




Fig. 1 Sample run, with true and estimated target acceleration in the horizontal plane; $\sigma_a^2 = (3g)^2$.




Fig. 2 Acceleration estimation error in the horizontal plane; $\sigma_a^2 = (3g)^2$.

the vertical plane. The variance σ_a^2 affects mainly the acceleration estimate. It has a minor effect on the velocity and position estimate.

Summary

A simple structure for a high-performance three-dimensional tracking filter is derived in the rotating sensor frame assuming high measurement rates and small angular deviations during target tracking. The resulting three-dimensional estimator consists of three independent filters, whose state includes position, velocity, and acceleration along each of the axes in the rotating sensor frame. Because line-of-sight models for the target acceleration are identical in all three axes and independent, the replacement of the basic exponentially correlated acceleration model by more complicated models preserves the filter structure. Such sophisticated models are based on the multiple model approach.

References

- ¹Fitts, J. M., "Aided Tracking as Applied to High Accuracy Pointing Systems," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-9, No. 3, 1973, pp. 350–368.
- ²Baheti, R. S., "Efficient Approximation of Kalman Filter for Target Tracking," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-22, No. 1, 1986, pp. 8–14.
- ³Pearson, J. B., and Stear, E. B., "Kalman Filter Application in Airborne Radar Tracking," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-10, No. 3, 1974, pp. 319–329.

⁴Singer, R. A., "Estimating Optimal Tracking Filter Performance for Manned Maneuvering Targets," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-6, No. 4, 1970, pp. 473–483.

Covariance Control for Sensor Management in Cluttered Tracking Environments

Michael Kalandros* and Lucy Y. Pao†

University of Colorado, Boulder, Colorado 80309-0425

I. Introduction

SENSOR management schemes have recently been proposed to reduce the tracking demands on multisensor systems while minimizing the loss of tracking performance by selecting only enough sensing resources to maintain a desired covariance level for each target.^{1,2} However, covariance control algorithms to date have not addressed the presence of clutter measurements and the need for data association in those cases. This Note presents a method of reducing the effects of probabilistic data association on covariance control algorithms through the addition of a scalar loss of information parameter. Whereas single-sensor management systems have been designed for the probabilistic data association filter (PDAF),^{3,4} including one that uses the loss of information parameter to approximate the covariance calculation,⁴ no sensor managers to date have been designed around the multisensor extension of the PDAF.^{5,6} Monte Carlo simulations show that, without this parameter, the covariance control system is unable to maintain the desired covariance, resulting in a much larger actual covariance level. Use of the loss of information parameter generally restores system performance.

The Note is organized as follows: The multisensor PDAF and the resulting loss of information is summarized in Sec. II. The sensor management algorithm that estimates this loss of information is derived in Sec. III, and simulation results of these algorithms are presented in Sec. IV. Finally, conclusions are given in Sec. V.

II. Sequential Multisensor PDAF

The basic tracking system architecture includes a set of sensors that observe the environment, including the targets to be tracked, and a central system for processing the outputs of those sensors. These sensors are assumed to receive false measurements in addition to returns from targets of interest. These clutter measurements can be due to noise in the sensing system itself, multipath reflections from the target, or returns from objects that are not targets of interest. Clutter measurements are typically assumed to be uniformly distributed, and thus, the number of false measurements for a given volume will be characterized by a Poisson distribution parameterized by the clutter density λ . Additionally, it is possible that not all targets may be detected by a specific sensor during each sampling interval, which further complicates the tracking process. Several tracking algorithms have been designed to address these challenges, including the PDAF.⁷ The PDAF replaces individual

Received 26 August 2002; revision received 10 October 2003; accepted for publication 30 December 2003. Copyright © 2004 by Michael Kalandros and Lucy Y. Pao. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/04 \$10.00 in correspondence with the CCC.

*Research Assistant, Department of Electrical and Computer Engineering; currently Senior Staff, Air Defense Systems Department, Johns Hopkins Applied Physics Laboratory, Laurel, MD 20723-6099; Michael.Kalandros@jhuapl.edu. Member AIAA.

†Associate Professor, Department of Electrical and Computer Engineering; pao@colorado.edu.

measurements with weighted averages of both the target-originated and the clutter measurements. The weights are a function of the normalized distance of each measurement from the predicted target location.

When it is assumed that the tracking system has N_s sensors, there are 2^{N_s} possible combinations or subsets of those sensors that can be selected by the multisensor manager. The i th possible subset is defined as Φ_i , and N_{s_i} is the number of sensors in that combination. The input from each selected sensor is used to update the state estimate of the target.

The sequential multisensor PDAF (SMSPPDAF) is an extension of the original probabilistic data association algorithm⁷ that combines inputs from multiple sensors monitoring stochastic or linearized systems to form an estimate in a state-space representation. This approach begins with a linear representation of the target motion and measurement processes,⁷

$$x_t(k) = F_t x_t(k-1) + w_t(k-1) \quad (1)$$

$$z^t(k, j) = H_j x_t(k) + v_j^t(k), \quad j = 1, \dots, N_{s_i} \quad (2)$$

where $x_t(k)$ is the state of target t at time k , F_t is the state transition matrix that models the motion of target t between time $k-1$ and k , $z^t(k, j)$ is a measurement of target t from the j th sensor in Φ_i , and H_j is the measurement matrix that creates $z^t(k, j)$. Here $w_t(k)$ is a variable representing process noise or higher-order motion not modeled by F_t , and $v_j^t(k)$ is a variable representing measurement noise in sensor j . Both $w_t(k)$ and $v_j^t(k)$ are assumed to have zero-mean, white, Gaussian probability distributions.

Because $w_t(k)$ and $v_j^t(k)$ are zero-mean noise processes, the target states and measurements in the next time interval can be predicted by

$$\hat{x}_t(k|k-1) = F_t \hat{x}_t(k-1|k-1) \quad (3)$$

$$\hat{z}^t(k, j) = H_j \hat{x}_t(k|k-1), \quad j = 1, \dots, N_{s_i} \quad (4)$$

The difference between the predicted measurement and the actual measurements, $z_{j,\ell}(k)$,

$$v_t(k, j, \ell) = z_{j,\ell}(k) - \hat{z}^t(k, j) \quad (5)$$

is known as the innovation. The Mahalanobis distance (see Ref. 7) of each innovation is used to gate or eliminate measurements that are unlikely to have originated from the target.

The PDAF derives an estimate for the state by creating a weighted sum of the gated innovations, known as the combined innovation,

$$v_t(k, j) = \sum_{\ell=0}^{m_{k,j}} v_t(k, j, \ell) \beta_{t,j,\ell}(k) \quad (6)$$

where $\beta_{t,j,\ell}(k)$ is the probability that measurement $z_{j,\ell}(k)$ is the true measurement of target t from sensor j , $m_{k,j}$ is the number of gated measurements from sensor j at time k , and $\ell=0$ signifies the event that none of the gated measurements is the true target measurement.

The covariance of the state and the innovation predictions are

$$P_t(k|k-1) = F_t P_t(k-1|k-1) F_t' + Q_t(k-1) = P(k|k, 0) \quad (7)$$

$$S_t(k, j) = H_j P_t(k|k, j-1) H_j' + R_j(k), \quad j = 1, \dots, N_{s_i} \quad (8)$$

respectively, where $Q_t(k)$ is the process noise covariance and $R_j(k)$ is the measurement noise covariance for the j th sensor. $P_t(k|k, j)$ for $j = 1, \dots, N_{s_i}$ is the updated covariance resulting from sensor j as defined in Eq. (11).

The sequential algorithm runs a separate filter for each sensor in the combination, propagating its estimate to the next filter,^{5,8}

$$\hat{x}_t(k|k, 1) = \hat{x}_t(k|k-1) + K_t(k, 1) v_t(k, 1)$$

$$\hat{x}_t(k|k, j) = \hat{x}_t(k|k, j-1) + K_t(k, j) v_t(k, j), \quad j = 2, \dots, N_{s_i}$$

$$\hat{x}_t(k|k) = \hat{x}_t(k|k, N_{s_i}) \quad (9)$$

where

$$K_t(k, j) = P_t(k|k, j-1) H_j' S_t^{-1}(k, j), \quad j = 1, \dots, N_{s_i} \quad (10)$$

The state covariance is updated for each filter by

$$P_t(k|k, j) = P(k|k, j-1)$$

$$- (1 - \beta_{t,j,0}(k)) K_t(k, j) H_j P_t(k|k, j-1) + \tilde{P}_{t,j}(k) \quad (11)$$

$$P_t^i(k|k) = P_t(k|k, N_{s_i})$$

$$\tilde{P}_{t,j}(k) = K_t(k, j) \left[\sum_{\ell=1}^{m_k} \beta_{t,j,\ell}(k) v_t(k, j, \ell) v_t'(k, j, \ell) - v_t(k, j) v_t'(k, j) \right] K_t'(k, j) \quad (12)$$

Once the state and covariance estimates have been updated, they are fed back into the algorithm, and the entire process is repeated for the new set of measurements at the next time step.

The presence of $\beta_{t,j,0}$ and $\tilde{P}_{t,j}(k)$ in Eq. (11) changes the calculation of the covariance from a deterministic equation into a stochastic one (because both terms depend on the actual data received). It can be shown, however, that the expectation of the covariance for a single-sensor PDAF can be computed by⁹

$$E[P(k|k)|Z^{k-1}, P(k|k-1)] = [I - q_2 K(k) H(k)] P(k|k-1) \quad (13)$$

where Z^{k-1} is the set of all measurements before time k and q_2 is a scalar parameter that varies between 0 and 1. Because it reduces the effectiveness of the sensor measurements, q_2 is known as the loss of information parameter. If this loss of information due to data association is ignored in covariance control techniques, the sensor manager will overestimate the effect of each sensor on the state estimate covariance and the sensor selections will consistently fail to meet the covariance goals. The covariance control algorithm presented in the next section uses the q_2 parameter and Eq. (13) to estimate the actual performance of each sensor, which allows the sensor manager to make better sensor-to-target assignments.

The q_2 parameter is a function of the expected number of gated clutter measurements [which in turn is a function of the innovations covariance $S(k)$] and the probability of detection p_D (Ref. 9) and can be calculated using the closed-form numerical approximation of q_2 described in Ref. 4. Note that, although the expectation in Eq. (13) is exact for the first sensor in the sensor combination Φ_i , it is only an approximation for the sensors that follow because each depends nonlinearly on the updated covariance resulting from the previous sensors. A method called hybrid approximation^{5,8,10} is a more exact representation, but is more computationally demanding and, thus, is not used in this application. Without the hybrid approximation, it is expected that the value of q_2 will be underestimated for all but the first sensor because the algorithm only calculates the q_2 parameter based on the original prediction covariance. It does not account for the smaller a priori covariance $[P(k|k, j-1)]$ in Eq. (11) seen by all of the sensors after the first one.

III. Covariance Control Approach for the Multisensor PDAF

The algorithm begins by using the covariance of each target state prediction (provided by the tracking algorithm) and calculating the innovations covariance for each potential target-to-sensor pairing using Eq. (8). It then estimates the loss of information parameter by calculating the expected number of clutter measurements, $\lambda V(k)$, where $V(k)$ is the volume of the gate, and using that information to approximate q_2 as described in the preceding section. Once the loss of information parameters have been calculated, they are passed to the covariance controller (Fig. 1).




Fig. 1 Block diagram of the augmented covariance controller for the multisensor PDAF.

The covariance controller ranks each target based on its “need” $n(t)$, a function of the minimum eigenvalue of the covariance difference,^{1,2}

$$n(t) = -\min\{\text{eig}[P_d(t) - P_t(k|k-1)]\}(10 - t_p) \quad (14)$$

where t_p is the priority of the target. Target priority ranges from 1 (high priority) to 9 (low priority), although Eq. (14) can easily be modified to reflect other prioritization schemes. Note that the negative sign means that a need that is greater than zero implies a negative eigenvalue, which in turn means that the desired covariance has not been achieved. The target with the largest need is then selected and the covariance that would result from using each available sensor is calculated. (Note that sensors are assumed to have a finite number of targets that they can track in a single sampling period. Once the number of targets assigned to a sensor is equal to its maximum capacity, that sensor is no longer available.) The controller uses Eq. (13) to estimate the effect of each sensor on the updated target covariance. The sensor that maximizes

$$\min\left(\text{eig}\left\{P_d(t) - E\left[P_t(k|k, j)|Z^{k-1}, P_t(k|k-1)\right]\right\}\right) \quad (15)$$

is selected, the covariance is updated $P_t(k|k, j) \rightarrow P_t(k|k-1)$ and the need is updated, $n(t) = -\min\{\text{eig}[P_d(t) - P_t(k|k, j)]\}(10 - t_p)$. Then the remaining neediest target is selected, and the process repeated until all of the desired covariances have been achieved, that is, the need of each target is zero or negative, or the total tracking capacity of all sensors has been reached.

IV. Simulation Results of Covariance Control with q_2 Parameter

This section presents the results from Monte Carlo simulations using the covariance control technique for the SMSPDAF discussed in the preceding section. The tracking situation presented is two targets moving in the $x-y$ plane in nominally straight lines corrupted by acceleration or control noise. This noise is zero-mean, white, and Gaussian with covariance matrices of 0.05 times the identity matrix. The target states to be tracked (estimated) consist of both the position and velocity of each target, that is, $[x \dot{x} y \dot{y}]'$. The desired covariance for target 1 is diagonal with a position variance of 0.02 and a velocity variance of 1. The desired covariance for target 2 is also diagonal with a position variance of 0.02 but has a smaller velocity variance of 0.5. The targets have equal priority. Each simulation is run for 500 scans.

For the purpose of these simulations, the targets are assumed to be noninteracting, which means that the distance between the two is large enough that the measurement of one will not lie in the gated region of the other. Although this is not always the case in actual tracking applications, interacting targets result in a nonuniform clutter density that is not modeled by the q_2 parameter. Covariance control techniques for tracking interacting targets are discussed in Ref. 11.

The simulations compare the sensor management approach previously described in Ref. 2 (without the q_2 parameter) to the new approach described in Sec. III with the q_2 parameter applied in Eq. (13). Each sensor suite consists of $N_s = 8$ sensors measuring both the x and y position of each target. The probability of detection for each sensor is $p_D = 1$ for each target, which means that each target will generate a valid measurement for each sensor used in each scan, and the gate size is $\gamma = 16$. The tracking capacity of

each sensor is greater than the number of targets; thus, it is possible for both targets to have all eight sensors applied to them in the same sampling period. (In our simulation, each sensor is allowed at most one sensor dwell per target per sampling period.) Although the sensors' accuracy in a given direction varies, they have roughly equal overall abilities. The clutter density λ is also the same for each sensor in a given simulation. Clutter density was varied across simulations from $\lambda = 0.2$ to $\lambda = 1.0$ in steps of 0.2, and the results are averaged over 500 runs each. This clutter density range leads to the expected number of gated clutter measurements varying from 5 to 27 measurements for target 1 and from 3 to 15 measurements for target 2, assuming that the desired covariance goals are met for each target. Note that the estimates for clutter measurements hold for the first sensor applied to each target only. Subsequent sensors will see a lower a priori covariance and, thus, fewer clutter measurements.

The simulations record the actual minimum eigenvalue of the difference between the desired and updated covariance, as well as that of the difference predicted by the sensor selection algorithm for each system. The smallest eigenvalue of the difference will be positive if the algorithm achieves the covariance goal and be negative if it does not. The covariance metric is only evaluated while the target is actually being tracked and does not reflect covariance trends after the target is lost. Track loss is declared when either 1) target-originated measurements have not been gated for 5 consecutive scans, 2) the actual rms position error exceeds 10 times the standard deviation of the Kalman filter's position error estimate (without the effects of data association) for 5 consecutive scans, or 3) more than 500 measurements from a single sensor are gated by a single target.

Figures 2 and 3 show the average minimum eigenvalue of the covariance difference $P_d - P(k|k)$, normalized by the minimum




Fig. 2 Average normalized minimum eigenvalue of the covariance differences when the sensor management algorithm does not take into account the effects of data association.




Fig. 3 Average normalized minimum eigenvalue of the covariance differences when the sensor management algorithm uses the loss of information parameter q_2 to model the effects of data association.




Fig. 4 Combined track loss for both targets with and without the q_2 parameter.

eigenvalue of the desired covariance matrix, for both targets with and without the use of the q_2 parameter (indicated by the solid lines). When q_2 is not used, the minimum eigenvalue of the actual covariance is 40–170 times that of the desired covariance, which means that the covariance goal is not achieved at any clutter level. Furthermore, the minimum eigenvalue of the predicted covariance difference (indicated by the dashed lines) is always positive, which means that the controller “thinks” that it will achieve the desired covariance at every scan. Because the probability of track loss is a function of the number of gated measurements^{9,10} and, thus, a function of the prediction covariance and the clutter density, the failure to maintain a stable covariance leads to dramatically reduced track life, as shown in Fig. 4, the cumulative number of tracks lost at each time scan for each clutter density level.

When the q_2 factor is used (Fig. 3), the desired covariance goals are easily achieved, resulting in a positive definite covariance difference, in all but the highest clutter levels. The error between the predicted difference and the actual difference is always negative, which means that the system is consistently underestimating the sensor performance, as predicted in Sec. II. This improved covariance tracking performance also leads to significantly longer track lifetimes (Fig. 4) over that seen when the q_2 parameter is not used.

V. Conclusions

This Note augments a previously proposed multisensor covariance control technique through the addition of a scalar loss of information parameter, which allows the algorithm to evaluate the reduced effectiveness of each sensor when tracking in cluttered environments. Monte Carlo simulations show that without this parameter, the covariance control system is unable to maintain the desired covariance, resulting in a much larger actual covariance level. Use of the loss of information parameter restores this performance, which allows the control system generally to achieve the desired covariance goals.

Acknowledgments

This work is supported in part by the Office of Naval Research (Young Investigator Award Grant N00014-97-1-0642 and Grant N00014-02-1-0136), the Data Fusion Corporation, the Colorado Advanced Software Institute, and a University of Colorado Faculty Fellowship.

References

- Kalandros, M., “Managing Multiple Sensor Resources Using Covariance Control Techniques for Tracking Systems with Data Association,” Ph.D. Dissertation, Dept. of Electrical and Computer Engineering, Univ. of Colorado, Boulder, CO, Dec. 2000.
- Kalandros, M., and Pao, L. Y., “Covariance Control for Multisensor Systems,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 38, No. 4, 2002, pp. 1138–1157.

³Ahmeda, S. S., Harrison, I., and Woolfson, M. S., “Tracking Multiple Targets with a Phased Array Radar,” *Radar 97: Proceedings of Radar Systems*, IEE, London, 1997, pp. 601–605.

⁴Kershaw, D. J., and Evans, R. J., “Waveform Selective Probabilistic Data Association,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 33, Oct. 1997, pp. 1180–1188.

⁵Frei, C. W., and Pao, L. Y., “Alternatives to Monte Carlo Simulation Evaluations of Two Multisensor Fusion Algorithms,” *Automatica*, Vol. 34, No. 1, 1998, pp. 103–110.

⁶Pao, L. Y., “Centralized Multisensor Fusion Algorithms for Tracking Applications,” *Control Engineering Practice*, Vol. 2, No. 5, 1994, pp. 875–887.

⁷Bar-Shalom, Y., and Fortmann, T., *Tracking and Data Association*, Academic Press, San Diego, CA, 1988, Chap. 6.

⁸Frei, C. W., “A Comparison of Parallel and Sequential Implementations of a Multisensor Multitarget Tracking Algorithm,” M.S. Thesis, Dept. of Electrical and Computer Engineering, Northwestern Univ., Evanston, IL, May 1995.

⁹Fortmann, T. E., Bar-Shalom, Y., Scheffe, M., and Gelfand, S., “Detection Thresholds for Tracking in Clutter—A Connection between Estimation and Signal Processing,” *IEEE Transactions on Automatic Control*, Vol. AC-30, No. 3, 1985, pp. 221–228.

¹⁰Li, X. R., and Bar-Shalom, Y., “Stability Evaluation and Track Life of the PDAF for Tracking in Clutter,” *IEEE Transactions on Automatic Control*, Vol. 36, No. 5, 1991, pp. 588–602.

¹¹Kalandros, M., and Pao, L. Y., “Sensor Management of Tracking Interacting Targets,” *Proceedings of the Workshop on Estimation, Tracking, and Fusion: A Tribute to Yaakov Bar-Shalom*, Omnipress, 2001, pp. 221–248.

Mars and Mercury Missions Using Solar Sails and Solar Electric Propulsion

Christian Circi*

University of Rome “La Sapienza,” 00184 Rome, Italy

Introduction

FROM the beginning of space travel, numerous types of propulsion systems, alternative to the original chemical engine, have been proposed. Today the low-thrust engines are a reality and the Deep Space 1 mission, launched by NASA in October 1998, demonstrated this in practice. Different types of missions obtain advantages from this technology. Scheel and Conway¹ and Kechichian² applied low thrust for transfer from low earth orbit (LEO) to geostationary orbit, and Kechichian³ analyzed optimal steering for north–south stationkeeping of geostationary spacecraft. With regard to the lunar mission, Kluever and Pierson⁴ and Herman and Conway⁵ have studied the transfer from LEO to low lunar orbit in the three-dimensional case. Interplanetary missions have also been studied, particularly the rendezvous trajectories from Earth to Mars.^{6,7} Solar pressure gives the possibility of using another propulsion system, the solar sail. The use of a solar sail is a wonderful prospect for interplanetary travel because no propellant mass is required, but only the weight of the sail.⁸ Minimum-time for Earth–Mars transfer has been studied by Powers and Coverstone⁹ and Otten and McInnes,¹⁰ whereas McInnes et al.¹¹ have presented an extensive investigation of the use of solar sail propulsion for both Mercury orbiter and Mercury sample return missions.

Both solar sails and solar electric propulsion (SEP) use solar light, the first to push the sail and the second to produce energy for the engine. The present work compares the rendezvous missions to Mars

Received 24 September 2003; revision received 19 January 2004; accepted for publication 20 January 2004. Copyright © 2004 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/04 \$10.00 in correspondence with the CCC.

*Assistant Professor, Scuola di Ingegneria Aerospaziale, Department of Aerospace and Astronautical Engineering, Via Eudossiana 18.