2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Risk-Based Sensor Management for Integrated Detection and Estimation

Y. Wang*, I. I. Hussein*, and R. S. Erwin†

Abstract— **In this paper, we develop a risk-based sensor management scheme for unknown object detection and process estimation under limited sensory resources. Bayesian sequential detection and estimation methods are utilized for risk analysis. The objective is to find every object of interest in the mission domain and satisfactorily estimate the associated process dynamics with minimum risks. Two types of costs are taken into account for risk evaluation, i.e., the cost of making an erroneous decision regarding object existence or its estimates, and the cost of taking more observations for a possibly better decision. The Rényi information divergence is investigated to measure the information loss in making a suboptimal sensor allocation decision, which is used to formulate the observation cost. A set of simulation results are provided to confirm the effectiveness of the proposed sensor management scheme.**

I. Introduction

Object search, detection and state estimation are challenging tasks due to the uncertainty in sensor perception. In these cases, it is crucial to manage sensors such that competing tasks are effectively assigned across multiple objects to be detected and processes to be tracked. This is especially true when the sensing resources are limited (e.g., sensory range and number of sensors) compared to the size of the mission domain and the large number of objects.

In the literature, sensor management and task allocation have been mainly studied under no resources constraints. In [1], a distributed sequential auction scheme is presented for a multi-robot search and destroy operation. The control goal is to allocate a vehicle to each object and complete the task in minimum time. In [2], the author proposes a Bayesian-based multi-target multi-sensor management scheme. The approximation strategy maximizes the expected number of targets. In [3], the authors seek to maximize the probability of finding a target with some foreknown location information in the presence of uncertainty. Coordinated search and tracking in probabilistic frameworks has been focused mainly on optimal path planning and estimation. In [4], a cooperative control scheme based on Fischer information measure is proposed for the optimal path planning of a team of uninhabited aerial vehicles (UAVs) in a ground target tracking problem. In [5], the pursuit-evasion game and map building problems are combined in a probabilistic game theoretic framework. Sub-optimal pursuit policies are presented to minimize the expected capture time. In [6], a dynamic strategy is developed to control the relative configuration of sensor teams based on particle filtering. The goal is to get optimal estimates for target tracking through sensor fusion. Besides, the development of a unified framework for search and tracking/localization problems has been studied in several work [7]–[9]. A recursive Bayesian estimator is used to fuse observations with uncertainty in sensor perception. However, none of the above work consider the problem of competing tasks and apply to cases when there is only a single target or the sensing resources are not limited. Hence, there is need to develop sensor management schemes that are able to choose optimally between competing tasks and make minimum-risk decisions under limited sensing resources.

For object detection, Bayesian sequential detection method formulated by Wald and Wolfowitz in [10] provides a strong theoretical background for risk analysis. It is a sequential hypothesis testing for a stationary discrete random variable, which allows the number of observations to vary in order to achieve optimal decisions. For process estimation, Bayesian sequential detection is extended to Bayesian sequential estimation [11]. In Bayesian sequential settings, the goal is to minimize Bayes risk, which include two types of costs: 1) the cost of making a wrong decision without taking any further observation, and 2) the cost of all the observations taken up to the decision time. Rènyi information measures is used to model the information loss in making a suboptimal sensor allocation decision and formulate the observation cost [12], [13].

We seek to develop an integrated framework for probabilistic object detection and process estimation treated as tasks competing for the same limited sensory resources via Bayesian sequential analysis.

The paper is organized as follows. A review of Bayesian sequential detection for discrete random variables is provided in Section II. Its extension to Bayesian sequential estimation for continuous random variables is developed in Section III. In Section IV, we extend the Bayesian sequential detection and estimation methods to multiple elements (cells or objects) over the entire mission domain. An integrated risk-based sensor management scheme for the detection and estimation of multiple elements is developed in Section V. We also discuss measures of expected information gain for both detection (discrete random variables) and estimation (continuous random variables). The paper is concluded with a summary of both current and future work in Section VII.

*Mechanical Engineering Department, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609, USA {yuewang,ihussein}@wpi.edu

†Air Force Research Laboratory, Space Vehicles Directorate, Kirtland Air Force Base, NM 87117, USA

978-1-4577-0079-8/11/$26.00 ©2011 AACC

II. Sequential Detection

A. State Description

Assume that we are given a domain Ω, where the unknown objects of interest to be detected and estimated are located. The domain is descritized into cells, each denoted by $\mathbf{c}_j,\ \ j = 1, \cdots, N$. The discretization is assumed to be fine enough such that at most one object can exist within a cell. Let the actual state of object existence at $\mathbf{c}_j$ be denoted by X_j, which is equal to 1 if object exists (or, more generally, there is a process at $\mathbf{c}_j$ to be estimated) and 0 if no object exists (or, equivalently, there is no process at $\mathbf{c}_j$).

B. Sensor Model

Here we employ a sensor model with a Bernoulli distribution. For cell $\mathbf{c}_j$, at time t, the sensor is assumed to give the output $Y_{j,t} = 1$ (a "positive observation", i.e., "object present") or $Y_{j,t} = 0$ (a "negative observation", i.e., "object absent"). The sensor model is given by the following general conditional probability matrix,

$$B = \begin{bmatrix} \text{Prob}(Y_{j,t} = 0|X_j = 0) & \text{Prob}(Y_{j,t} = 0|X_j = 1) \\ \text{Prob}(Y_{j,t} = 1|X_j = 0) & \text{Prob}(Y_{j,t} = 1|X_j = 1) \end{bmatrix},$$

where $\text{Prob}(Y_{j,t} = i|X_j = l),\ i, l = 0, 1$ describes the probability of measuring $Y_{j,t} = i$ given $X_j = l$. Let β be the detection probability, we have $\text{Prob}(Y_{j,t} = 0|X_j = 0) = \text{Prob}(Y_{j,t} = 1|X_j = 1) = \beta$ and $\text{Prob}(Y_{j,t} = 1|X_j = 0) = \text{Prob}(Y_{j,t} = 0|X_j = 1) = 1 - \beta$.

C. Bayesian Update Equations

As measurements are made with time, the probability of object existence at $\mathbf{c}_j$ is updated according to Bayes' rule. At time $t_0 = 0$ (before taking any measurements) we assume that we have an estimate of the prior distribution of object existence given by $\{\bar{p}_{j,0}, 1 - \bar{p}_{j,0}\}$. At time t, the sensor makes an observation $Y_{j,t}$ at a cell $\mathbf{c}_j \in \Omega$. Let $Y_{j,1:t}$ be the set of all measurements taken at cell $\mathbf{c}_j$ from time $t = 1$ up to time t. The prediction step gives the predicted probability of object existence at cell $\mathbf{c}_j$ at time $t + 1$ given all measurements made up to and including time t:

$$\begin{aligned} \bar{p}_{j,t+1} &\equiv \text{Prob}(X(\mathbf{c}_j) = 1; t + 1|Y_{j,1:t}) \\ &= \text{Prob}(X_j = 1; t|Y_{j,1:t}) = \hat{p}_{j,t}. \end{aligned} \tag{1}$$

At time t, the update step is as follows:

$$\hat{p}_{j,t} = \frac{\bar{p}_{j,t}\text{Prob}(Y_{j,t}|X(\mathbf{c}_j) = 1; t)}{\text{Prob}(Y_{j,t}|Y_{j,1:t-1})}, \tag{2}$$

where the denominator in the last expression in Equation (2) is given by

$$\begin{aligned} &\text{Prob}(Y_{j,t}|Y_{j,1:t-1}) \\ &= \begin{cases} (1-\beta)(1-\bar{p}_{j,t}) + \beta\bar{p}_{j,t} & \text{for } Y_{j,t} = 1 \\ \beta(1-\bar{p}_{j,t}) + (1-\beta)\bar{p}_{j,t} & \text{for } Y_{j,t} = 0 \end{cases}. \end{aligned}$$

Substituting this into Equation (2), for the update equation we finally obtain

$$\hat{p}_{j,t} = \begin{cases} \frac{\beta\bar{p}_{j,t}}{(1-\beta)(1-\bar{p}_{j,t})+\beta\bar{p}_{j,t}} & \text{if } Y_{j,t} = 1 \\ \frac{(1-\beta)\bar{p}_{j,t}}{\beta(1-\bar{p}_{j,t})+(1-\beta)\bar{p}_{j,t}} & \text{if } Y_{j,t} = 0 \end{cases}. \tag{3}$$

Equations (1) and (3) constitute the belief prediction and update equations for the probability mass function (p.m.f.) of object presence at cell $\mathbf{c}_j$ at time t.

D. Bayesian Sequential Detection for A Single Cell

For a single cell $\mathbf{c}_j$, the goal of Bayesian sequential detection is to determine the actual state X_j with minimum risk given a sequence of observations up to time t. In this work, we will consider two types of costs when making a decision: 1) the expected cost of making an erroneous decision, and 2) the expected cost of taking future new observations for a possibly better decision. When the risk of not taking further measurements is lower, the sensor will stop and make a decision regarding object existence at cell $\mathbf{c}_j$.

1) Decision Cost Assignment: We first introduce the hypotheses: $\mathcal{H}_{j,0}$: the null hypothesis that $X_j = 0$; and $\mathcal{H}_{j,1}$: the alternative hypothesis that $X_j = 1$. Define the cost of accepting hypothesis $\mathcal{H}_{j,i}$ when the actual state at $\tilde{\mathbf{c}}$ is $X_j = k$ as C_{ik}. Using a uniform cost assignment (UCA), the decision cost matrix is given by

$$C_{ik} = \begin{cases} 0 & \text{if } i = k \\ c_d^\tau & \text{if } i \neq k \end{cases},\ \tau \geq 0,$$

where $i = 0, 1$ correspond to accepting $\mathcal{H}_{j,0}$ and $\mathcal{H}_{j,1}$, and $k = 0, 1$ correspond to the true state $X_j = 0$ and $X_j = 1$. Here, $c_d^\tau > 0$ is the cost of making the wrong decision at time $\tau \geq 0$ indicating the number of observations.

A detection estimator is a map $\hat{X}_{j,t+\tau}$ that maps a sequence of observations $Y_{j,1:t+\tau}$ into a decision to accept $\mathcal{H}_{j,0}$ or $\mathcal{H}_{j,1}$, $\tau \geq 0$. Let the notation $C(\hat{X}_{j,t+\tau}(Y_{j,1:t+\tau}), X_{j,t+\tau})$ denote the cost of using estimator $\hat{X}$ given that the actual state at $\mathbf{c}_j$ at time $t + \tau$ is $X_{j,t+\tau}$.

2) Detection Decision-Making: At $\mathbf{c}_j$, if a decision regarding object existence is made without taking any further observations, i.e., the observation number $\tau = 0$, we define a Bayes risk r_j as the expected cost of accepting the wrong hypothesis over all possible states conditioned on all previous observations. The Bayes risk associated with accepting $\mathcal{H}_{j,0}$ is given by:

$$r_j(\hat{X}^1_{j,t}, \tau = 0) = E_{X_{j,t}|Y_{j,1:t}}[C(\hat{X}^1_{j,t}, X_{j,t})] = c_d^0\hat{p}_{j,t}. \tag{4}$$

Similarly, the Bayes risk associated with accepting $\mathcal{H}_{j,1}$ is given by:

$$r_j(\hat{X}^2_{j,t}, \tau = 0) = c_d^0(1 - \hat{p}_{j,t}). \tag{5}$$

Next, we derive the risk associated with delaying the decision and keeping taking a measurement at $\mathbf{c}_j$. Since we do not have a measurement at time $t+1$ yet, define the conditional risk, $R_{X_{j,t+1}}(\hat{X}_{j,t+1}(Y_{j,t+1}), \tau = 1)$, over all possible measurement realizations at time $t + 1$ associate with the estimator $\hat{X}_{j,t+1}(Y_{j,t+1})$ when the state is $X_{j,t+1}$ as

$$\begin{aligned} &R_{X_{j,t+1}}(\hat{X}_{j,t+1}(Y_{j,t+1}), \tau = 1) \\ &= E_{Y_{j,t+1}|X_{j,t+1}}[C(\hat{X}_{j,t+1}(Y_{j,t+1}), X_{j,t+1})] + c_{j,\text{obs}}, \end{aligned}$$

where $c_{j,\text{obs}}$ is the cost for taking one more observation at $\mathbf{c}_j$ at the next time step $t + 1$.

In this case, the Bayes risk r_j is define as the expected

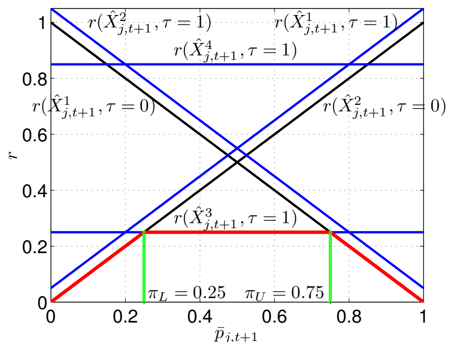

Fig. 1. Minimum Bayes risk curve (red segments) with $\beta = 0.8$, $c_{j,\text{obs}} = 0.05$ and $c_d^0 = c_d^1 = 1$.

conditional cost $R_{X_{j,t+1}}$ over all possible true states:

$$r_j(\hat{X}_{j,t+1}(Y_{j,t+1}), \tau = 1)$$
$$= E_{X_{j,t+1}|Y_{j,1:t}}[R_{X_{j,t+1}}(\hat{X}_{j,t+1}(Y_{j,t+1}), \tau = 1)].$$

When $\tau = 1$, there are 4 possible estimators available. We assume that the decision cost associated with taking one more observation ($\tau = 1$) is $c_d^1 > 0$.

1) Estimator 1 ($\hat{X}^1_{j,t+1}(Y_{j,t+1})$): Accept $\mathcal{H}_{j,0}$ whether observing $Y_{j,t+1} = 0$ or 1.
$$r_j(\hat{X}^1_{j,t+1}(Y_{j,t+1}), \tau = 1) = c_d^1 \bar{p}_{j,t+1} + c_{j,\text{obs}}. \quad (6)$$
2) Estimator 2 ($\hat{X}^2_{j,t+1}(Y_{j,t+1})$): Accept $\mathcal{H}_{j,1}$ whether observing $Y_{j,t+1} = 0$ or 1.
$$r_j(\hat{X}^2_{j,t+1}(Y_{j,t+1}), \tau = 1)$$
$$= c_d^1(1 - \bar{p}_{j,t+1}) + c_{j,\text{obs}}. \quad (7)$$
3) Estimator 3 ($\hat{X}^3_{j,t+1}(Y_{j,t+1})$): Accept $\mathcal{H}_{j,i}$ when observing $Y_{j,t+1} = i,\ i = 0, 1$.
$$r_j(\hat{X}^3_{j,t+1}(Y_{j,t+1}), \tau = 1) = c_d^1(1 - \beta) + c_{j,\text{obs}}. \quad (8)$$
4) Estimator 4 ($\hat{X}^4_{j,t+1}(Y_{j,t+1})$): Accept $\mathcal{H}_{j,i}$ when observing $Y_{j,t+1} = 1 - i,\ i = 0, 1$.
$$r_j(\hat{X}^4_{j,t+1}(Y_{j,t+1}), \tau = 1) = c_d^1 \beta + c_{j,\text{obs}}. \quad (9)$$

The optimal decision is the one that gives the minimum risk:

$$r_j^*(\bar{p}_{j,t+\tau}) = \min_{\hat{X}_{j,t+\tau},\tau} r_j(\hat{X}_{j,t+\tau}(Y_{j,t+\tau}), \tau).$$

3) Simulation Results: Figure 1 shows the minimum Bayes risk curve r_j^* as a function of $\bar{p}_{j,t+1}$ with $\beta = 0.8$, $c_{j,\text{obs}} = 0.05$ and $c_d^0 = c_d^1 = 1$ at a cell $\mathbf{c}_j \in \Omega$. The Bayes risk functions (4), (5), and (6)-(9) that construct r_j^* are also shown in the figure.

Note that the intersection probabilities $\pi_L = 0.25$ and $\pi_U = 0.75$ give the threshold points of whether to take one more observation or not. If $\bar{p}_{j,t+1} < \pi_L$, accepting $\mathcal{H}_0$ with $\tau = 0$ results in the minimum risk, and if $\bar{p}_{j,t+1} > \pi_U$, accepting $\mathcal{H}_1$ with $\tau = 0$ results in the minimum risk. If $\tau^* \neq 0$, the sensor will postpone making a decision and take one more observation.

III. Sequential Estimation

A. System Model: Single Sensor and a Single Process

In this section we develop risk analysis tools for Bayesian sequential estimation. A linear system is assumed for the process associated with a detected object at $\mathbf{c}_j$ (i.e., if we accept that $X_j = 1$, omitting the process index j):

$$\mathbf{x}_{t+1} = \mathbf{F}_t \mathbf{x}_t + \mathbf{v}_t,$$
$$\mathbf{y}_t = \mathbf{H}_t \mathbf{x}_t + \mathbf{w}_t,$$

where $\{\mathbf{x}_t \in \mathbb{R}^n, t \in \mathbb{N}\}$ is the state sequence, $\mathbf{F}_t \in \mathbb{R}^{n\times n}$ is the state matrix, $\{\mathbf{v}_t \in \mathbb{R}^n, t \in \mathbb{N}\}$ is the i.i.d. Gaussian process noise sequence with zero mean and positive semi-definite covariance $\mathbf{Q}_t \in \mathbb{R}^{n\times n}$, $\{\mathbf{y}_t \in \mathbb{R}^m, t \in \mathbb{N}\}$ is the measurement sequence, $\mathbf{H}_t \in \mathbb{R}^{m\times n}$ is the output matrix, and $\{\mathbf{w}_t \in \mathbb{R}^m, t \in \mathbb{N}\}$ is the i.i.d. Gaussian measurement noise sequence with zero mean and positive definite covariance $\mathbf{R}_t \in \mathbb{R}^{m\times m}$. The initial condition for the state is assumed Gaussian with mean $\bar{\mathbf{x}}_0$ and positive definite covariance $\mathbf{P}_0 \in \mathbb{R}^{n\times n}$. We will assume that the initial state, process noise, and measurement noise are all uncorrelated.

B. Sequential State Estimation

In sequential estimation decision-making (i.e., the estimation plus the decision whether to take more measurements or not), we assume that the estimation problem is solved (here the Kalman filter is used for estimation) and the only question to be addressed is whether we accept the estimate as the true state (and, hence, stop taking additional measurements) or we take (at least) one more measurement. Hence, the list of decisions are: (1) stop, and take no more measurements, and (2) take one more measurement.

C. The State Estimation Problem

For the estimation problem, we use the Kalman filter since it is the optimal filter for linear Gaussian systems. At time step t, the state and error covariance matrix prediction equations are given by [14]

$$\bar{\mathbf{x}}_t = \mathbf{F}_{t-1}\hat{\mathbf{x}}_{t-1},$$
$$\bar{\mathbf{P}}_t = \mathbf{Q}_{t-1} + \mathbf{F}_{t-1}\hat{\mathbf{P}}_{t-1}\mathbf{F}_{t-1}^{\mathrm{T}}.$$

The current posterior state estimate and error covariance matrix are given by:

$$\hat{\mathbf{x}}_t = \bar{\mathbf{x}}_t + \mathbf{K}_t\left(\mathbf{y}_t - \mathbf{H}_t\bar{\mathbf{x}}_t\right),$$
$$\hat{\mathbf{P}}_t = \left(\mathbf{I} - \mathbf{K}_t\mathbf{H}_t\right)\bar{\mathbf{P}}_t,$$

where

$$\mathbf{K}_t = \bar{\mathbf{P}}_t\mathbf{H}_t^{\mathrm{T}}\left(\mathbf{H}_t\bar{\mathbf{P}}_t\mathbf{H}_t^{\mathrm{T}} + \mathbf{R}_t\right)^{-1}.$$

1) Estimation Error Cost assignment: The cost of choosing a specific estimate $\mathbf{x}_t^e(\mathbf{y}_t)$ (we will omit the dependence on $\mathbf{y}_t$ for notational brevity) when the actual state is $\mathbf{x}_t$ is denoted by $C(\mathbf{x}_t^e, \mathbf{x}_t)$. We can set $C(\mathbf{x}_t^e, \mathbf{x}_t) = c_e^\tau \left\|\mathbf{x}_t^e - \mathbf{x}_t\right\|^2$ (quadratic cost with $c_e^\tau > 0$ being some τ-dependent cost value and $\tau \geq 0$ indicating the number of future observations), or UCA:

$$C(\mathbf{x}_t^e, \mathbf{x}_t) = \begin{cases} 0 & \|\mathbf{x}_t^e - \mathbf{x}_t\| \leq \epsilon \\ c_e^\tau & \|\mathbf{x}_t^e - \mathbf{x}_t\| > \epsilon \end{cases}, \quad (10)$$

where $\epsilon > 0$ is some preset small interval. In this work, for $\mathbf{x}_t^e$, we use the updated Kalman Filter estimate $\hat{\mathbf{x}}$.

2) Estimation Decision-Making: At time t, after making a measurement $\mathbf{y}_t$, if we decide not to take any more measurements, as in the sequential detection approach, the

Bayes risk is defined as the expected value (over all possible true states, conditioned on all previous measurements) of the cost of choosing the estimate $\hat{\mathbf{x}}_t$:

$$r(\hat{\mathbf{x}}_t, \tau = 0) = E_{\mathbf{x}_t|\mathbf{y}_{1:t}}[C(\hat{\mathbf{x}}_t, \mathbf{x}_t)]. \quad (11)$$

If we assume a quadratic cost assignment, we have

$$r(\hat{\mathbf{x}}_t, \tau = 0) = c_e^0 \text{Tr}\left[\hat{\mathbf{P}}_t\right],$$

where $c_e^0 > 0$ is the estimation cost when the sensor does not take an observation (i,e., $\tau = 0$).

We also need to compute the (expected) risk associated with taking more observations ($\tau \geq 1$). Since we do not have measurements over time period $t+1 : t+\tau$ yet, define the conditional risk, $R_{\mathbf{x}_{t+1:t+\tau}}(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+1:t+\tau}), \tau)$ over all possible measurement realizations over $t+1 : t+\tau$ given the state $\mathbf{x}_{t+\tau}$ at time $t+\tau$ as

$$R_{\mathbf{x}_{t+1:t+\tau}}(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+1:t+\tau}, \tau))$$
$$= E_{\mathbf{y}_{t+1:t+\tau}|\mathbf{x}_{t+1:t+\tau}}[C(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+1:t+\tau}), \mathbf{x}_{t+\tau})] + \kappa\tau c_{\text{obs}},$$

where c_{obs} is the cost of taking an observation and $\kappa > 0$ is some scaling parameter. The Bayes risk is defined as the weighted conditional risk $R_{\mathbf{x}_{t+1:t+\tau}}$, weighted by the predicted density function $p(\mathbf{x}_{t+1:t+\tau}|\mathbf{y}_{1:t})$ at time $t+1 : t+\tau$:

$$r(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+1:t+\tau}), \tau)$$
$$= E_{\mathbf{x}_{t+1:t+\tau}|\mathbf{y}_{1:t}}\left[R_{\mathbf{x}_{t+1:t+\tau}}(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+1:t+\tau}), \tau)\right].$$

If we choose a quadratic error cost assignment, the expected Bayes risk if we take more measurements is given by

$$r(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+1:t+\tau}), \tau) = c_e^\tau \text{Tr}\left[\hat{\mathbf{P}}_{t+\tau}\right] + \kappa\tau c_{\text{obs}}.$$

If we choose UCA, the computation of r can be performed using Monte-Carlo approximation. If the state dimension is 1, the expected Bayes risk is given by

$$r(\hat{\mathbf{x}}_{t+\tau}(\mathbf{y}_{t+\tau}), \tau) = c_e^\tau \left(1 - \text{Erf}\left(\frac{\epsilon}{2\sqrt{2\hat{\mathbf{P}}_{t+\tau}}}\right)\right) + \kappa\tau c_{\text{obs}}, \quad (12)$$

where $\tau = 0, 1$ and

$$\text{Erf}(\cdot) = \frac{2}{\sqrt{\pi}} \int_0^{(\cdot)} e^{-t^2} \text{d}t$$

is the error function and ϵ is an error bound as indicated in Equation (10).

The optimal decision corresponds to a particular observation number τ^* that yields minimum Bayes risk:

$$\tau^* = \text{argmin}_\tau r(\hat{\mathbf{x}}_{t+\tau}, \tau).$$

IV. Extension to Multiple Elements

We now extend the Bayes risk formulation for a singe cell in Section II and a single object in Section III to the case when a sensor chooses among multiple elements (cell or object).

Let us first consider the Bayes detection risks at a cell $\mathbf{c}_j$. The risks associated with making a detection decision at $\mathbf{c}_j$ at the current time step t do not change in multi-element case because this is the decision associated with cell $\mathbf{c}_j$ itself. Hence, they are the same as Equations (4) and (5). The Bayes risk r_k associated with observing element $\mathbf{c}_k$ (including the possibility of choosing $\mathbf{c}_j$ again) at the next time step $t+1$ given that the sensor is observing $\mathbf{c}_j$ at t is define as:

$$r_k(\hat{X}_{k,t+1}(Y_{k,t+1}), \tau = 1)$$
$$= E_{X_{k,t+1}|Y_{k,1:t}}[R_{X_{k,t+1}}(\hat{X}_{k,t+1}(Y_{k,t+1}), \tau = 1)],$$

where the conditional risk is given by:

$$R_{X_{k,t+1}}(\hat{X}_{k,t+1}(Y_{k,t+1}), \tau = 1)$$
$$= E_{Y_{k,t+1}|X_{k,t+1}}[C(\hat{X}_{k,t+1}(Y_{k,t+1}), X_{k,t+1})] + c^k_{j,\text{obs}},$$

where $c^k_{j,\text{obs}}$ is the observation cost assigned for cell $\mathbf{c}_j$ if it decides to take an observation at element $\mathbf{c}_k$ at the next time step $t+1$. The optimal decision is then to choose a combination of $\hat{X}_{j,t+\tau}$, $\tau = 0, 1$, element $\mathbf{c}_k$ and observation number τ that minimizes Bayes risk:

$$r^*_{j,\min} = \min_{\hat{X}_{j,t+\tau},k,\tau}$$
$$\left(r_j(\hat{X}_{j,t}, \tau = 0), r_k(\hat{X}_{k,t+1}(Y_{k,t+1}), \tau = 1)\right).$$

For the process estimation of a detected object $\mathbf{c}_j$, the Bayes risk of not taking any more measurement is the same as Equation (11). The (expected) risk of taking one more measurement associated with some element $\mathbf{c}_k$:

$$r_k(\hat{\mathbf{x}}_{k,t+1}(\mathbf{y}_{k,t+1}), \tau = 1) =$$
$$\int\int C(\hat{\mathbf{x}}_{k,t+1}(\mathbf{y}_{k,t+1}), \mathbf{x}_{k,t+1}) p(\mathbf{x}_{k,t+1}|\mathbf{y}_{k,1:t+1}) \text{d}\mathbf{x}_{k,t+1}$$
$$p(\mathbf{y}_{k,t+1}|\mathbf{y}_{k,1:t}) \text{d}\mathbf{y}_{k,t+1} + \kappa c^k_{j,\text{obs}}.$$

If under a quadratic cost assignment, the expected Bayes risk is given by

$$r_k(\hat{\mathbf{x}}_{k,t+1}(\mathbf{y}_{k,t+1}), \tau = 1) = c_e^{1,k} \text{Tr}\left[\hat{\mathbf{P}}_{k,t+1}\right] + \kappa c^k_{j,\text{obs}},$$

where $c_e^{1,k} > 0$ is the estimation cost with 1 observation associated with element k. If under UCA and assuming a 1 dimensional state, the expected Bayes risk is given by

$$r^k_j(\hat{\mathbf{x}}_{j,t+1}(\mathbf{y}_{j,t+1}), \tau = 1)$$
$$= c_e^{1,k}\left(1 - \text{Erf}\left(\frac{\epsilon}{2\sqrt{2\hat{\mathbf{P}}_{j,t+1}}}\right)\right) + \kappa c^k_{j,\text{obs}}.$$

The Bayesian sequential estimation method finds a particular combination of element $\mathbf{c}_k$ and observation number τ that yields the decision with minimum Bayes risk $r^*_{j,\min}$ for each given observation $\mathbf{y}_{j,t+1}$.

$$r^*_{j,\min} = \min_{k,\tau}\left(r_j(\hat{\mathbf{x}}_{j,t}, \tau = 0), r_k(\hat{\mathbf{x}}_{k,t+1}, \tau = 1)\right).$$

V. Risk-based Sensor Management for Integrated Detection and Estimation

A. Problem Statement

In this section, we develop an integrated sensor management scheme based on Bayesian sequential detection and estimation introduced in Sections II and III and their extension to multiple-element case in Section IV.

B. Detection and Estimation Sets

Let $Q_D(t) \subseteq \Omega$ be the set of cells for which no detection decision has been made up to time t and that are expected to be within the sensor's coverage area at the next time step $t+1$. Let $Q_T(t)$ be the set of detected objects that still need further measurements before an estimation decision can be made and that will be within the sensor's coverage area at the next time step $t+1$. Let $Q(t) = Q_D(t) \cup Q_T(t)$. Let $E(t)$ be the set of all cells in which it has been decided that no objects exist up to time t. Let $T(t)$ be the set of all processes that have the minimum Bayes risk which no further measurements are required.

C. Decision List

At some time t, a sensor makes one of two types of measurements of an element $\mathbf{c}_j$: (1) a detection measurement or (2) an estimation measurement. Based on the decisions made, an element $\mathbf{c}_j \in Q(t)$ can transition between the above mentioned sets as shown in Figure 2. Note that each arrow is associated with a Bayes risk and the optimal decision is made with overall minimum risk. The dashed arrows represent possible transitions of an already decided element and are not included in the risk computations in this paper.

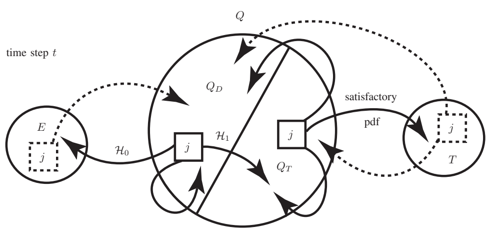

Fig. 2. Transition of an element $\mathbf{c}_j \in Q(t)$ at time t.

D. Observation Decision Costs

Define the cost of taking an observation at element $\mathbf{c}_j$ as

$$\mathbf{c}_{j,\mathrm{obs}} = E[I_{j^*}] - E[I_j], \quad (13)$$

where $E[I_j]$ is the expected information gain when measuring $\mathbf{c}_j$ and $\mathbf{c}_j^*$ is the element with the highest value of expected information gain. Hence, $\mathbf{c}_{j,\mathrm{obs}}$ models the information loss in making a suboptimal sensing allocation. We use the Rényi information divergence in computing the gain in information.

a) Rényi Information Divergence for Discrete Random Variables: For detection, the divergence is computed between two binary discrete random variables: the expected posterior p.m.f. $\{\hat{p}_{j,t+1}, 1-\hat{p}_{j,t+1}\}$ and the predicted probability mass function p.m.f. $\{\bar{p}_{j,t}, 1-\bar{p}_{j,t}\}$ [12]:

$$I_{j,\alpha}\left(\{\hat{p}_{j,t+1}, 1-\hat{p}_{j,t+1}\} \,|\, \{\bar{p}_{j,t}, 1-\bar{p}_{j,t}\}\right) = \frac{1}{\alpha-1}\log_2\left(\frac{\hat{p}_{j,t+1}^{\alpha}}{\bar{p}_{j,t+1}^{\alpha-1}} + \frac{(1-\hat{p}_{j,t+1})^{\alpha}}{(1-\bar{p}_{j,t+1})^{\alpha-1}}\right).$$

Since we need the Rényi information divergence to be a proper metric, we have to use $\alpha = 0.5$.

If we let $I_{j,\alpha;Y_{j,t+1}=1}$ and $I_{j,\alpha,Y_{j,t+1}=0}$ denote the Rényi information gain for the two possible types of sensor outputs at time $t+1$, the expected Rényi information gain is then given by

$$E_{Y_{j,t+1}|Y_{j,1:t}} I_{j,\alpha;Y_{j,t+1}}(\hat{p}_{j,t+1}\|\bar{p}_{j,t+1}) = \sum_{i=0}^{1} \mathrm{Prob}(Y_{j,t+1}=i|Y_{j,1:t}) I_{j,\alpha;Y_{j,t+1}=i}.$$

b) Rényi Information Divergence for Continuous Random Variables: For continuous random variables (e.g., for estimation), the Rényi information divergence at time t is computed between two continuous random variables: (a) the expected posterior probability density function (p.d.f.) $p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t+1})$ after another (unknown) measurement $\mathbf{y}_{j,t+1}$ is made, and (b) the predicted p.d.f. $p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t})$ given the measurements up to $\mathbf{y}_{j,t}$ [12], [13]

$$I_{j,\alpha}\left(p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t+1})|p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t})\right) = \frac{1}{\alpha-1}\log_2 \int p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t})\left(\frac{p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t+1})}{p(\mathbf{x}_{j,t+1}|\mathbf{y}_{j,1:t})}\right)^{\alpha} \mathrm{d}\mathbf{x}_{j,t+1}.$$

For linear Guassian models combined with a Kalman filter, we have [13]:

$$E_{y_{j,t+1}|y_{j,1:t}} I_{j,\alpha}(p_j(\mathbf{x}_{t+1}|\mathbf{y}_{1:t+1})\|p_j(\mathbf{x}_{t+1}|\mathbf{y}_{1:t})) = \frac{1}{2(1-\alpha)}\log\left(\frac{|\alpha\mathbf{R}^{-1}\mathbf{H}\bar{\mathbf{P}}_{t+1}\mathbf{H}^{\mathrm{T}}+\mathbf{I}|}{|\mathbf{R}^{-1}\mathbf{H}\bar{\mathbf{P}}_{t+1}\mathbf{H}^{\mathrm{T}}+\mathbf{I}|^{\alpha}}\right) + \frac{1}{2}\mathrm{Tr}\left[\mathbf{I} - \left(\alpha\mathbf{R}^{-1}\mathbf{H}\bar{\mathbf{P}}_{t+1}\mathbf{H}^{\mathrm{T}}+\mathbf{I}\right)^{-1}\right]$$

E. Solution Approach

Figure 3 summaries the solution approach as a general flow chart. At time step t, the sensor takes an observation at the current element $\mathbf{c}_j \in Q(t-1)$. Based on this real-time observation and the prior probability, the updated (posterior) probability and the predicted probability are obtained via a recursive implementation. The predicted probability is treated as the prior probability at the next time step $t+1$. We then compute the corresponding Bayes risk, where the updated probability is used to compute the Bayes risk $r_j(\tau=0)$ of making a direct detection/estimation decision without taking any further observations, and the predicted probability is used to compute the Bayes risk $r_j(\tau=1)$ associated with taking one more observation. Bayesian sequential decision-making is then employed as follows. If the minimum Bayes risk $r^*_{j,\min}$ is giving by taking future observations, then the sensor will take an observation at some element $\mathbf{c}_k \in Q(t)$ (including the possibility of choosing $\mathbf{c}_j$) at the next time step $t+1$. Otherwise ($\tau=0$), the sensor makes a detection/estimation decision at $\mathbf{c}_j$, and moves to some $\mathbf{c}_k \in Q(t)\backslash\{\mathbf{c}_j\}$ that minimizes the Bayes risk and takes an observation at that element at the next time step $t+1$. This process is repeated until a detection/tracking decision can be made at every element in $Q(t)$.

VI. Simulation Results

Assume that we have $N=10$ cells initially, among which there are 7 processes (Cell 1-7) to be detected and estimated. The initial predicted probability $\bar{p}_{j,t=0}$ for $j=3$ is set to be 0.1 and that for all the other cells is 0.5. The sensor detection probability β follows a Gaussian distribution with mean 0.6 and variance 0.1. The states are assumed to be time-invariant Gaussian processes with zero mean and positive definite covariance 0.1. The processes parameters are: $\mathbf{F}=1, \mathbf{H}=1, \mathbf{R}=1, \mathbf{Q}=0.1$. UCA is assumed and ϵ is set to be 0.1. The decision costs for detection and estimation are $c_d^0=1, c_d^1=0.3, c_e^0=$

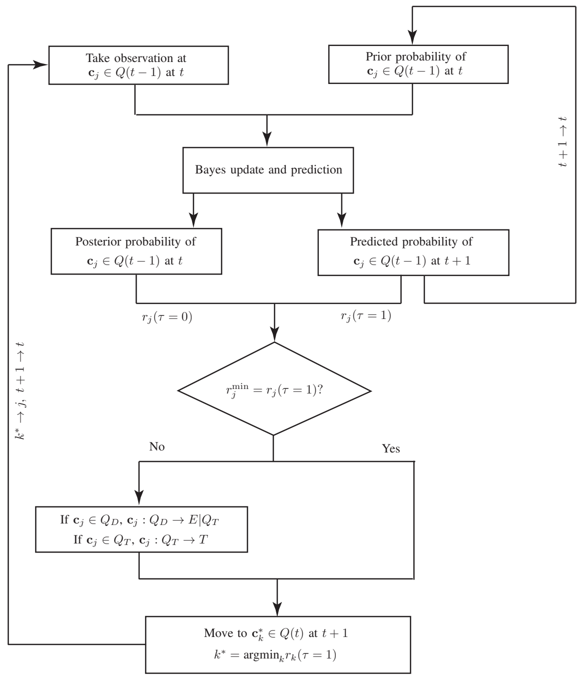

Fig. 3. Decision flowchart.

$1, c_e^1 = 0.16$. The information gain scaling parameter κ is chosen to be 0.06.

All the objects have been detected and have their processes estimated except that there is a missed detection at Cell 4. Figures 4(a), 4(b) and 4(c) show the detection and estimation results as examples. Figure 4(d) enlarges the estimation performance of Cell 1 during time period 1-300. Figure 5 shows the assigned observing cell at each time step according to the proposed scheme and the dots represent the detection stopping time for each object. Since there is a missed detection at Cell 4, no estimation is performed after the detection decision $X_4 = 0$ is made at time step 137.

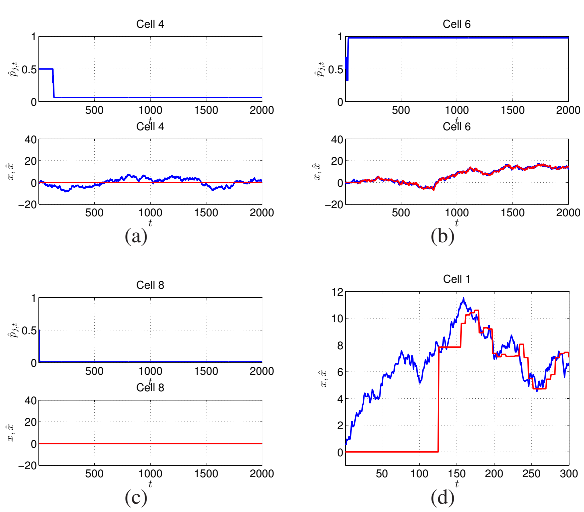

Fig. 4. Updated probability $\hat{p}_{j,t}$, and actual state $\mathbf{x}$ (blue) and state estimate $\hat{\mathbf{x}}$ (red) for (a) Cell 4, (b) Cell 6, and (c) Cell 8, and (d) Enlarged estimation performance for Cell 1 during time period 1-300.

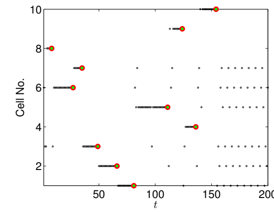

Fig. 5. Observed cell at each time step for a 10-cell detection and estimation problem (green dots: detection stopping time for each object).

VII. Conclusion and Future Work

In this paper, we develop an integrated sequential detection and estimation approach for optimal sensor allocation schemes. Rényi information gain measures is also introduced to model the observation cost. Future work will extend the above framework to nonlinear system estimation. Object mobility will be taken into account by using a probabilistic (Markov chain) model with non-identity probability transitions.

References

[1] P. B. Sujit and R. Beard, "Distributed Sequential Auctions for Multiple UAV Task Allocation," *Proceedings of 2007 American Control conference*, pp. 3955–3960, July 2007.

[2] R. Mahler, "Objective Functions for Bayeisan Control-Theoretic Sensor Management, I: Multitarget First-Moment Approximation," *Proceedings of IEEE Aerospace Conference*, 2003.

[3] M. Flint, M. Polycarpou, and E. Fernández-Gaucherand, "Cooperative Control for Multiple Autonomous UAV's Searching for Targets," *Proceedings of the 41st IEEE Conference on Decision and Control*, December 2002.

[4] A. Sinha, T. Kirubarajan, and Y. Bar-Shalom, "Autonomous Ground Target Tracking by Multiple Cooperative UAVs," *IEEE Aerospace Conference*, 2005.

[5] R. Vidal, O. Shakernia, H. J. Kim, D. H. Shim, and S. Sastry, "Probabilistic Pursuit-Evasion Games: Theory, Implementation and Experimental Evaluation," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, 2002.

[6] J. R. Spletzer and C. J. Taylor, "Dynamic Sensor Planning and Control for Optimally Tracking Targets," *The International Journal of Robotics Research*, pp. 7–20, January 2003.

[7] T. Furukawa, F. Bourgault, B. Lavis, and H. F. Durrant-Whyte, "Recursive Bayesian Search-and-Tracking using Coordinated UAVs for Lost Targets," *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, May 2006.

[8] T. Furukawa, H. F. Durrant-Whyte, and B. Lavis, "The Element-based Method - Theory and its Application to Bayesian Search and Tracking," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2807–2812, October 2007.

[9] J. Tisdale, A. Ryan, Z. Kim, D. Törnqvist, , and J. K. Hedrick, "A multiple UAV system for vision-based search and localization," *American Control Conference*, June 2008.

[10] A. Wald and J. Wolfowitz, "Bayes Solutions of Sequential Decision Problems," *The Annals of Mathematical Statistics*, vol. 21, no. 1, pp. 82–99, 1950.

[11] M. Ghosh, N. Mukhopadhyay, and P. K. Sen, *Sequential Estimation*, vol. 117 of *Wiley Series in Probability and Statistics - Applied Probability and Statistics Section Series*. Wiley-interscience, 2 ed., 1997.

[12] A. Rényi, "On Measures of Information and Entropy," *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability*, pp. 547–561, 1961.

[13] T. Hanselmann, M. Morelande, B. Moran, and P. Sarunic, "Sensor Scheduling for Multiple Target Tracking and Detection using Passive Measurements," *11th International Conference on Information Fusion*, pp. 1–8, 2008.

[14] A. Gelb, ed., *Applied Optimal Estimation*. Cambridge Mass: MIT press, 1974.