k-1)] + Q(k-1). \end{aligned}$$

The measurement prediction is

$$\hat{z}(k|k-1) = h[k, \hat{x}(k|k-1)] + 1/2 \sum_{i=1}^{n_x} e_i \text{Tr}[h_{xx}^i(k)P(k|k-1)].$$

The innovation is

$$\begin{aligned} \nu(k) &= z(k) - \hat{z}(k|k-1) \\ &= h[k, \tilde{x}(k|k-1)] + w(k), \end{aligned}$$

and the innovation's covariance is

$$S(k) \triangleq h_x(k)P(k|k-1)h'_x(k) + 1/2 \sum_{i=1}^{n_z} \sum_{j=1}^{n_z} e_i e'_j [h_{xx}^i(k)P(k|k-1) \cdot h_{xx}^j(k)P(k|k-1)] + R(k)$$

where

$$h_x(k) \triangleq [\nabla_x h'(k, x)]'_{x=\hat{x}(k|k-1)}$$

and

$$h_{xx}^i(k) \triangleq [\nabla_x \nabla'_x h^i(k, x)]'_{x=\hat{x}(k|k-1)}.$$

The EKF gain is

$$W(k) \triangleq P(k|k-1)h'_x(k)S^{-1}(k).$$

The filtered estimate is then

$$\hat{x}(k|k) = \hat{x}(k|k-1) + W(k)\nu(k) \quad (6)$$

and the filtered estimate error covariance is

$$\begin{aligned} P(k|k) &= P(k|k-1) - P(k|k-1)h'_x(k)S^{-1}(k)h_x(k)P(k|k-1) \\ &= [I - W(k)h_x(k)]P(k|k-1). \end{aligned}$$

3.3 Gating

Gating is a method of determining whether or not it is likely that a particular measurement came from the object being tracked. Since we are assuming that the measurement at each time step is normally distributed around the truth measurement, it is possible to define a region in the measurement space where there is a high probability of finding the measurement

$$\tilde{V}_k(\gamma) \triangleq \{z : \nu'(k)S^{-1}(k)\nu(k) \leq \gamma\}. \quad (7)$$

The probability of gating the truth measurement is

$$P_g = \int_0^\gamma \frac{\gamma^{n_z/2-1} \exp\{-\gamma/2\}}{2^{n_z/2} \Gamma(n_z/2)} d\gamma, \gamma \geq 0$$

where $\Gamma(\cdot)$ is the gamma function and n_z is the dimension of the measurement space. Because the measurement error is assumed Gaussian, this norm is chi-square distributed. If the measurement error is not Gaussian, such as in the case of accounting for the underlying geometry of the system in the measurement, gating can still be carried out with some modification [2].

As defined in Sections 1 and 2, at each time step there is exactly one measurement of unknown origin that is received. If the measurement falls within the gate, it is passed on for use in the EKF. If the measurement falls outside the gate, no measurement is passed to the EKF, and the EKF simply predicts forward one time step.

3.4 Probabilistic Data Association

Probabilistic Data Association has been used when multiple measurements are received at each time step in the tracking of an object in a region with uniformly distributed environmental clutter [1]. Given this assumption (uniformly distributed clutter), the probability β_i of each measurement z_i being the true measurement from the object of interest is computed. The PDA is used in conjunction with a filtering algorithm, such as the EKF, where updated estimates $\hat{x}_i(k|k)$ due to each measurement z_i are computed and the overall updated state estimate is

$$\hat{x}(k|k) = \sum_{i=0}^{m_k} \hat{x}_i(k|k)\beta_i(k), \quad (8)$$

where each $\hat{x}_i$ is of the form (6), m_k is the number of measurements at time k , and

$$\sum_{i=0}^{m_k} \beta_i(k) = 1,$$

i.e., the probabilities are mutually exclusive as well as exhaustive. When tracking in cluttered environments where multiple measurements are received at each time step, gating is often used to limit the computational complexity. When gating is used, m_k in (8) is the number of gated measurements. Thus for the system considered in this report, $m_k = 1$ always.

When tracking in cluttered environments with some types of sensors (e.g., radar or sonar), the probability P_D of detecting the object of interest is often less than 1. That is, at each time step, the probability P_D that one of the multiple measurements received is due to the object of interest is less than 1. For this report, $P_D = \text{UI}$.

The $i = 0$ case in (8) represents the hypothesis that no true measurement from the object of interest is received at time k . This implies that for the type of sensor described in Section 2, if the one measurement per time step is determined to not be from the object of interest, the update (8) will be entirely based on the $i = 0$ case.

The error covariance associated with (8) is

$$P(k|k) = \beta_0(k)P(k|k-1) + [1 - \beta_0(k)]P^c(k|k) + \tilde{P}(k)$$

where $P(k|k-1)$ is the same as is in the EKF,

$$\tilde{P}(k) = W(k) \left[\sum_{i=1}^{m_k} \beta_i(k) \nu_i(k) \nu_i'(k) - \nu(k) \nu'(k) \right] W'(k)$$

where

$$\begin{aligned} \nu_i(k) &= z_i(k) - \hat{z}(k|k-1), \\ \nu(k) &= \sum_{i=1}^{m_k} \beta_i(k) \nu_i(k), \end{aligned}$$

and

$$P^c(k|k) = [I - W(k)H(k)]P(k|k-1).$$

Clutter measurements are assumed to be uniformly distributed throughout the tracking volume with density

$$\lambda = \frac{m_k}{V_k},$$

where V_k is the volume of the validation region

$$V_k = c_{n_z} |\gamma S(k)|^{1/2} = c_{n_z} \gamma^{n_z/2} |S(k)|^{1/2}$$

and n_z is the dimension of the measurement and c_{n_z} is the volume of the n_z -dimensional unit hypersphere.

The $\beta_i(k)$ probabilities are computed as [1]

$$\begin{aligned} \beta_i(k) &= \frac{e_i}{b + \sum_{j=1}^{m_k} e_j}, \quad i = 1, \dots, m_k \\ \beta_0(k) &= \frac{b}{b + \sum_{j=1}^{m_k} e_j} \end{aligned}$$

where

$$e_i \triangleq \exp \left[-\frac{\nu_i'(k) S^{-1}(k) \nu_i(k)}{2} \right]$$

and

$$\begin{aligned} b &\triangleq \lambda |2\pi S(k)|^{1/2} (1 - P_D P_G) / P_D \\ &= (2\pi/\gamma)^{n_z/2} \lambda V_k / c_{n_z} (1 - P_D P_G) / P_D. \end{aligned}$$

3.5 Kolmogorov-Smirnov Tests

Kolmogorov-Smirnov (KS) tests provide another approach for determining the probability of whether a measurement in a set originated from the object of interest [5]. Consider $\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n$, independent observations of a random variable with unknown cumulative distribution function (CDF) $F(\bar{x})$. If the hypothesis is

$$H : F(\bar{x}) = F_0(\bar{x}),$$

then any test of this hypothesis is a goodness-of-fit test [4]. This hypothesis states that the theoretical CDF of the random variable is equal to the CDF from which actual measurements are drawn. The KS tests and many simple variants are goodness of fit tests. Empirical CDFs are formed for a *window* of n measurements and some metric is then applied to measure the distance between theoretical and empirical CDFs.

For n ordered samples

$$\bar{x}_{(1)} \leq \bar{x}_{(2)} \leq \dots \leq \bar{x}_{(n)},$$

the empirical cumulative distribution function is [4]

$$S_n(\bar{x}) = \begin{cases} 0, & \bar{x} < \bar{x}_{(1)} \\ r/n, & \bar{x}_{(r)} \leq \bar{x} < \bar{x}_{(r+1)} \\ 1, & \bar{x}_{(n)} \leq \bar{x} \end{cases} \quad (9)$$

If $F_0(\bar{x})$ is the true, fully specified theoretical cumulative distribution function from which the samples are drawn, then from the strong law of large numbers

$$\lim_{n \rightarrow \infty} P\{S_n(\bar{x}) = F_0(\bar{x})\} = 1.$$

Define the metric used for measuring the separation between theoretical and empirical CDFs to be

$$A_n = \left| \int_0^\infty (F_0(\bar{x}) - S_n(\bar{x})) dx \right|. \quad (10)$$

This is different from the usual measure of deviation

$$D_n = \sup_{\bar{x}} |S_n(\bar{x}) - F_0(\bar{x})| \quad (11)$$

and its variants that are well developed for the KS test. The measure of deviation (11) proved to not be sensitive enough for the data association problem considered here. The metric (10) is more sensitive to incorrect measurements. This results from the fact that (10) largely takes advantage of how a single incorrect measurement affects the shape of the empirical CDF. When using (10) distinct peaks appear in this value for windows of n measurements that contain false measurements. Figure 1 shows these peaks in error for the two satellite system for which results are presented in Section 6.




Figure 1: Peaks in error between empirical and theoretical CDFs around incorrect measurements




Figure 2: Satellites in close proximity with measurement mixing occurring

There has been a large body of theory developed using D_n in (11), demonstrating many properties, such as the distribution of D_n is independent of the distribution $F_0(\bar{x})$ for continuous CDFs [4]. Critical values or thresholds $d_n(\alpha)$ have also been established [4, 5] such that

$$P\{D_n < d_n(\alpha)\} = 1 - \alpha$$

where $1 - \alpha$ is the confidence level. That is, if $D_n < d_n(\alpha)$, then with probability $1 - \alpha$, the empirical CDF $S_n(\bar{x})$ is formed from n samples drawn from $F_0(\bar{x})$. However, there is currently no theoretical basis for specifying critical values for A_n in (10), and this is an area of future work. Section 6 will show the performance of using (11) on the example discussed in Section 5.

4 Computational Complexity

The batch filter has historically found wide use throughout the aerospace tracking community. A probable reason for this is that the batch filter converges to zero RMS error with the assumption of a perfect model. The batch filter has been used classically for problems where there is no measurement mixing. The capability of associating data is not directly built into the batch filter like it is in gating, the KS tests, and PDA. The batch filter can be used for the data association problem, but at a large cost in computational complexity.

To illustrate the heavy computational complexity of applying the batch filter to the data association problem, consider the situation pictured in Figure 2. There are two satellites in similar orbits, i.e., they are in close proximity. A single ground based sensor produces three measurements after three time steps. The first and third measurements come from satellite A and the second from satellite B, which is assumed to not be of interest. The method by which the batch filter accomplishes data association is to minimize the covariance associated with the estimate over the set of possible hypotheses of measurement origin. The set of possible hypotheses for the two satellite system in Figure 2 is

Hypothesis	Measurement		
	1	2	3
a	Y	Y	Y
b	Y	Y	N
c	Y	N	Y
d	Y	N	N
e	N	Y	Y
f	N	Y	N
g	N	N	Y
h	N	N	N

Thus, the batch filter associates the three measurements by testing each of these hypotheses (i.e., running the measurements through a batch filter with the corresponding assumptions) to find the one that minimizes the covariance of the estimate. Obviously for this case the covariance will be minimized by hypothesis *c*.

For the example above, there are only two satellites and three measurements taken over three time steps. The batch filter must be run eight separate times to check the hypotheses of whether or not each individual measurement originated from the object of interest or not. In general, the computational complexity of using a batch filter at each individual time step is 2^k , where k denotes the current number of time steps that have elapsed since measurements started to be taken. For the example above, $k = 3$. It is easy to see that as the number of time steps and hence measurements goes up, the computational complexity associated with using the batch filter for data association becomes intractable very quickly.

In contrast to this, the computational complexity of hypothesis testing does not show up in recursive methods like the EKF and PDAF. These methods only take the current measurement as input, independent of time elapsed since starting to take measurements (i.e., $m_k = 1$, the number of measurements used for filtering is always one at each time step). Thus there is a constant amount of computational complexity associated with these methods at each time step.

5 Orbital Description

To actually run simulations for the purposes of testing these data association techniques, orbits for the satellites must be described. The orbits are always described in terms of the six classical orbital elements ($a, e, i, \Omega, \omega, \nu$) denoting semi-major axis, eccentricity, inclination, right ascension of the ascending node, argument of perigee, and true anomaly. The sum of results for this report will be given for two satellites in equatorial plane orbits, one in a true circular orbit and the other with a slight perturbation in the eccentricity.

Choosing a Cartesian basis corresponding to the ECI frame provides us with the most natural choice of coordinates for describing orbits. In simulation, the six orbital parameters are used to obtain position and velocity initial conditions in Cartesian coordinates aligned with the ECI frame. Thus, the dynamics describing the orbits are written down in this frame. The method by which the dynamics are solved is a straightforward approach using Lagrangian techniques. The resulting equations are simple 2-body dynamics; the same equations can be found in any standard astro-dynamics text, for example [6].

The equations for describing the dynamics are obtained by using *Euler-Lagrange formulations*. The Lagrangian is defined as follows

$$L = KE - PE. \quad (12)$$

For satellites in the ECI frame, the kinetic energy is just $KE = 1/2 M_{satellite_i} (\dot{x}^2 + \dot{y}^2 + \dot{z}^2)$ and the potential energy $PE = M_{satellite_i} G_{earth} \sqrt{x^2 + y^2 + z^2}$.

The dynamics can then be directly obtained from the *Euler-Lagrange equations*

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}} - \frac{\partial L}{\partial q} = 0 \quad (13)$$

where q is the state vector $q = (\dot{x}, x, \dot{y}, y, \dot{z}, z)$. The resulting equations are

$$\begin{aligned}\ddot{x}(t) &= \frac{\mu x(t)}{(x(t)^2 + y(t)^2 + z(t)^2)^{3/2}} \\ \dot{x}(t) &= \frac{d}{dt}x(t) \\ \ddot{y}(t) &= \frac{\mu y(t)}{(x(t)^2 + y(t)^2 + z(t)^2)^{3/2}} \\ \dot{y}(t) &= \frac{d}{dt}y(t) \\ \ddot{z}(t) &= \frac{\mu z(t)}{(x(t)^2 + y(t)^2 + z(t)^2)^{3/2}} \\ \dot{z}(t) &= \frac{d}{dt}z(t)\end{aligned}$$

where $\mu = 3.986 \cdot 10^5 \text{ km}^3\text{s}^{-2}$. These equations are integrated in simulation to produce the “truth” orbits, i.e., the actual orbit with initial conditions specified by the six orbital parameters. At this point there is no noise added yet.

The sensor is located on the surface of the earth, and the SEZ frame whose origin is aligned with the sensor’s location is the frame in which measurements are assumed to be “taken.” To generate “measurements” for the simulation, the positions of the orbits in ECI are first sampled at a specified rate. Process noise is added to these samples because the dynamics are written down in the ECI frame and the process noise is the noise associated with inadequacies in this model (see equation (1)). After adding process noise, a representation of the resulting “noisy” ECI samples is obtained in SEZ coordinates centered at the sensor’s location through the use of rigid body transformations.

Obtaining the SEZ representations of the orbits (with process noise added) is accomplished using a composition of rigid body transformations. The resulting transformation has the form

$$q_3 = g_{31} \cdot q_1$$

where q_1 is the homogeneous representation of the position vector $q_1 = (x, y, z, 1)$ in frame 1, “ $\cdot$ ” is standard matrix multiplication, and $g = (p, R) \in SE(3)$. In general, g can be written as

$$g = \begin{bmatrix} R & p \\ 0 & 1 \end{bmatrix}$$

and is such that

$$g_{31} = g_{32} \cdot g_{21} \tag{14}$$

where R is a standard Euler rotation matrix about the x , y , or z axes, p is the position vector (x, y, z) , and “ $\cdot$ ” is again standard matrix multiplication. The notation g_{31} defines a rigid body transform g that allows us to transform objects represented in frame 1 to frame 3. The right hand side of equation (14) is a composition of rigid body transformations: the resulting transformation from frame 1 to frame 3 is just the transformation from frame 1 to 2 followed by the transformation from frame 2 to 3. For a more thorough description of rigid body transformations see [4].

To transition between ECI and SEZ coordinates it is only necessary to know the earth’s rotational rate, the latitude and longitude of the sensor’s position on the earth, and the radius of the earth. These four quantities will describe four rigid body transformations. The first transformation is a time varying rotation about the earth’s rotation axis (which in ECI is aligned with the z -axis) at a rate equal to the earth’s rotation rate, and aligns the frame such that it is not rotating with respect to the sensor. The second transformation is another rotation about the earth’s rotational axis, which is constant and corresponds to the sensor’s longitude. The next transformation is another rotation, this one about the y -axis of the frame by the correct amount of latitude corresponding to the sensor position. Finally, the last transformation is just a translation in the z -direction by the earth’s radius. Thus the resulting rigid body transformation that takes points represented in ECI coordinates into SEZ coordinates is

$$g_{ECI \rightarrow SEZ} = g(p_{er}, R(0)) \cdot g(0, R_y(\psi)) \cdot g(0, R_z(\phi)) \cdot g(0, R_z(\theta(t))) \quad (15)$$

where $p_{er} = (0, 0, 6378.135 \text{ km})$, R_y an Euler rotation matrix corresponding to a rotation about the y -axis, ψ the latitude at which the sensor is located, R_z an Euler rotation matrix about the z -axis, ϕ the longitude at which the sensor is located, and $\theta(t)$ the earth's rotational rate.

All of the samples produced in ECI coordinates undergo this transformation and we are left with (x, y, z) descriptions of the orbits in the SEZ frame. These transformations can clearly be seen in the short movie located at <http://robotics.colorado.edu/wiki/index.php/Travers:Satellites>.

The next step in obtaining measurements for the simulation is to then take the noisy samples expressed in an SEZ coordinate frame and represent them in R/Az/El coordinates. This transformation is defined by

$$(R, Az, El) = (\sqrt{x^2 + y^2 + z^2}, \tan^{-1}(-x/y), \sin^{-1}(z/\sqrt{x^2 + y^2 + z^2})). \quad (16)$$

Because the actual sensors take measurements in R/Az/El, measurement noise distributions are assumed to be represented in R/Az/El. This fact implies that measurement noise is added to the samples in R/Az/El.

These measurements in R/Az/El represented in a sensor local frame then need to be transformed back into ECI coordinates so that they are in the form of equation (2). This is accomplished by doing the previously stated transformations in reverse order. Note that the inverse transformation for equation (16) is

$$(x, y, z) = (-R \cos(Az) \cos(El), R \sin(Az) \cos(El), R \sin(El)), \quad (17)$$

and the rigid body transformation from SEZ coordinates to ECI coordinates is the matrix product on the right hand side of (15) in reverse order.

The fact that measurements are taken in R/Az/El coordinates from an SEZ coordinate frame implies that the resulting measurement error covariance is most naturally represented in R/Az/El, i.e., the measurement noise distributions in all three directions R/Az/El are Gaussian. The measurement error covariance, $R_{R/Az/El}(k)$, must be represented in the ECI frame to correspond with the measurement and system model. The variance in each direction can be thought of as a member of the corresponding tangent space at each point, thus the transformation of the covariance has the form

$$\left(\frac{\partial x}{\partial y}(k) \right) \cdot (R_{R/Az/El}(k)) \cdot \left(\frac{\partial x}{\partial y}(k) \right)^T, \quad (18)$$

where $\partial x/\partial y$ is the Jacobian. In R/Az/El coordinates the measurement error covariance description is constant, under this transformation the measurement error covariance becomes time varying. A movie of how the three dimensional measurement error distribution changes as an object traverses an orbit, when viewed from ECI coordinates, can be found at <http://robotics.colorado.edu/wiki/index.php/Travers:Satellites>.

6 Results

Results are presented for a simulated system consisting of two satellites. A number of system parameters are varied to highlight performance over two specific objectives: tracking capability and correct identification of individual measurements as being from the satellite being tracked or not.

Figure 3 displays the results of each of the filters described in Section 3 in terms of RMS position tracking error. The RMS error is the error between the filtered estimate of position at each discrete time and the actual position of the simulated satellite in orbit at each corresponding time. The cumulative average of the norms (standard Euclidean norm) of theses differences up to each time step are the values of the RMS calculation.

The EKF without false measurements as well as the Batch filter are run on a list of measurements that do not include mixing (all the measurements are from the same satellite). These results are included as a basis for comparison. The same can be said about the "Samples Alone" plot. This plot is obtained by comparing the value of the position measurement of a single satellite expressed in ECI coordinates (with process noise and measurement noise added in the appropriate frames) to truth at each time step, then calculating the new RMS value including the cumulative normed difference up to that time step. This plot can be seen as




Figure 3: Comparison of various filter performances. The horizontal axis is time in seconds and the vertical RMS error in kilometers. The upper plot presents results for the case where the two satellites are further away ($\epsilon_f = 0.01$) from each other than the lower ($\epsilon_c = 0.0004$). The UI=95%.

how well the measurements from a single satellite do without any filtering. In terms of RMS error of tracking a single satellite where there is no measurement mixing, the batch filter performs better than the EKF. Both methods offer an improvement in tracking, measured in RMS error, over the measurements considered alone.

The single satellite for which these results are stated is in a circular equatorial orbit. The radius of this orbit is $r = 7078.14$ km. The initial conditions are always such that the satellite starts on the x -axis of the ECI frame. The results of the other filters evaluated here are all for a two satellite system. The uniformity index is 95%, meaning that 5% of the total measurements are from satellite 2, while it is satellite 1 that we are trying to track. Proximity for these two satellites will be defined in terms of a perturbation in eccentricity. The first satellite is in the same orbit described above, i.e., a circular equatorial orbit with $r = 7078.14$ km. The second satellite is again equatorial, except that it is slightly elliptical with semi-major axis $a = 7078.14$ km. For the “Far” proximity case, $\epsilon_f = 0.01$ and for the “Close” case $\epsilon_c = 0.0004$. Over the measurement period of $2s$, ϵ_f corresponds to an average separation distance of 70.78 km and ϵ_c corresponds to an average separation of 8.9 km. The process noise associated with each component of the position is described by the normal distribution $N(0, (0.7)^2 \text{km}^2)$. The process noise associated with each component of the velocity is described by the distribution $N(0, (3)^2 (\text{km/s})^2)$. The distributions describing measurement noise in range, azimuth, and elevation are $N(0, (0.02)^2 \text{km}^2)$, $N(0, (0.01)^2 \text{radians}^2)$, and $N(0, (0.01)^2 \text{radians}^2)$ correspondingly. Lastly, the time step is $dt = 0.001$ s.

The PDA results are in general better than the EKF and converge faster than the Batch filter. The EKF with Gating does better than the EKF for the no mixing case as well as the EKF run on the results of the KS-Test. When the two satellites are far from each other, the EKF with Gating does noticeably better than the other two EKF results, while the EKF with KS does basically about as good as the EKF evaluated without any incorrect measurements. The EKF with Gating results do better because at this distance they not only gate out all of the incorrect measurements, but they also gate out the true measurements that are in effect “bad measurements.” The EKF run on KS results and the EKF on true measurements alone are essentially the same because the KS test only removes the incorrect measurements, meaning that after the KS test is run, the resulting list of measurements is exactly the same as the list of measurements used to evaluate the EKF alone, except that in this case 5% of the measurements are removed by the test.

When the two satellites are in close proximity the EKF results change slightly. The EKF run on true




Figure 4: Comparison of PDA with gating performance as UI is varied. RMS error in kilometers are plotted against time in seconds. The satellites considered are both in equatorial orbits with $a = 7078.14$. The satellite being tracked has eccentricity, $\varepsilon_t = 0.1$ and the satellite not being tracked has eccentricity, $\varepsilon_{fl} = 0.11$.

measurements obviously does not change. Both the EKF with Gating and EKF with KS test do worse as the satellites come into closer proximity. The EKF with gating still does marginally better than the EKF on truth, but the KS test results are actually worse. This can be explained by looking at Figure 5. At the separation distance of 8.9 km, the KS test is almost completely failing. Thus most of the incorrect measurements are being left in the list of measurements and the filtering performance is adversely affected.

Figure 4 illustrates the effect of varying the uniformity index (UI) on PDA filtering performance. As the UI goes down, filtering performance degrades. For the case plotted, two satellites in equatorial orbits are considered with $a = 7078.14$ km for each. The satellite being tracked has eccentricity, $\varepsilon_t = 0.1$, and the “follower” satellite has eccentricity, $\varepsilon_{fl} = 0.11$.

Described in Section 3, both gating and the KS test variant are methods of identifying individual measurements as being from the satellite being tracked or not. Figure 5 displays simulated results for the two satellite system where the separation distance between the two satellites is varied and the corresponding percentage of incorrect measurements correctly identified is calculated for each method. The separation distance is calculated from the average separation between satellites over the total measurement period. In simulation the separation distance is actually varied by varying the eccentricity of the orbit of satellite 2, other than this the system parameters are the same as above with $UI = 95\%$. The KS test starts to have degraded performance at a higher separation distance than does gating, although above about 40 km both




Figure 5: Comparison of Gating and KS methods at correctly identifying false measurements as the separation distance between the two satellites is varied; $UI=95\%$.




Figure 6: Comparison of Gating and KS variant as a function of the discrepancy between actual and modeled measurement error covariances. The UI=95%. The two satellites considered are in the “far” configuration (i.e., $\varepsilon_f = 0.01$).

methods perfectly identify all of the false measurements.

Correctly identifying incorrect measurements carries a tradeoff in gating: incorrectly identifying true measurements as false. The intent of this work is not only to correctly identify false measurements as false, but also to correctly identify true measurements as true while doing so. Figure 6 shows a particular situation in which gating struggles to correctly identify true measurements. For the case considered, UI = 95% and the two satellites are in the far configuration (i.e., the two satellites are in equatorial orbits, the satellite being tracked in a circular orbit and the satellite not being tracked having identical orbital element values except for a difference in eccentricity, $\varepsilon_f = 0.01$). This figure plots the percentage of true measurements correctly identified as true as a function of how well the modeled and actual measurement error covariances match each other. Thus, at Sigma=1 the modeled and actual covariance’s match each other perfectly. At Sigma=2 the actual measurement error covariance is twice the modeled, and so on. As the discrepancy between actual and modeled measurement error covariance increases, the performance of gating can be seen to degrade quickly, while the KS test maintains its performance throughout. This is explained by the empirical basis from which the KS test can form “theoretical” CDFs. The KS test does not rely on an accurate model of the measurement error covariance. Rather, when the satellite of interest is not close to other space objects, “clean” unmixed measurements can be used to empirically determine the “theoretical” CDF for the satellite. Hence, the KS test also does not depend on Gaussian distributions of the measurement errors.

7 Conclusions & Future Works

This report has highlighted several key issues facing tracking and correct data association for multiple satellite systems where there is measurement mixing. Initial solutions to these issues were presented, and simulation results given.

The first problem addressed was that of computational complexity, a common concern in typical tracking schemes for multi-body systems. The historical tendency to use batch filters for the purposes of filtering single body systems turns out to be computationally burdensome for multi-body systems with measurement mixing. As stated, the batch approach to data association grows as 2^k where k is the number of time steps. For one measurement (of unknown origin) at each time step while tracking one satellite, the recursive methods such as the EKF and PDAF outlined in Section 3 have constant complexity. For this reason the recursive filtering methods are preferable for filtering satellite systems where there is measurement mixing.

Simulation results were presented for a simple two satellite system with measurement mixing. The results were grouped into filtering performance and correct data association performance. The results show that of the data association and filtering algorithms considered, PDA with gating has the best performance although PDA by itself still offers a dramatic improvement over the standard EKF used by itself or in conjunction with the gating or KS techniques.

The data association problem was considered from two separate standpoints, the first being the identification of false measurements and the second the correct identification of true measurements while identifying false measurements. In terms of finding a minimum separation distance between satellites for which the identification of incorrect measurements is possible, gating is the preferred method of data association over the KS test variant. In terms of correctly identifying true measurements, the performance of gating contains tradeoffs and in general is not very robust to modeling errors in the system parameters. This was highlighted in the simulation results where the difference between modeled and actual measurement error covariances for the simulated system was varied. The empirical basis of the KS test variant allowed for perfect performance as the modeled covariance deviated further from the truth, while the gating method performance degraded. These results show that both gating and the KS test variant have favorable characteristics in different situations. For this reason, one proposed avenue of future work is the development of a hybrid method of these two techniques such that the strongest qualities of each are used to produce a better data association method for systems with mixed measurements.

Another focus of future work is to add various levels of complexity and fidelity to the two satellite, single sensor system considered in the Section 6. The first goal is to implement a network of sensors around the earth that take measurements of the satellites as they travel into and out of each individual sensor's range (which is limited to line of sight). The work presented in this report assumes that line of sight is not required for the sensor to make measurements (i.e., the single sensor on the surface of the earth is assumed to be able to track a satellite through an entire orbit). The network of sensors is a more realistic scheme and in theory could present added complications to the data association problem, especially in regions where the object being tracked switches from one sensor's range to another. Adding maneuvering to the individual satellites is another area of future work. With maneuvering, the data association problem can be studied for systems where two satellites are originally not in close proximity, but then one of the satellites moves toward the other ("chasing" it in a sense). In addition to a network of sensors and maneuvering, the effect of adding more satellites to the system will be considered in future work. The results presented in this report considered two satellites with proximity defined in a specific way (namely the two by one ellipse that results from a perturbation in eccentricity of two satellites in the same orbit with the same initial conditions). The techniques presented in this report are extendable to multi-satellite systems, but the overall level of complexity that results from having many satellites in varied degrees of proximity may require more advanced data association techniques.

References

- [1] Y. Bar-Shalom and T.E. Fortmann, Tracking and Data Association, Academic Press, Inc., San Diego, CA, 1988.
- [2] T. Bailey, B. Upcroft, H. Durrant-Whyte, "Validation Gating for Non-Linear Non-Gaussian Target Tracking," in Ninth International Conference on Information Fusion, Florence, Italy, 2006, pp. 1-6.
- [3] Y. Kosuge and T. Matsuzaki, "The Optimum Gate Shape and Threshold for Target Tracking," SICE 2003 Annual Conference, Fukui, Japan, 2003, pp. 2152-2157.
- [4] A. Stuart, K. Ord, and S. Arnold, Kendall's Advanced Theory of Statistics, Oxford University Press, 1999.
- [5] L.H. Miller, "Table of Percentage Points of Kolmogorov Statistics," J. American Statist. Assoc., 51(273):111-121, Mar. 1950.
- [6] V.A. Chobotov, Orbital Mechanics, Third Edition, AIAA Educational Series, 2002.
- [7] R.M. Murray, Z. Li, S.S. Sastry, A Mathematical Introduction to Robotic Manipulation, CRC Press, 1994.