

Sensor Management using Discrimination Gain and Interacting Multiple Model Kalman Filters

Wayne Schmaedeke and Keith Kastella
Lockheed-Martin Tactical Defense Systems - Eagan

Abstract: This paper describes an algorithm using a discrimination-based sensor effectiveness metric for sensor assignment in multisensor multitarget tracking applications. The algorithm uses Interacting Multiple Model Kalman Filters to track airborne targets with measurements obtained from two or more agile-beam radar systems. Each radar has capacity constraints on the number of targets it can observe on each scan. For each scan the expected discrimination gain is computed for the sensor target pairings. The constrained globally optimum assignment of sensors to targets is then computed and applied. This is compared to a fixed assignment schedule in simulation testing. We find that discrimination based assignment improves track accuracy as measured by both the root-mean-square position error and a measure of the total covariance.

1. Introduction: The goal of sensor management systems can be regarded as selecting sensors and sensor dwells to increase the information that a data fusion system contains about a surveillance region. Discrimination [Blahut] is related to the notions of information and entropy in probability distributions. A sharply peaked distribution has high discrimination gain and low entropy. Discrimination measures the relative increase in information between two probability distributions. In light of this, discrimination gain is a natural measure for use in sensor management systems. Expected discrimination gain is a measure of sensor effectiveness that has been used in a wide variety of model applications including multisensor / multitarget assignment problems [Sch93], minimizing error correlation between close targets [Sch94], and single and multi sensor detection / classification problems [Kas96b, Kas97a] and joint detection, tracking and identification [Kas96c, Kas97b]. It compares favorably to other methods [Mus96]. In the approach presented here this entails predicting the expected discrimination gain for each sensor dwell, determining the optimal global assignment of sensors to targets given sensor constraints and then applying this solution.

To understand how expected discrimination gain can be computed in Kalman filter based tracking systems, recall that Kalman filters maintain estimates of both the target state and

Submitted to IEEE Transactions on Systems, Man and Cybernetics, January, 1998.

This work was supported by the Lockheed-Martin internal Research and Development program, the Minnesota Center for Industrial Mathematics under Air Force Office of Scientific Research contract F49620-96-1-0382 and Data Fusion Corporation under the Air Force Research Laboratory's Small Business Innovative Research Program.

the state covariance. The target state consists of the target position and velocity. It may also include higher order kinematics terms such as turn rate and acceleration. The covariance is a measure of the uncertainty in these state estimates. In addition to the state and covariance estimates computed after each measurement update, the Kalman filter also includes a mechanism to predict the state at future times based on assumptions about the target motion (the "motion model"). The uncertainty in this prediction can also be computed. Given the predicted covariance, the expected discrimination gain can then be evaluated. A surprising feature of conventional Kalman filters is that the covariance estimate depends only on the uncertainty of the measurements, the state uncertainty and the target motion model, but it does not depend directly on the measurements or the target state itself. (In an extended Kalman filter, the uncertainty can depend on the measurements indirectly through the linearization process.) Thus, in conventional Kalman filters the uncertainty and the expected discrimination gain are insensitive to whether a target is maneuvering or not. This undesirable feature reduces the effectiveness of sensor management systems based on them.

The insensitivity of the Kalman filter to target maneuvers has long been recognized and a wide variety of methods have been developed to recognize maneuvers and change the filter accordingly. Recently, a new approach called the Interacting Multiple Model Kalman Filter (IMMKF) has been developed that seems to offer superior performance to other approaches [Blom88, Blom92]. The IMMKF uses several target motion models and has been successfully applied to large multisensor multitarget applications [Yed]. For example, it may use one model for straight and level flight and different models for turns. The turning models can use a fixed turn rate [Li] or can estimate the turn rate as an additional state in the filter [Kas96]. Older approaches to filtering switch between models. In contrast to this, the IMMKF always maintains all of the models and blends their outputs with weights that are computed probabilistically. In addition to the state estimates for each motion model the IMMKF maintains an estimate of the probability that the target is moving in accordance with each model. With the IMMKF the prediction covariance increases during target maneuvers. As a result the expected discrimination gain goes up when a sensor is directed at a target that is perceived to be maneuvering. Another factor that contributes to the expected discrimination gain with the IMMKF is its estimate of which target motion model is active. As a result, additional sensor resource is directed toward the target as it is initiating a maneuver.

The natural coupling between the IMMKF and sensor management for single target applications was recognized in [Watson] and exploited using measures of effectiveness based on maintaining a desired level of position error covariance. Coordination across multiple sensors is desirable in multitarget/multisensor applications. Typically, each sensor can only observe a few targets on each sensing cycle. By coordinating sensing across platforms, sensing effort is not wasted by having several sensors expend all of their capacity on one target while other targets go unobserved.

The optimization problem considered here is one of assigning sensors to targets in such a way that the one step ahead discrimination gain of the track set is maximized. To do this,

the one step ahead expected discrimination gain G_{st} of a track t after it is updated with covariance data from sensor s is determined for all pairs t,s . This can be formulated as a Linear Programming problem with suitable constraints. Situations where more than one sensor may be assigned to the same target can also be treated [Sch93]. As has been previously pointed out [Nash] one method is to construct pseudo sensors comprised of combinations of the basic sensors. This allows any combination of sensors to be assigned to a single target simply by considering that a single pseudo sensor has been assigned to it. The number of "sensors" is thus increased from S to $2^S - 1$. The problem is now to make an assignment of these pseudo sensors to the targets in an optimal way.

There are however constraints on this assignment. One of the most important is the maximum tracking capacity of a sensor. That is, given a specified time interval, some sensors can scan a certain volume of space and also track a specified number of targets. They cannot exceed this maximum tracking capacity in the specified time period. To examine the impact of coordination between the sensors, we studied three cases. In the first case, each sensor simply selects targets at random for each scan cycle. This is the worst case approach in that it assumes that the sensing is completely uncoordinated. This can be improved upon by assuming that the sensing is coordinated using a round-robin sensing schedule. The result of these approaches is compared to the global discrimination optimum. To establish a baseline for comparison, all of the scheduling schemes maintain central tracks updated via report fusion as soon as measurements are available. Although this may be unrealistic for some target applications it enables comparison of schemes on an even footing.

This paper is organized as follows. Section 2.1 provides a brief overview of the Interacting Multiple Model Kalman Filter used here. The general theory of the IMMKF is presented in [Bar]. A more detailed description of this filter is contained in [Kas96a]. Section 2.2 shows how the expected discrimination gain can be approximately computed for such filters. Section 2.3 presents details on the optimization scheme used to compute the globally optimum assignment of sensors to targets. Section 3 describes a test scenario and presents simulation results. Section 4 presents conclusions and suggestions for further work.

2.1 Interacting Multiple Model Kalman Filters

This section provides a brief overview of Interacting Multiple Model Kalman Filters (IMMKF) and defines notation required in discrimination gain calculations. The IMMKF blends motion models matched to different flight regimes to achieve improved performance against maneuvering targets and minimize latency due to model switching. As part of the design process models are selected to approximate different flight regimes for the target. The IMMKF employs a soft switching scheme that probabilistically blends the output of the different filters. The probability $\mu_j(k)$ that model j is active at time k is estimated as part of the IMMKF. For this feasibility study we use an IMMKF based on one model for uniform straight and level flight with small random accelerations and another model tuned for a constant rate turn, together with random changes in the turn

rate. For more realistic applications, it may be appropriate to include additional models for extreme acceleration maneuvers. Model blending is achieved assuming that the target undergoes transitions between flight modes modeled by a Markov chain. The estimates for the different models are combined with the residual errors for the models to estimate the posterior probability that each mode is active.

The IMMKF is based on a stochastic hybrid system model with a state evolution equation of the form

$$\mathbf{X}(k) = f[k-1, \mathbf{X}(k-1), m(k)] + g[m(k), k-1, \mathbf{X}(k-1), \mathbf{v}(k, m(k))],$$

where k is the time index; $\mathbf{X}(k)$ is a vector describing the state of the system at time k (e.g. position, velocity, turn-rate); f is the state transition function; g is the process noise function; $m(k)$ is the system mode during the interval prior to time k (e.g. uniform motion, or constant-rate turn modes); $\mathbf{v}(k, m(k))$ is a process noise vector. The process noise g generally depends on the mode m . For example, in straight and level flight, the process noise models the effect of small inaccuracies in guidance while for a constant rate turn, the process noise also models changes in the turn rate. The system mode m is modeled as a random event with Markov dynamics. The IMMKF used here has two modes: m_U for nearly uniform motion and m_T for a nearly constant rate turn. For example, $m_U(k-1)$ denotes the event where a target is undergoing uniform motion at time $k-1$ and $m_T(k)$ denotes the event where it is turning at time k .

A summary of the IMMKF and its notation is shown in Table 1 and its data flow is shown in Figure 2. The basic operation of the two-model filter comprised of a uniform motion filter and a constant turn rate filter is summarized as follows.

Table 1
Interacting Multiple Model Kalman Filter Summary

Estimate Mixing (for model j):

model transition probability for time interval T : $\pi_{ji}(T)$

probability of model j at time $k-1$: $\mu_j(k-1)$

predicted model probability: $\mu_j^- = \sum_i \pi_{ji} \mu_i(k-1)$

model mixing probability: $\mu_{ij} = \pi_{ji} \mu_i(k-1) / \mu_j^-$

mixed state estimate for model j :

$$\hat{\mathbf{X}}_{0j}(k-1|k-1) = \sum_i \hat{\mathbf{X}}_i(k-1|k-1) \mu_{ij}$$

mixed state variance estimate for model j :

$$\mathbf{P}_{0j} = \sum_i \mu_{ij} \left\{ \mathbf{P}_i(k-1|k-1) + V_{ij} \right\}$$

$$V_{ij} = \left[\hat{\mathbf{X}}_i(k-1|k-1) - \hat{\mathbf{X}}_{0j}(k-1|k-1) \right] \left[\hat{\mathbf{X}}_i(k-1|k-1) - \hat{\mathbf{X}}_{0j}(k-1|k-1) \right]^T$$

Filtering (for each model j):

state prediction: $\hat{\mathbf{X}}_j(k|k-1) = \mathbf{F}_j(k-1) \hat{\mathbf{X}}_{0j}(k-1|k-1)$

covariance prediction:

$$\mathbf{P}_j(k|k-1) = \mathbf{F}_j(k-1) \mathbf{P}_{0j}(k|k-1) \mathbf{F}_j(k-1)^T + \mathbf{G}_j(k-1) \mathbf{Q}_j(k-1) \mathbf{G}_j(k-1)^T$$

measurement residual: $v_j(k) = \mathbf{z}(k) - \mathbf{H}_j(k-1) \hat{\mathbf{X}}_j(k|k-1)$

residual covariance estimate: $\mathbf{S}_j(k) = \mathbf{H}_j(k-1) \mathbf{P}_j(k|k-1) \mathbf{H}_j(k-1)^T + \mathbf{R}_j(k-1)$

filter gain: $\mathbf{W}_j(k) = \mathbf{P}_j(k|k-1) \mathbf{H}_j(k-1)^T \mathbf{S}_j(k)^{-1}$

measurement update: $\hat{\mathbf{X}}_j(k|k) = \hat{\mathbf{X}}_j(k|k-1) + \mathbf{W}_j(k) v_j(k)$

covariance update: $\mathbf{P}_j(k|k) = \mathbf{P}_j(k|k-1) - \mathbf{W}_j(k) \mathbf{S}_j(k) \mathbf{W}_j(k)^T$

Model probability update:

model likelihood function: $\Lambda_j = (2\pi)^{-1} \left| \mathbf{S}_j(j) \right|^{-1/2} \exp\left(-\frac{1}{2} v_j(k)^T \mathbf{S}_j(j)^{-1} v_j(k)\right)$

model probability: $\mu_j = \mu_j^- \Lambda_j / \sum_i \mu_i^- \Lambda_i$

Estimate combination:

state estimate: $\hat{\mathbf{X}}(k|k) = \sum_j \hat{\mathbf{X}}_j(k|k) \mu_j$

covariance estimate:

$$\mathbf{P}(k|k) = \sum_j \mu_j \left\{ \mathbf{P}_j(k|k) + \left[\hat{\mathbf{X}}(k|k) - \hat{\mathbf{X}}_j(k|k) \right] \left[\hat{\mathbf{X}}(k|k) - \hat{\mathbf{X}}_j(k|k) \right]^T \right\}$$

1. Estimate Mixing: To initiate IMMKF processing of a measurement, the previous state estimates for the models are blended to produce mixed state estimates (see Table 1). This results in a modification of the target prediction equations and accounts for mode switching between the time of the last track update and the generation of the current report. For example, in this process some of the turning model's estimate is included in the input to the uniform motion filter, with appropriate weighting. This accounts for the possibility that a target that was turning at the conclusion of the previous measurement cycle transitions to uniform motion for the current measurement cycle. The key element in this computation is the model mixing probability μ_{ij} . This is the probability that the target was in mode i for the previous scan, given that it is now in mode j . This is used to construct the mixed state estimate $\hat{\mathbf{X}}_{0j}$ and mixed covariance estimate \mathbf{P}_{0j} for each model j (see Table 1). In the mixed covariance estimate, the factor V_{ij} models increased uncertainty due to disagreement between the model estimates.

2. Filtering: Each mixed state estimate is updated with the current radar report. The standard Kalman filter equations are used with the appropriate target motion models and plant noise models. This produces target state estimates conditioned on the target model.

3. Model Probability Updating: After each model has been updated with the radar plot, new model probabilities μ_j are computed from the model likelihood's Λ_j given in

Table 1 and the predicted model probability μ_j^- computed during the preceding mixing step.

4. Estimate Combination: The state and covariance estimates for all of the models are combined using the updated model probability weights μ_j . There is an extra contribution to the covariance $\mathbf{P}(k|k)$ generated by disagreement between the models from the term $\left[\hat{\mathbf{X}}(k|k) - \hat{\mathbf{X}}_j(k|k) \right] \left[\hat{\mathbf{X}}(k|k) - \hat{\mathbf{X}}_j(k|k) \right]^T$, similar in form to the mixed state covariance estimate $\mathbf{P}_{0j}(k|k)$. Each separate model state is maintained internally by the system. The combined estimate is only required for output of the system track to other elements of the air defense system.

The important quantity required for discrimination gain calculations is the IMMKF combined state density at the conclusion of each measurement cycle given by $p(\mathbf{X}, j|Z) = \mu_j p(\mathbf{X}|Z)$ where

$$p(\mathbf{X}|Z) = |2\pi\mathbf{P}(k|k)|^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{X} - \hat{\mathbf{X}}(k|k))^T \mathbf{P}(k|k)^{-1}(\mathbf{X} - \hat{\mathbf{X}}(k|k))\right)$$

with the other quantities defined in Table 1.

2.2 Discrimination Gain and the IMMKF

To compute the expected discrimination gain using an IMMKF filter, suppose that we have a set of measurements Z made up to time t_{k-1} . The basic idea is to compute the expected discrimination between the predicted density when no observation is made and the density obtained when a particular sensor is used. For multivariate Gaussian $q_0(\mathbf{X})$ and $q_1(\mathbf{X})$ with means \mathbf{X}_0 and \mathbf{X}_1 and variances Σ_0 and Σ_1 , the discrimination q_1 with respect to q_0 is [Blahut]

$$\begin{aligned} L(q_1; q_0) &= \int d\mathbf{X} q_1(\mathbf{X}) \log(q_1(\mathbf{X}) / q_0(\mathbf{X})) \\ &= \frac{1}{2} \text{tr} \left[\Sigma_0^{-1} \left(\Sigma_1 - \Sigma_0 + (\mathbf{X}_1 - \mathbf{X}_0)(\mathbf{X}_1 - \mathbf{X}_0)^T \right) \right] - \frac{1}{2} \log \frac{|\Sigma_1|}{|\Sigma_0|} \end{aligned}$$

To evaluate the relative utility of an observation made at time t_k , first note that if no observation is made then the target state density is obtained from the IMMKF prediction equations alone. Let $p_0(\mathbf{X}, j|Z) = \mu_j^- p_0(\mathbf{X}|Z)$ denote the predicted density. Then

$$p_0(\mathbf{X}|Z) = |2\pi\mathbf{P}(k|k-1)|^{-1/2} \exp \left(-\frac{1}{2} \left(\mathbf{X} - \hat{\mathbf{X}}(k|k-1) \right)^T \mathbf{P}(k|k-1)^{-1} \left(\mathbf{X} - \hat{\mathbf{X}}(k|k-1) \right) \right)$$

where $\mathbf{P}(k|k-1) = \sum_j \mu_j^- \mathbf{P}_j(k|k-1)$, $\hat{\mathbf{X}}(k|k-1) = \sum_j \mu_j^- \hat{\mathbf{X}}_j(k|k-1)$ and μ_j^- is given in Table 1.

Now if a new measurement z obtained with a particular sensor is received at time t_k , the new observation set is $Z' = \{z\} \cup Z$. The IMMKF computes the joint conditional density for a target to be in state \mathbf{X} and flight mode j , $p(\mathbf{X}, j|Z) = \mu_j p(\mathbf{X}|Z)$. The discrimination of $p(\mathbf{X}, j|Z)$ with respect to $p_0(\mathbf{X}, j|Z)$ is

$$L = \sum_j \int d\mathbf{X} p(j, \mathbf{X}|Z') \ln(p(j, \mathbf{X}|Z') / p_0(j, \mathbf{X}|Z))$$

Note that this depends on the particular observation z . For sensor management, we require the expected value of this quantity with respect to the conditional density $p(z|Z)$. This is

$$\mathcal{E}L = \sum_j \int d\mathbf{X} dz p(z|Z) p(j, \mathbf{X}|Z') \ln(p(j, \mathbf{X}|Z') / p_0(j, \mathbf{X}|Z))$$

The integral dz over the new observation density is complicated. However $p(z|Z)$ is sharply peaked about the expected value of the observation for scan k

$\hat{z}^{k|k-1} = H \sum_j \mu_j^- \mathbf{X}_j^{k|k-1}$. To lowest order we may approximate this distribution by a delta function at $\hat{z}^{k|k-1}$. Then defining $\hat{Z}' = \{\hat{z}^{k|k-1}\} \cup Z$, we can approximate

$$\varepsilon L \approx \sum_j \int d\mathbf{X} p(j, \mathbf{X} | \hat{Z}') \ln \left(p(j, \mathbf{X} | \hat{Z}') / p_0(j, \mathbf{X} | Z) \right).$$

Since $p(\mathbf{X} | Z)$ and μ_j are separately normalized

$$\varepsilon L = \sum_j \mu_j \ln(\mu_j / \mu_{0j}) + \int d\mathbf{X} p(\mathbf{X} | \hat{Z}') \ln \left(p(\mathbf{X} | \hat{Z}') / p_0(\mathbf{X} | Z) \right)$$

Using the fact that $p(\mathbf{X}, j | Z)$ and $p_0(\mathbf{X}, j | Z)$ are Gaussian, we obtain

$$\begin{aligned} \varepsilon L = & \sum_j \left\{ \mu_j \ln(\mu_j / \mu_{0j}) \right. \\ & + \frac{1}{2} \text{tr} \left[\mathbf{P}(k|k-1)^{-1} \left(\mathbf{P}(k|k) - \mathbf{P}(k|k-1) + (\hat{\mathbf{X}}(k|k) - \hat{\mathbf{X}}(k|k-1)) (\hat{\mathbf{X}}(k|k) - \hat{\mathbf{X}}(k|k-1))^T \right) \right] \\ & \left. - \frac{1}{2} \log \frac{|\mathbf{P}(k|k)|}{|\mathbf{P}(k|k-1)|} \right\}. \end{aligned}$$

The gain εL depends on which sensor is used through the sensor covariance dependence of $\mathbf{P}(k|k)$. This expression, evaluated for each target sensor pair, is the discrimination gain used here as a sensor effectiveness metric.

2.3 Optimization

Given the gains εL computed above, we must now assign sensors to targets. The sensors are indexed $s = 1, \dots, S$ and the targets are indexed $t = 1, \dots, T$. Each sensor has a fixed capacity τ_s . This is the maximum number of targets that can be sensed on each sensor scan. The discrimination gain when sensor s is assigned to target t is denoted G_{st} . Our objective is to maximize the total gain across all of the targets.

The optimization problem considered in this paper is one of assigning sensors to targets in such a way that the one step ahead discrimination gain of the track set is maximized. To do this, the one step ahead predicted discrimination gain G_{ij} of a track j after it is updated with covariance data from sensor i is determined for all pairs i, j . This can be formulated as a Linear Programming problem with suitable constraints. The next step in the formulation is to consider the situation where more than one sensor may be assigned to the same target. As has been previously pointed out [Nash], one method is to construct pseudo sensors comprised of combinations of the basic sensors. This allows any

combination of sensors to be assigned to a single target simply by considering that a single pseudo sensor has been assigned to it. The number of "sensors" is thus increased from S to $2^S - 1$. The problem is now to make an assignment of these pseudo sensors to the targets in an optimal way.

There are however constraints on this assignment. One of the most important is the maximum tracking capacity of a sensor. That is, given a specified time interval, some sensors can scan a certain volume of space and also track a specified number of targets. They cannot exceed this maximum tracking capacity in the specified time period. If this maximum tracking capacity is known for each of the basic sensors, then this must be accounted for when the pseudo sensors are assigned. Previous workers have not handled this accounting correctly [Nash] when they define additional maximum tracking capacities for the pseudo sensors. Surprisingly, The original maximum tracking capacities are all that is needed.

The constraints can be handled exactly as follows. Let the basic sensors be numbered from 1 to S . Let the pseudo sensors be numbered from $S + 1$ up to $2^S - 1$. For each basic sensor k , let $J(k)$ be the set of integers consisting of k and the integer numbers of the pseudo sensors which contain sensor k in their combination. There will be 2^{S-1} integers in each set $J(k)$. These sets $J(k)$ will appear in the constraints that are given for the Linear Programming formulation to our optimal assignment problem.

Example 1.

For example, with three basic sensors, $s = 3$, and $2^3 - 1 = 7$. Let $S1, S2, S3, S4, S5, S6$, and $S7$ be the designations of the sensors with $S4 = \{S1, S2\}$, $S5 = \{S1, S3\}$, $S6 = \{S2, S3\}$, and $S7 = \{S1, S2, S3\}$. The integer sets $J(k)$ then contain $2^{3-1} = 4$ integers and are then: $J(1) = \{1, 4, 5, 7\}$, $J(2) = \{2, 4, 6, 7\}$, $J(3) = \{3, 5, 6, 7\}$.

We can now state the linear programming problem.

$$\text{maximize } C = \sum_{i=1}^{2^s-1} \sum_{j=1}^t G_{i j} x_{i j} \tag{2.3.1}$$

subject to the constraints

$$\sum_{i=1}^{2^s-1} x_{i j} \leq 1 \quad \text{for } j = 1, L, t \tag{2.3.2}$$

$$\sum_{i \in J(k)} \sum_{j=1}^t x_{i j} \leq \tau_k \quad \text{for } k = 1, L, s \tag{2.3.3}$$

$$x_{i j} \geq 0 \quad \text{for all pairs } i j \tag{2.3.4}$$

and where τ_k is the maximum tracking capacity of the basic sensor k . In the LP solution, each x_{st} will be either 0 or 1. When $x_{st} = 1$, sensor s is assigned to target t .

Equation (2.3.2) fixes each j and sums over the sensors including the pseudo sensors. The fact that this is constrained to be less than or equal to one means that only one or none of the sensors (recall, the pseudo sensors are combinations) will be assigned to target j . Equation (2.3.1) clearly sums the G_{ij} for which there is a sensor assigned because the x_{ij} are zero or one. Equation (2.3.3), by virtue of the sets $J(k)$, insures that the number of targets tracked by the assignment sensors does not exceed the maximum tracking capacity of the basic sensor k . If there are fewer targets than the maximum sensor tracking capacity, then some pseudo sensors will be assigned to individual targets. If there are more targets than the total tracking capacity, then Equation (2.3.2) allows some targets to not have sensors assigned.

3. Test Results Typical test results for a 2 sensor, 8 target test case are shown in Figure 1-3. For this test case, the two sensors are at fixed positions. Sensor locations and other system parameters are shown in Table 2. Two scenarios were examined. The first test simulates the situation where two sensors must track 8 closing targets. For the second case, the sensor geometry is changed to reflect targets transiting along a political boundary. The sensors provides measurements to the central fusion center at 1 second intervals, referred to as a scan. In each scan, we assume that the sensor can observe two targets, so half the targets are observed on each scan. The targets are numbered 1-8, with target 1 initially the southern-most target and target 8 the northern-most. Target 4 undergoes a standard rate turn. The remaining targets are in straight line flight. Figure 3 shows the true turn rate and the turn rate estimate for Target 4 averaged over 10 runs of the scenario. There is some lag in the filter and it correctly estimates the peak turn rate at about .03 radians/sec.

Figure 3 shows the expected discrimination gain vs. scan k averaged over the same 10 runs of the test scenario. The upper curve is the expected gain of the optimal sensor for the maneuvering target. The lower curve is the optimal gain averaged over the 7 non-maneuvering targets. As expected, the maneuvering target generally has greater uncertainty and has greater expected discrimination gain as a result. This is confirmed by the fact that the maneuvering target is observed nearly 60% of the time, while the other targets are observed somewhat less than 50% of the time. Note that even though there is significant turn-rate lag as shown in Figure 2, the expected gain begins to grow as soon as the target initiates its maneuver in scan 20.

To quantify the efficacy of discrimination gain as a measure of sensor effectiveness we computed the RMS position error and total error covariance for three alternative sensor management schemes. These are given in Tables 3 and 4. The rows labeled "DG" use discrimination gain with the simplex optimization. The rows labeled "Fixed" use a fixed round-robin sensing schedule. There are many possible alternatives. We chose to use the schedule where sensor 1 observes target pair (1,5), then (2,6), (3,7) while sensor 2 observes targets (3,7), then (4,8) and so on. This is simple to implement and spreads the observations out among the targets. This coordinates sensing but does not adaptively compute the sensor effectiveness. For the rows labeled "Rand" each sensor selects targets

at random for each scan cycle. This is the worst case approach in that it assumes that the sensing is completely uncoordinated.

In order to establish a baseline for comparison, we assumed that all of the scheduling schemes maintain central tracks updated via report fusion as soon as measurements are available. Although this may be unrealistic for some applications, it enables simple direct comparison of the three schemes.

We examined two measures of effectiveness for this sensor management study. The first is simply the RMS position error for each target after each measurement update. For scans where a particular target is not updated, its predicted position $\hat{\mathbf{X}}(k|k-1)$ is computed and compared with ground truth to obtain the RMS error.

For some tracking applications other quantities such as the velocity accuracy are of primary interest. A simple metric that is also sensitive to this more global aspect of the problem can be obtained from the covariance

$$\mathbf{C} = (\mathbf{X}(k) - \hat{\mathbf{X}}(k|k))(\mathbf{X}(k) - \hat{\mathbf{X}}(k|k))^T,$$

which is a 5×5 matrix. The determinant has units $[\mathbf{C}] = m^8 / s^6$ and is proportional to the volume of the uncertainty ellipsoid of the target estimate. The volume of total uncertainty ellipsoid for the 8-target problem (in a 40-dimensional space) is proportional to the product of the single target volumes.

Global discrimination optimization yields the best performance with 80 m RMS position error averaged across all the targets while random assignment yields the worst with 90 m average RMS error. Also, the greatest gain was obtained against the maneuvering target where discrimination yielded 93 m RMS position error while the random approach suffered from a 109 m RMS error for the maneuvering target. Similar results were obtained for both test scenarios. These qualitative results are reflected in the covariance determinants for these scheduling methods, as well. Interestingly, the round-robin sensor scheduling was able to provide significant performance improvement relative to the random scheduling, although its performance still fell below that of discrimination gain. This serves to emphasize the utility of multisensor scheduling.




Figure 1 – XY Plot (meters) of test scenario for a collection of targets closing with a pair of sensor platforms S₁ and S₂ showing target trajectories and sensor detections. All targets enter from the left. Target 4 (indicated) executes a standard rate turn during the scenario.




Figure 2 - IMMKF Turn Rate Estimate Turn rate (dashed) and turn rate estimate for target 4 averaged over 10 scenario runs.




Figure 3 - Expected Discrimination Gain with Global Optimization Expected discrimination is greater for maneuvering target leading to more sensor resource allocation and improved tracking performance. Expected discrimination for maneuvering target (target 4) - upper curve (solid line). Average expected discrimination for non-maneuvering targets - lower curve (dashed line). Both curves are average results obtained over 10 runs of the closing target scenario. Note that the gain for the maneuvering target increases relative to the non-maneuvering targets as soon as the maneuver is initiated at scan 20.

$\sigma_{\theta} = 1^{\circ}$ Sensor 1 location (x,y)	11 km -5 km
Sensor 2 location - closing targets (x,y)	11 km 11 km
Sensor 2 location - transiting targets (x,y)	-5 km -5 km
Sensor update rate	1 sec ⁻¹
Target speed	100 m/sec
σ_{range}	100 m
σ_{θ}	1°

Table 2 – Sensor and target parameters used to test sensor effectiveness measure.

	Tgt1	Tgt2	Tgt3	Tgt4	Tgt5	Tgt6	Tgt7	Tgt8	Avg
DG RMS (m)	85	76	83	93	76	74	73	75	80
DG Det	25	14	19	1470	14	14	454	10	252
Fixed RMS	90	76	81	110	78	76	79	84	85
Fixed Det	32	14	19	2824	19	14	525	25	434
Rand RMS	82	90	83	109	83	95	95	80	90
Rand Det	19	19	19	2293	19	32	902	19	415

Table 3 -- Test results obtained for 10 runs of 2-sensor closing target scenario. Use of discrimination gain (DG) to select the optimal sensor assignment reduces the average track error as measured by both the RMS position error and the determinant of the average covariance (highlighted column at right). DG is compared to a fixed sensor schedule and to a random sensor schedule (Rand) that simulates the situation where there is no coordination between the sensors. DG provides the most benefit against target 4 which maneuvers during the scenario.

	Tgt1	Tgt2	Tgt3	Tgt4	Tgt5	Tgt6	Tgt7	Tgt8	Avg
DG RMS (m)	77	81	80	102	87	73	80	88	84
DG Det	11	20	14	1850	25	7.6	454	24	300
Fixed RMS	91	86	85	111	78	77	85	91	89
Fixed Det	32	19	19	2824	14	11	19	32	371
Rand RMS	85	94	92	116	86	105	93	90	95
Rand Det	19	32	25	2824	19	52	792	19	472

Table 4 – Test results obtained for 10 runs of transiting target scenario are qualitatively similar to results for closing target scenario (Table 3).

4. Discussion To summarize, we have demonstrated that the formalism of the IMMKF can be used as a framework for computing the expected information gain when multiple sensors observe a collection of airborne targets. When combined with the IMMKF, discrimination gain is able to recognize target maneuvers and respond by allocating additional sensor resource to the maneuvering target. The sensor target assignment that optimizes the total discrimination across all of the targets in a surveillance volume can be readily computed. Application of this optimal assignment results in improved tracking performance relative to random assignment of sensors to targets or fixed round-robin scheduling.

The IMMKF used in this assessment was designed for an air traffic control application. For use in air-to-air tracking applications, different versions of the IMMKF must be developed to treat target motion in 3-dimensions and to incorporate range-rate measurements. Another topic that must be examined is the use of simpler filters that can still support discrimination gain evaluation. For example, the adaptive single model Kalman filter can also be used in this application but was not studied here. It may be that for some applications such as tracking low-priority targets, the computational simplicity of the single model filter outweighs its disadvantages relative to the IMMKF. This study has assumed that the target detection probability is 1 and that there is no clutter. Furthermore, it was assumed that the track is a confirmed target. However, if the track is a tentative target, then additional information is obtained when the target is observed with the sensor. These effects can be included in the discrimination gain calculations.

One might expect the improved performance for the maneuvering target to come at the expense of some loss of performance for the non-maneuvering targets. However, this is not the case. By improving sensor allocation across all of the targets, average performance is improved for both the maneuvering and non-maneuvering targets. Finally, this study has addressed sensor effectiveness independent of the tactical or strategic utility of knowledge about the targets. In order to address this aspect of the problem, discrimination must be combined with a means of assessing preferences as well.

References:

1. [Bar] Y. Bar-Shalom, and X. R. Li, *Estimation and Tracking : Principles, Techniques, and Software* Artech House Radar Library, Boston, 1993.
2. [Blahut] R. E. Blahut, *Principles and Practice of Information Theory*, Addison-Wesley, Reading, MA, 1987.
3. [Blom88] Blom, H. A. P., and Bar-Shalom, Y., (Aug 1988), "The interacting multiple model algorithm for systems with Markovian switching coefficients", *IEEE Trans. Aut. Con.*, vol AC-33, pp 780-783.

4. [Blom92] Blom, H. A. P., Hogendoorn, R. A., and van Doorn, B. A.,(1992), "Design of a Multisensor Tracking System for Advanced Air Traffic Control", in *Multitarget-Multisensor Tracking: Applications and Advances, Vol II*, Y. Bar-Shalom, Ed., Norwood, MA: Artech House, pp 129.
5. [Kas96a] K. Kastella and M. Biscuso, *Tracking Algorithms for Air Traffic Control Applications*, Air Traffic Control Quarterly, pp. 19-43, Vol. 3 (1), Jan. 1996.
6. [Kas96b] K. Kastella and S. Musick, *The Search for Optimal Sensor Management*, *SPIE Proceedings*, Vol. 2759, pp. 318-329, April, 1996.
7. [Kas96c] K. Kastella, "Discrimination Gain for Sensor Management in Multitarget Detection and Tracking", *Proc. 1996 IMACS Multiconference on Computational Engineering in Systems Applications (CESA '96); Symposium on Control, Optimization and Supervision, Vol. 1*, Lille, France, July 9-12, 1996. pp. 167-172.
8. [Kas97a] K. Kastella, "Discrimination gain to optimize detection and classification", *Proc. IEEE Trans. Sys. Man. Cyb., Part A*, vol. 27 no. 1, Jan. 1997, pp. 112-116.
9. [Kas97b] K. Kastella, "Joint multitarget probabilities for detection and tracking", *SPIE Proceedings, Acquisition, Tracking and Pointing XI*, 21-25 April, 1997, Orlando, FL.
10. [Li] X. R. Li and Y. Bar-Shalom, "Design of an Interacting Multiple Model Algorithm for Air Traffic Control Tracking", *IEEE Trans. Cont. Sys. Tech.*, vol 1, pp. 186-194, Sep 1993.
11. [Mus96] S. Musick and K. Kastella, "Comparison of Sensor Management Strategies for Detection and Classification", Proceedings of the 9th National Symposium on Sensor Fusion 11-13 March 1996 Naval Postgraduate School Monterey, CA.
12. [Nash] J.M. Nash. Optimal allocation of tracking resources. *Proceedings IEEE Conference on Decision and Control* pp. 1177-1180, 1977.
13. [Press] W. H. Press, B. P. Flannery, S. A Teukolsky, and W. T. Vetterling, *Numerical Recipes in C*. Cambridge: Cambridge Univ. Press, 1988.
14. [Sch93] W. W. Schmaedeke, "Information-based Sensor Management", *SPIE Proceedings Vol. 1955, Signal Processing, Sensor Fusion, and Target Recognition II*, April, 1993.

15. [Sch93b] W. Schmaedeke and K. Kastella, "Information Theory and Sensor Management", Basic Research Group of the Joint Directors of Laboratories 1993 Symposium on Command and Control Research, National Defense University, Fort McNair, Washington, D.C. June 28 - July 1, 1993.
16. [Sch94] W. Schmaedeke and K. Kastella, "Event-averaged maximum likelihood estimation and information based sensor management", *SPIE Proceedings*, Vol. 2232, Signal Processing, Sensor Fusion, and Target Recognition III, April, 1994, pp. 91-96.
17. [Watson] G. A. Watson and W. D. Blair, "IMM Algorithm for Tracking Targets that Maneuver Through Coordinated Turns", *Proc. SPIE Signal and Data Processing of Small Targets*, 1992, Orlando, FL, Apr. 1992, pp 236-247; "Revisit Control of a Phased-Array Radar for Tracking Maneuvering Targets when Supported by a Precision Electronic Support Measures Sensor", *Proc. SPIE Signal and Data Processing of Small Targets*, 1994, Orlando, FL, Apr. 1994, pp. 448-459.
18. [Yed] M. Yeddanapudi, Y. Bar-Shalom, and K. R. Pattipati, "IMM Estimation for Multitarget-Multisensor Air Traffic Surveillance", *Proc. IEEE*, Vol 85, No. 1 Jan, 1997, pp. 80-94.