Particle filter approximations for general open loop and open loop feedback sensor management

Edson Hiroshi Aoki[1], Arunabha Bagchi[1], Pranab Mandal[1], and Yvo Boers[2]

[1]Department of Applied Mathematics, University of Twente, Enschede, The Netherlands.
[2]Thales Nederland B. V., Hengelo, The Netherlands.
[1]{e.h.aoki, a.bagchi, p.k.mandal}@ewi.utwente.nl
[2]yvo.boers@nl.thalesgroup.com

Memorandum 1952 (September 2011). ISSN 1874-4850. Available from: http://www.math.utwente.nl/publications
Department of Applied Mathematics, University of Twente, Enschede, The Netherlands

***Abstract*—Sensor management is a stochastic control problem where the control mechanism is directed at the generation of observations. Typically, sensor management attempts to optimize a certain statistic derived from the posterior distribution of the state, such as covariance or entropy. However, these statistics often depend on future measurements which are not available at the moment the control decision is taken, making it necessary to consider their expectation over the entire measurement space.**

Though the idea of computing such expectations using a particle filter is not new, so far it has been applied only to specific sensor management problems and criterions. In this memorandum, for a considerably broad class of problems, we explicitly show how particle filters can be used to approximate general sensor management criterions in the open loop and open loop feedback cases. As examples, we apply these approximations to selected sensor management criterions.

As an additional contribution of this memorandum, we show that every performance metric can be used to define a corresponding estimate and a corresponding task-driven sensor management criterion, and both of them can be approximated using particle filters. This is used to propose an approximate sensor management scheme based on the OSPA metric for multi-target tracking, which is included among our examples.

Keywords: Sensor management, entropy, Kullback-Leibler divergence, Rényi divergence, OSPA metric, particle filter.

I. Introduction

Sensor management is a control problem associated with partially observed systems, where the control action aims to influence the generation of observations (direct feedthrough) and not the state of the system, generally with the goal of obtaining the best possible estimation quality of the state given limited sensing resources.

Typically the sensor management problem is formulated in terms of minimization of a risk function related to the error between the true state and the estimated state; this is the so-called "task-driven" sensor management. An alternative is instead attemping to improve (in some sense) the "information content" of the distribution. This "information-driven" sensor management consists of choosing the control decision that maximizes some notion of information gain (or, alternatively, minimizes some notion of uncertainty).

As it is well-known, when the goal function (i.e. the risk function to be minimized or the reward function to be maximized) does not depend on the actual measurement, there is no need to take the expectation over the measurement space. This is the case, for instance, when the system is linear-Gaussian and the criterion in question is the covariance of the MMSE estimate. For general systems and criterions, however, the goal function would depend on future measurements, which are not known at the time the control decision is taken.

One approach, illustrated by Williams, Fisher and Wilsky [1], is to use a linearized Gaussian approximation, which results in covariance being assumed as measurement-independent. This implies that a goal function based on conditional entropy would also be measurement-independent due to the relationship betweeen covariance and entropy in the Gaussian case.

Another approach, shown by Zhao, Shin and Reich [2], consists on applying an heuristic that uses an estimated measurement in the computation of the goal function. Combined with a grid-based discretization method, this approach does not require linearization, neither imposes obvious restrictions on the criterion to be chosen.

To the best of our knowledge, Doucet et al. [3] were the first to propose the use of particle filters on the evaluation of sensor management goals, which in this case was the Kullback-Leibler divergence, although others were suggested. Simulated measurements were used to address the continuity of the measurement space. This approach fundamentally differs from the previously described ones in the sense that it considers the expectation of the goal function over the measurement space, and it is thus optimal save for the inherent errors in the particle approximation.

Subsequently, Kreucher, Kastella and Hero [4] proposed using particle filters to implement the Rényi divergence criterion, with the implementation explicitly described for discrete measurement spaces. Another work of Kreucher, Kastella and Hero [5] considered a sensor management criterion based on maximization of the marginalized posterior density, with the purpose of being compared with the Rényi divergence in terms of performance.

This work is organized as follows. Section II describes the sensor management problem for which the particle approximations proposed in this memorandum are valid. Section III gives a short introduction to sensor management criterions. Section IV describes how a general performance metric can be used to define a corresponding estimate and a sensor management criterion; this is used to define a criterion based on the OSPA metric for multi-target tracking. Section V demonstrates how particle filters can be used to approximate generic goals for both continuous and discrete measurement spaces, for a reasonably broad class of sensor management problems. We consider only optimal solutions (i.e. with no errors other than those resulting from particle approximations) for long- or short-term, open loop (feedback) sensor management. Finally, Section VI, as examples, shows the application of this method to the previously discussed criterions.

II. Mathematical formulation of the sensor management problem

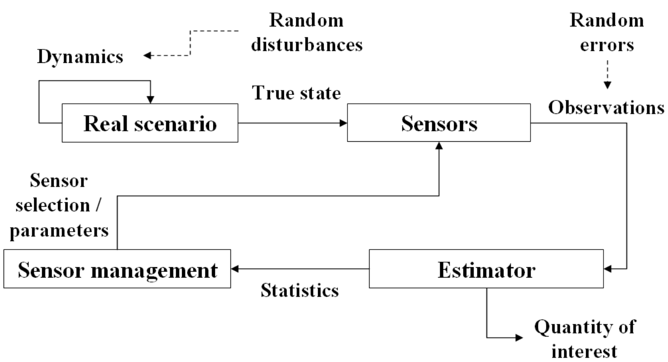

Figure 1. Sensor management as a stochastic control problem

A representation of sensor management as a stochastic control problem is shown on Figure 1. We consider a (static or dynamic) scenario described by a state X, observed by a measurement device composed of one or more sensors, with the sensor observations Y corrupted by random errors. These measurements are used as input to an estimator, which obtains an estimate $\hat{\Psi}$ of a quantity of interest Ψ, in general a function of X.

A feedback occurs through a sensor management device, which uses $\hat{\Psi}$ or other statistics computed by the estimator to select a control decision (or "sensing action") U that affects the generation of subsequent observations. Typically, U is chosen by minimization (or maximization) of the expectation of a risk (or reward) function $\gamma(X, Y, U)$.

As we can see from Figure 1, the difference between sensor management and the standard control problem is that, in the first, the control decision aims to affect the generation of observations (i.e. direct feedthrough input), while in the second, it aims to affect the true state.

We will now assume certain properties of the class of problems we are considering. At an arbitraty time k, let X_k be the scenario state, Y_k be the set of measurements, U_k be the control decision (prior to generating Y_k), W_k be a random disturbance in process and V_k be a random measurement error. Then, the assumed properties are

1) $X_{k+1} = f(k, X_k, W_k)$,
2) $Y_k = g(k, X_k, V_k, U_k)$,
3) The sequences $\{W_k\}$ and $\{V_k\}$ contain independent elements, are mutually independent, and are independent from the initial state X_0.

We consider this problem a subset of the Partially Observed Controlled Markov Process (POCMP) problem. We will, however, avoid describing the problem as a Partially Observed Markov Decision Process (POMDP) as done by other authors, since the term rigorously only applies when the decision space is discrete [6], and it does not distinguish the sensor management problem from the regular control problem.

At time $k+1$, the information state (the available information prior to selecting control decision U_{k+1}) is given by $Z_k = \{Y_1, \ldots, Y_k, U_1, \ldots, U_k\}$. Hence we have

$$\begin{aligned} U_{k+1} &= \eta_k(Z_k) \\ &= \eta_k(Y_1 \ldots Y_k, U_1 \ldots U_k) \end{aligned} \tag{1}$$

where η_k is some function.

For this problem, it is possible to show that a sufficient statistic is the density function given by $p(x_k|z_k)$ (usually referred as "filtering density"), and that we can infer the following additional properties:

$$p(x_{k+1}|x_k, z_k, u_{k+1}) = p(x_{k+1}|x_k), \tag{2}$$

$$p(x_k|z_k, u_{k+1}) = p(x_k|z_k) \text{ and} \tag{3}$$

$$p(y_k|x_k, u_k, z_{k-1}) = p(y_k|x_k, u_k). \tag{4}$$

Finally, for an initial time k_0 and a given time horizon H, the goal of sensor management is to select a control law η that minimizes (or maximizes)

$$\tilde{J}(k_0, z_{k_0}, H, \eta) \triangleq E\left[\sum_{i=k_0+1}^{k_0+H} \gamma(i, X_i, Y_i, U_i) \middle| \eta, z_{k_0}\right] \tag{5}$$

where $\gamma(i, X_i, Y_i, U_i)$ is a goal function associated with time i, and $\eta = \{\eta_{k_0+1}, \ldots, \eta_{k_0+H}\}$ is a feasible control law for the system. We will from now on refer to $\tilde{J}$ as "expected goal" to avoid confusion with the goal function γ.

A. Open and closed loop sensor management

Let us consider an initial time k_0, an horizon $H \geq 1$, and a control law η. In this situation, we can use one of three possible approaches for sensor management: closed loop control (CLC), open loop control (OLC), and open loop feedback control (OLFC).

1) **Closed loop control**: minimization of (5) is done over all η such that $U_{k+1} = \eta_k(Z_k)$ for $k_0 \leq k < k_0 + H$; therefore, it results in the so-called globally optimal control law.
2) **Open loop control:** minimization of (5) considers only η such that $U_{k+1} = \eta_k(z_{k_0})$, i.e. the control law η does not consider any information that become available after the initial time k_0. OLC is a much simpler problem than CLC because since z_{k_0} is known, U_k is not random but deterministic. This means that instead of attempting to find a control law η, we may just search the decision space for the optimal values of U_k.

3) **Open loop feedback control:** basically a compromise between the other two. At time k_0, we find the open loop control law $\eta^{(k_0)}$ and make $U_{k_0+1} = \eta^{(k_0)}_{k_0}(z_{k_0})$. For time k_0+2, instead of making $U_{k_0+2} = \eta^{(k_0)}_{k_0+1}(z_{k_0})$ as in OLC, we instead search for a new optimal open loop control law $\eta^{(k_0+1)}$ that uses the new available information z_{k_0+1}, i.e. we have $U_{k_0+2} = \eta^{(k_0+1)}_{k_0+1}(z_{k_0+1})$, and the process is repeated.

Although only CLC guarantees the globally optimal control law, its practical use is restricted. In principle, dynamic programming allows the CLC problem to be represented as a system of equations, but in practice, a solution can only be obtained for very specific classes of problem (see [6]). For these reasons, the CLC problem is not further considered in this memorandum; interesting approximations, however, can be found e.g. in [7], [8].

As for the other two approaches, one can prove that the performance of the optimal OLFC is no worse than the optimal OLC [9], but the difference in performance between OLFC and CLC can be arbitrarily large. Nevertheless, OLFC is popular approach because it is tractable as long as the OLC problem is tractable. Some nice examples of application of OLFC can be found in [1], [10].

III. Sensor management criterions

Different ways of selecting the goal function of the sensor management problem have been proposed in the literature. We distinguish two: task-driven and information-driven sensor management. From now on, we are going to use the notation $\mu_{A|B}$ to refer to the distribution of a random variable A conditioned on a random variable or control decision B.

A. Task-driven sensor management

For the sensor management problem described in Section II, let Ψ (a function of state X) be a quantity of interest to the operator of the system; hence the estimate $\hat{\Psi}$ is the external output of the estimator shown in Figure 1. Let now $\epsilon(\Psi, \hat{\Psi})$ be a performance metric for our estimator, corresponding to some measure of error between the true quantity Ψ and its estimate $\hat{\Psi}$.

In what we call task-driven sensor management, we directly attempt to optimize the chosen performance metric, i.e. the goal function is given by

$$\gamma(k, X_k, Y_k, U_k) = \epsilon(\Psi_k, \hat{\Psi}_k). \tag{6}$$

Naturally, we can define task-driven criterions that do not precisely have form (6), such as when we use a performance metric that is not function of the true state (for instance, the maximum posterior probability suggested in [5]).

B. Information-driven sensor management

In information-driven sensor management, we attempt to maximize the "information content" of the posterior, i.e. its capacity of yielding (in some sense) useful information to the operator, rather than attempting to maximize the quality of the estimate. In this case, for the sensor management problem described in Section II, the goal function may be given by

$$\gamma(k, Y_k, U_k) = f(\mu_{X_k|Y_k,U_k,Z_{k-1}}) \tag{7}$$

where f is some measure of information content of the posterior distribution $\mu_{X_k|Y_k,U_k,Z_{k-1}}$. Typically γ, although a function of the distribution of X, is not a function of X itself. Alternatively, instead of just looking at the posterior distribution, we may attempt to maximize some notion of information gain between prior and posterior distributions. In this case, we have

$$\gamma(k, Y_k, U_k) = f(\mu_{X_k|Y_k,U_k,Z_{k-1}}, \mu_{X_k|Z_{k-1}}) \tag{8}$$

where f is some measure of information gain obtained by moving from prior distribution $\mu_{X_k|Z_{k-1}}$ to posterior distribution $\mu_{X_k|Y_k,U_k,Z_{k-1}}$. Also, observe that $\mu_{X_k|Z_{k-1}} = \mu_{X_k|U_k,Z_{k-1}}$ due to properties (2) and (3).

1) The Kullback-Leibler (KL) divergence: The relative Shannon entropy, or Kullback-Leibler (KL) divergence is a measure of difference between two distributions. Consider a pair of distributions μ and ν which respectively admit densities p and q with respect to a dominating σ-finite measure ρ. The KL divergence from μ to ν (or alternatively, from p to q) is given by

$$D_{\text{KL}}(p\|q) \triangleq \int p(x) \log \frac{p(x)}{q(x)} \rho(dx) \tag{9}$$

where we apply the conventions $\log \frac{p(x)}{q(x)} = 0$ for $p(x) = 0$ and $q(x) = 0$, and $a/0 = \infty$ for $a > 0$. Note that the measure is asymmetric in the sense that $D_{\text{KL}}(p\|q) \neq D_{\text{KL}}(q\|p)$.

In information-driven sensor management, divergences are used as criterions of form (8), i.e. they are considered to represent a notion of information gain between the prior and the posterior distribution. Since the KL divergence is assymetric w.r.t. its arguments, one may ask which order of the arguments shall be used. Following the discussion in [11], if we consider that minimization of Shannon entropy is desirable, the "correct" order corresponds to $D_{\text{KL}}(\mu_{X_k|Y_k,U_k,Z_{k-1}} \| \mu_{X_k|Z_{k-1}})$.

2) The Rényi (α-) divergence: The Rényi divergence or α-divergence is a generalisation of the KL divergence. The α-divergence from μ to ν (or alternatively, from p to q) is given by

$$D_\alpha(p\|q) \triangleq \frac{1}{\alpha - 1} \log \int p^\alpha(x) q^{1-\alpha}(x) \rho(dx). \tag{10}$$

where we apply the conventions $p^\alpha(x)q^{1-\alpha}(x) = 0$ for $p(x) = q(x) = 0$, and $a/0 = \infty$ for $a > 0$. D_0 and D_1 are defined using the limits from right and left respectively, which makes D_1 the same as the D_{KL}. $D_{0.5}$ has the special property that is a true metric, in the sense that it is symmetric and obeys the triangle inequality.

IV. Relation between performance measurement, estimation and task-driven sensor management

In this section, we will show that performance measurement, estimation and task-driven sensor management are intrinsically related problems, i.e. that the choice of a performance metric leads to some corresponding optimal estimate, and those lead to some corresponding optimal (task-driven) sensor management criterion.

Let X be, without loss of generality, both the state and the quantity of interest, $\hat{X}$ a general estimate of X, and $\epsilon(X,\hat{X})$ a performance metric. Then, given the posterior distribution $\mu_{X_k|Y_k,U_k,Z_{k-1}}$, we can define an "optimal" estimate based on ϵ and some function f by

$$\hat{X}_k \triangleq \arg\min_{\hat{X}_k^*} E_{\mu_{X_k|Y_k,U_k,Z_{k-1}}}\left[f\left(\epsilon(X_k,\hat{X}_k^*)\right)\right] \quad (11)$$

where we can naturally consider any appropriate distribution of X instead of $\mu_{X_k|Y_k,U_k,Z_{k-1}}$.

Now, we can use ϵ, f and $\hat{X}_k$ to define a task-driven sensor management goal of form (6), which makes the optimal control law (according to (5)) to be given by

$$\eta = \arg\min_{\eta^*} E_{\mu_{X_{k_0+1,\dots,k_0+H},Y_{k_0+1,\dots,k_0+H}|\eta^*,z_{k_0}}}\left[\sum_{i=k_0+1}^{k_0+H} f\left(\epsilon(X_i,\hat{X}_i)\right)\right]. \quad (12)$$

Let us assume that $H=1$, i.e. we have short-term sensor management. In this case, since $U_{k_0+1}=\eta_{k_0}(z_{k_0})$, the control law always corresponds to open loop, and following our discussion in Section II-A, we may perform the minimization in the space of sensing actions U rather than in the space of control laws η. Hence, for $k=k_0+1$, we may rewrite (12) as

$$U_k = \arg\min_{U_k^*} E_{\mu_{X_k,Y_k|U_k^*,z_{k-1}}}\left[f\left(\epsilon(X_k,\hat{X}_k)\right)\right] \quad (13)$$

A. Example: RMS errors

As an example, let us take as performance metric the RMS errors of X, i.e.

$$\epsilon(X,\hat{X}) := \sqrt{\left(X-\hat{X}\right)'\left(X-\hat{X}\right)}. \quad (14)$$

If we choose $f(a) := a^2$, the corresponding estimate is then given by

$$\hat{X}_k = \arg\min_{\hat{X}_k^*} E_{\mu_{X_k|Y_k,U_k,Z_{k-1}}}\left[\left(X_k-\hat{X}_k^*\right)'\left(X_k-\hat{X}_k^*\right)\right] \quad (15)$$

which is the familiar MMSE estimate, and hence

$$\hat{X}_k = E_{\mu_{X_k|Y_k,U_k,Z_{k-1}}}\left[X_k\right]. \quad (16)$$

Using short-term, task-driven sensor management, the optimal control decision is obtained by

$$U_k = \arg\min_{U_k^*} E_{\mu_{X_k,Y_k|U_k^*,z_{k-1}}}\left[\left(X_k-\hat{X}_k\right)'\left(X_k-\hat{X}_k\right)\right]. \quad (17)$$

Observe that, if P_k is the covariance of estimate $\hat{X}_k$, this criterion is actually equivalent to minimizing the trace of P_k, since

$$\begin{aligned} &E_{\mu_{X_k,Y_k|U_k^*,z_{k-1}}}\left[\left(X_k-\hat{X}_k\right)'\left(X_k-\hat{X}_k\right)\right] \\ &= E_{\mu_{Y_k|U_k^*,z_{k-1}}}\Bigg[E_{\mu_{X_k|Y_k,U_k^*,z_{k-1}}}\left[\left(X_k-\hat{X}_k\right)'\left(X_k-\hat{X}_k\right)\right]\Bigg] \\ &= E_{\mu_{Y_k|U_k^*,z_{k-1}}}\Bigg[\text{tr}\left(E_{\mu_{X_k|Y_k,U_k^*,z_{k-1}}}\left[\left(X_k-\hat{X}_k\right)\left(X_k-\hat{X}_k\right)'\right]\right)\Bigg] \\ &= E_{\mu_{Y_k|U_k^*,z_{k-1}}}\left[\text{tr}\left(P_k\right)\right]. \end{aligned} \quad (18)$$

B. Example: OSPA metric

The Optimal Subpattern Assignment Metric (OSPA) [12] is a metric designed for the multi-target tracking problem, and has some nice properties including meaningful physical interpretation when the quantities compared have different cardinalities, and generation of the standard topology used in point process theory.

Consider the concatenated state $X=\left[X'^{(1)}\dots X'^{(T)}\right]'$, where $X^{(1)}\dots X^{(T)}$ denote respectively the individual states of targets $1\dots T$, and the corresponding estimate $\hat{X}=\left[\hat{X}'^{(1)}\dots\hat{X}'^{(\hat{T})}\right]'$, where the estimated number of targets $\hat{T}$ may be different from the actual number of targets T.

The OSPA metric is defined as follows. Let $1\le p<\infty$ be an order parameter that penalizes estimated objects far away from objects of the ground truth, and $c>0$ be a cut-off parameter that penalizes cardinality errors. Let also Π_k be the set of all permutations on $\{1\dots k\}$, and $d^{(c)}(a,b)$ be defined by

$$d^{(c)}(a,b) = \min\left(d(a,b),c\right). \quad (19)$$

where $d(a,b)$ is the Euclidean distance between a and b.

Then the OSPA metric parameterized by p and c is defined by

$$\epsilon_p^{(c)}(X,\hat{X}) \triangleq \left(\frac{1}{\hat{T}}\left(\min_{\pi\in\Pi_{\hat{T}}}\sum_{j=1}^{T} d^{(c)}\left(X^{(j)},\hat{X}^{(\pi(j))}\right)^p + c^p(\hat{T}-T)\right)\right)^{\frac{1}{p}} \quad (20)$$

for $T\le\hat{T}$, and $\epsilon_p^{(c)}(X,\hat{X}) \triangleq \epsilon_p^{(c)}(\hat{X},X)$ otherwise. By choosing $f(a) := a^p$, we may derive the corresponding estimate based on the OSPA metric

$$\hat{X}_k = \arg\min_{\hat{X}_k^*} E_{\mu_{X_k|Y_k,U_k,Z_{k-1}}}\left[\epsilon_p^{(c)}(X_k,\hat{X}_k^*)^p\right]. \quad (21)$$

This estimate was first derived by Guerriero et al. [13], which they called MMOSPA (Minimum Mean OSPA) estimate. In their paper, for a density defined in a set space (i.e. unordered density) and a class of densities defined in vector spaces (i.e. ordered densities) that are jointly equivalent to this unordered density, they verified that the MMSE estimate of one of the ordered densities is equal to the MMOSPA estimate with $p = 2$ and $c = \infty$, which is the same for the unordered and ordered densities. The equivalence to MMSE is important because it allows (21) to be more easily computed when the mean of a distribution can be easily computed.

Finally, we may derive the corresponding short-term, task-driven sensor management scheme, with the control decision obtained by

$$U_k = \arg\min_{U_k^*} E_{\mu_{X_k, Y_k | U_k^*, z_{k-1}}} \left[\epsilon_p^{(c)}(X_k, \hat{X}_k)^p \right]. \qquad (22)$$

V. Particle filter implementation of sensor management

In this section we discuss practical implementation of sensor management criterions for non-linear non-Gaussian systems, using particle filters. The reader is assumed to be familiar with the standard SIR particle filter.

We consider only open loop control (OLC), i.e. control law η is such that $U_{k+1} = \eta_k(z_{k_0})$ for all $k_0 \leq k < k_0 + H$. Naturally, all results obtained for OLC can be used for open loop feedback control (OLFC), as discussed in Section II-A. In the OLC case, the expected goal (5) is given by

$$\begin{aligned}
&\tilde{J}(k_0, z_{k_0}, H, \eta) \\
&= \int \cdots \int \left(\sum_{k=k_0+1}^{k_0+H} \gamma(k, x_k, y_k, u_k) \right) p(x_{k_0+1}, \ldots, x_{k_0+H} \\
&\quad , y_{k_0+1}, \ldots, y_{k_0+H} | z_{k_0}, u_{k_0+1}, \ldots, u_{k_0+H}) \\
&\quad \times dx_{k_0+1} \ldots dx_{k_0+H} dy_{k_0+1} \ldots dy_{k_0+H} \qquad (23)
\end{aligned}$$

where $u_k = \eta_{k-1}(z_{k_0})$. Now, considering the problem described in Section II, it is possible to see that, for k and l such that $k_0 < k \leq l$, we have the properties

$$p(x_k | x_{k_0}, \ldots, x_{k-1}, u_k, \ldots, u_l, z_{k-1}) = p(x_k | x_{k-1}), \quad (24)$$

$$p(x_{k_0} | u_{k_0+1}, \ldots, u_l, z_{k_0}) = p(x_{k_0} | z_{k_0}), \quad (25)$$

$$p(y_k | x_{k_0+1}, \ldots, x_k, u_k, \ldots, u_l, z_{k-1}) = p(y_k | x_k, u_k) \quad (26)$$

so after a few manipulations, the expected goal (23) becomes

$$\begin{aligned}
&\tilde{J}(k_0, z_{k_0}, H, \eta) \\
&= \int \cdots \int \left(\sum_{k=k_0+1}^{k_0+H} \gamma(k, x_k, y_k, u_k) \right) \\
&\quad \times \prod_{k=k_0+1}^{k_0+H} p(y_k | x_k, u_k) \prod_{k=k_0+1}^{k_0+H} p(x_k | x_{k-1}) \\
&\quad \times p(x_{k_0} | z_{k_0}) dx_{k_0} \ldots dx_{k_0+H} dy_{k_0+1} \ldots dy_{k_0+H}. \qquad (27)
\end{aligned}$$

If we are using a particle filter to estimate the target states, then the filtering density $p(x_{k_0} | z_{k_0})$ is approximated by a set of particles $\{x_{k_0}^{(i)}\}_{i=1}^N$, where N is the number of particles and we assume that the particles have identical weights. The expected goal (27) is therefore approximated as

$$\begin{aligned}
&\tilde{J}^N(k_0, z_{k_0}, H, \eta) \\
&= \frac{1}{N} \sum_{i=1}^{N} \int \cdots \int \left(\sum_{k=k_0+1}^{k_0+H} \gamma(k, x_k, y_k, u_k) \right) \\
&\quad \times \prod_{k=k_0+1}^{k_0+H} p(y_k | x_k, u_k) \prod_{k=k_0+2}^{k_0+H} p(x_k | x_{k-1}) \\
&\quad \times p(x_{k_0+1} | x_{k_0}^{(i)}) dx_{k_0+1} \ldots dx_{k_0+H} \\
&\quad \times dy_{k_0+1} \ldots dy_{k_0+H}. \qquad (28)
\end{aligned}$$

This approximation is by itself not particularly useful because we are still left with $2H$ integrals. What we do next depends on the type of problem. If the measurement space is continuous, or discrete with very large cardinality, we would need simulated measurements in order to approximate the expectations involving terms of form $p(y_k | x_k, u_k)$. In other cases (i.e. discrete measurement spaces with sufficiently small cardinality), we can analytically compute such expectations as long as we can compute the expectations involving the $p(x_k | x_{k-1})$ terms.

A. *Expectation approximation techniques*

We need to eliminate the integrals (or, in other words, approximate the expectations) involving the terms of form $p(x_k | x_{k-1})$ (and also $p(y_k | x_k, u_k)$ if simulated measurements are necessary) from the expected goal (28). An intuitive idea is to generate a new set of samples to eliminate each integral, i.e. for the expectation of a function $f(a)$ of some random variable A, taken over density $p(a)$, to make

$$\int f(a) p(a) da \approx \frac{1}{N_A} \sum_{i=1}^{N_A} f(a^{(i)}) \qquad (29)$$

where $\{a^{(i)}\}^{N_A}$ is a set of samples of A, sampled according to $p(a)$.

The problem of this approach is that the computational cost of the expected goal would be increased by a factor N_A for each integral to be eliminated. A simpler idea to keep only the initial set of particles $\{x_{k_0}^{(i)}\}_{i=1}^N$, and for $k_0 < k \leq k_0 + H$, repeatedly sample $x_k^{(i)} \sim p(x_k | x_{k-1}^{(i)})$ (and $y_k^{(i)} \sim p(y_k | x_k^{(i)}, u_k)$ as necessary) for each particle i. This is equivalent to using (29), but choosing $N_A = 1$.

While at first glance this approach may look too simplistic, the reader may notice that this sampling mechanism is exactly the same used in the importance sampling step of the particle filter. In fact, if the last importance sampling step was done using the Markov transition density $p(x_{k_0+1} | x_{k_0})$ as importance function, we may promptly use the resulting set of particles to eliminate the corresponding integral in (28).

We will now apply the described technique to approximate (28). If using simulated measurements is necessary, we start from (28) and repeately sample $x_k^{(i)} \sim p(x_k|x_{k-1}^{(i)})$ and $y_k^{(i)} \sim p(y_k|x_k^{(i)}, u_k)$, leading to the second approximation

$$\begin{aligned}&\tilde{J}^N(k_0, z_{k_0}, H, \eta)\\&= \frac{1}{N}\sum_{i=1}^{N}\sum_{k=k_0+1}^{k_0+H} \gamma(k, x_k^{(i)}, y_k^{(i)}, u_k).\end{aligned} \tag{30}$$

In other cases, we just need to sample $x_k^{(i)} \sim p(x_k|x_{k-1}^{(i)})$ since the expectation over the measurement space is not necessary, so the approximation is used is

$$\begin{aligned}&\tilde{J}^N(k_0, z_{k_0}, H, \eta)\\&= \frac{1}{N}\sum_{i=1}^{N}\int\cdots\int\sum_{k=k_0+1}^{k_0+H} \gamma(k, x_k^{(i)}, y_k, u_k)\\&\times\prod_{k=k_0+1}^{k_0+H} p(y_k|x_k^{(i)}, u_k)dy_{k_0+1}\dots dy_{k_0+H}.\end{aligned} \tag{31}$$

B. *Comments on practical implementation*

As discussed in Section II-A, for OLC the optimal $u_{k_0+1}, \dots, u_{k_0+H}$ are deterministic and can be obtained by searching the decision space, without the need to actually find the control law η. If the decision space is discrete this can be done by enumerating all possible decisions and evaluating the expected goal for each of them. In the continuous case we can use suitable minimization techniques such as gradient descent.

It is easy to see that the computational cost of searching the decision space can increase exponentially with the horizon H. Some heuristics to circunvent this problem can be found in [14] and [1].

We should also note that while the complexity of computing the expected goal for a single sensing action, for both (30) and (31), seems to be only $O(N)$, it can actually be more due to the cost resulting from the computation of $\gamma(k, x_k, y_k, u_k)$. In fact, for all examples in Section VI, the complexity of computing $\gamma(k, x_k, y_k, u_k)$ is at least $O(N)$, thus making the total cost at least $O(N^2)$. If this cost is too high, it can be reduced, as suggested in [3], by using only a subset (with cardinality $P < N$) of the original particles in the computation of the expected goal. This subset can be constructed by multinomial sampling.

Finally, we should remark that $p(x_k|x_{k-1})$ and $p(y_k|x_k, u_k)$ are not unique choices of densities that can be used to sample $\{x_k^{(i)}\}_{i=1}^N$ and $\{y_k^{(i)}\}_{i=1}^N$. In fact, Doucet et al. [3] mentioned that better approximations can be obtained by using different sampling functions.

VI. Examples

In this section, we will present some examples of particle approximations of sensor management goals. We will consider the previously discussed goal functions, simulated measurements, and short-term sensor management, i.e. we will use (30) with $H = 1$. After seeing these examples, the reader shall be able to derive the correspondings expressions for other cases (e.g. $H > 1$ and without simulated measurements) with relative ease.

In the short-term sensor management case, we can take (27) and make $k = k_0 + 1$, so the expected goal becomes

$$\tilde{J} = \int\int \gamma(k, x_k, y_k, u_k)p(y_k|x_k, u_k)p(x_k)dx_k dy_k \tag{32}$$

where we omit the conditioning on previous available information z_{k_0} for brevity. With $p(x_{k_0})$ approximated by the set of particles $\{x_{k_0}^{(i)}\}_{i=1}^N$, we sample $x_k^{(i)} \sim p(x_k|x_{k_0}^{(i)})$ and $y_k^{(i)} \sim p(y_k|x_k^{(i)}, u_k)$, and obtain the particle approximation

$$\tilde{J}^N = \frac{1}{N}\sum_{i=1}^{N}\gamma(k, x_k^{(i)}, y_k^{(i)}, u_k). \tag{33}$$

We will now derive expressions of the expected goal $\tilde{J}^N$ for different goal functions γ. We note that the approximation for the KL divergence has been provided by Doucet et al. [3], although we provide a slightly different derivation. The derivations of the other approximations are new, at least to the best of our knowledge.

A. *Rényi divergence*

For the α-divergence (10), there are two possible goal functions (given by the two possible order of arguments). The most commonly used order corresponds to

$$\begin{aligned}&\gamma(k, y_k, u_k)\\&= D_\alpha(p_{X_k|Y_k,U_k}\|p_{X_k})\\&= \frac{1}{\alpha-1}\log\int p(x_k|y_k, u_k)^\alpha p(x_k)^{1-\alpha}dx_k\\&= \frac{1}{\alpha-1}\log\int \frac{p(y_k|x_k, u_k)^\alpha}{p(y_k|u_k)^\alpha}p(x_k)dx_k.\end{aligned} \tag{34}$$

Note that the order is irrelevant when $\alpha = 0.5$, i.e. when the α-divergence is symmetric. Substituing (34) in particle approximation (33), we obtain

$$\begin{aligned}\tilde{J}^N &= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\alpha-1}\log\int \frac{p(y_k^{(i)}|x_k, u_k)^\alpha}{p(y_k^{(i)}|u_k)^\alpha}p(x_k)dx_k\\&= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\alpha-1}\log\left(\frac{1}{N}\sum_{j=1}^{N}\frac{p(y_k^{(i)}|x_k^{(j)}, u_k)^\alpha}{p(y_k^{(i)}|u_k)^\alpha}\right).\end{aligned} \tag{35}$$

Observe now that we can approximate $p(y_k^{(i)}|u_k)$ according to

$$\begin{aligned}p(y_k^{(i)}|u_k) &= \int p(y_k^{(i)}|x_k, u_k)p(x_k)dx_k\\&\approx \frac{1}{N}\sum_{l=1}^{N}p(y_k^{(i)}|x_k^{(l)}, u_k)\end{aligned} \tag{36}$$

and thus the expected goal becomes

$$\begin{aligned}\tilde{J}^N &= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\alpha-1}\log\sum_{j=1}^{N}\frac{p(y_k^{(i)}|x_k^{(j)},u_k)^\alpha}{\left(\sum_{l=1}^{N}p(y_k^{(i)}|x_k^{(l)},u_k)\right)^\alpha}\\ &= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{\alpha-1}\Bigg(\log\sum_{j=1}^{N}p(y_k^{(i)}|x_k^{(j)},u_k)^\alpha\\ &\quad -\alpha\log\sum_{l=1}^{N}p(y_k^{(i)}|x_k^{(l)},u_k)\Bigg). \end{aligned} \tag{37}$$

B. KL divergence

As we mentioned in Section III-B1, to ensure equivalence with the entropy criterion, there is a correct order of arguments to be chosen for the KL divergence, corresponding to

$$\begin{aligned}&\gamma(k,y_k,u_k)\\ &= D_{\mathrm{KL}}(p_{X_k|Y_k,U_k}\|p_{X_k})\\ &= \int p(x_k|y_k,u_k)\log\frac{p(x_k|y_k,u_k)}{p(x_k)}dx_k.\\ &= \int \frac{p(y_k|x_k,u_k)p(x_k)}{p(y_k|u_k)}\log\frac{p(x_k|y_k,u_k)}{p(x_k)}dx_k\\ &= \int \frac{p(y_k|x_k,u_k)p(x_k)}{p(y_k|u_k)}\log\frac{p(y_k|x_k,u_k)}{p(y_k|u_k)}dx_k. \end{aligned} \tag{38}$$

We could, as done for the Rényi divergence, substitute (38) in particle approximation (33). However, due to the presence of $p(y_k|u_k)$ in the denominator and the fact that the goal function does not depend on x_k, we can obtain a simpler expression by substituting (38) in the non-approximated expected goal (32):

$$\begin{aligned}\tilde{J} &= \int\int D_{\mathrm{KL}}(p_{X_k|Y_k,U_k}\|p_{X_k})p(y_k|x_k,u_k)p(x_k)dx_kdy_k\\ &= \int D_{\mathrm{KL}}(p_{X_k|Y_k,U_k}\|p_{X_k})\int p(y_k|x_k,u_k)p(x_k)dx_kdy_k\\ &= \int D_{\mathrm{KL}}(p_{X_k|Y_k,U_k}\|p_{X_k})p(y_k|u_k)dy_k\\ &= \int\int p(y_k|x_k,u_k)p(x_k)\log\frac{p(y_k|x_k,u_k)}{p(y_k|u_k)}dx_kdy_k \end{aligned} \tag{39}$$

where we can then apply the particle approximation

$$\begin{aligned}\tilde{J}^N &= \frac{1}{N}\sum_{i=1}^{N}\log\frac{p(y_k^{(i)}|x_k^{(i)},u_k)}{p(y_k^{(i)}|u_k)}\\ &= \frac{1}{N}\sum_{i=1}^{N}\log\frac{p(y_k^{(i)}|x_k^{(i)},u_k)}{\frac{1}{N}\sum_{j=1}^{N}p(y_k^{(i)}|x_k^{(j)},u_k)}. \end{aligned} \tag{40}$$

C. RMS errors

In order to apply the criterion based on RMS errors of the MMSE estimate described in Section IV-A, we can use the goal function

$$\begin{aligned}&\gamma(k,x_k,y_k,u_k)\\ &= (x_k-\hat{x}_k(y_k))'(x_k-\hat{x}_k(y_k)) \end{aligned} \tag{41}$$

where $\hat{x}_k$ is the MMSE estimate for time index k, and we have emphasized its dependence on the measurement y_k. We can now substitute (41) in (33) and obtain

$$\tilde{J}^N = \frac{1}{N}\sum_{i=1}^{N}\left(x_k^{(i)}-\hat{x}_k(y_k^{(i)})\right)'\left(x_k^{(i)}-\hat{x}_k(y_k^{(i)})\right) \tag{42}$$

where we can approximate $\hat{x}_k(y_k^{(i)})$ as follows:

$$\begin{aligned}\hat{x}_k(y_k^{(i)}) &= \int x_kp(x_k|y_k^{(i)},u)dx_k\\ &= \int x_k\frac{p(y_k^{(i)}|x_k,u)p(x_k)}{p(y_k^{(i)}|u)}dx_k\\ &\approx \frac{\sum_{j=1}^{N}x_k^{(j)}p(y_k^{(i)}|x_k^{(j)},u)}{\sum_{l=1}^{N}p(y_k^{(i)}|x_k^{(l)},u)}. \end{aligned} \tag{43}$$

D. Covariance log-determinant

As shown by (18), for sensor management purposes, minimization of the square of RMS errors (according to criterion (41)) is equivalent to minimization of the trace of the covariance matrix. However, we may think about minimizing its determinant instead, or more precisely, the logarithm of its determinant (a possible motivation is its equivalence to the KL divergence criterion in the Gaussian case – see [11]).

The goal function for the criterion based on the covariance log-determinant is given by

$$\begin{aligned}&\gamma(k,y_k,u_k)\\ &= \log\left|\int(x_k-\hat{x}_k(y_k))(x_k-\hat{x}_k(y_k))'p(x_k|y_k,u_k)dx_k\right|\\ &= \log\Bigg|\int(x_k-\hat{x}_k(y_k))(x_k-\hat{x}_k(y_k))'\\ &\quad\times\frac{p(y_k|x_k,u_k)p(x_k)}{p(y_k|u_k)}\Bigg|dx_k \end{aligned} \tag{44}$$

where $\hat{x}_k$ is the MMSE estimate. Substituting (44) in particle approximation (33), we obtain

$$\begin{aligned}\tilde{J}^N &= \frac{1}{N}\sum_{i=1}^{N}\log\Bigg|\int\left(x_k-\hat{x}_k(y_k^{(i)})\right)\left(x_k-\hat{x}_k(y_k^{(i)})\right)'\\ &\quad\times\frac{p(y_k^{(i)}|x_k,u_k)p(x_k)}{p(y_k^{(i)}|u_k)}dx_k\Bigg|\\ &= \frac{1}{N}\sum_{i=1}^{N}\log\Bigg|\frac{1}{\sum_{l=1}^{N}p(y_k^{(i)}|x_k^{(l)},u)}\\ &\quad\times\sum_{j=1}^{N}\left(x_k^{(j)}-\hat{x}_k(y_k^{(i)})\right)\left(x_k^{(j)}-\hat{x}_k(y_k^{(i)})\right)'\\ &\quad\times p(y_k^{(i)}|x_k^{(j)},u_k)\Bigg| \end{aligned} \tag{45}$$

where $\hat{x}_k(y_k^{(i)})$ can be computed according to (43).

E. OSPA

For the OSPA metric described in Section IV-B, we can use

$$\gamma(k, x_k, y_k, u_k) \\ = \epsilon_p^{(c)}(x_k, \hat{x}_k(y_k))^p \tag{46}$$

where $\epsilon_p^{(c)}$ is the OSPA metric defined by (20), and $\hat{x}$ is the corresponding estimate according to (21). Thus by substituting (46) in (33), we obtain

$$\tilde{J}^N = \frac{1}{N}\sum_{i=1}^{N} \epsilon_p^{(c)}(x_k^{(i)}, \hat{x}_k(y_k^{(i)}))^p \tag{47}$$

where $\hat{x}_k(y_k^{(i)})$ is approximated according to

$$\begin{aligned}\hat{x}_k(y_k^{(i)}) &= \arg\min_{\hat{x}_k^*} \int \epsilon_p^{(c)}(x_k, \hat{x}_k^*)^p p(x_k|y_k^{(i)}, u_k) dx_k \\ &= \arg\min_{\hat{x}_k^*} \int \epsilon_p^{(c)}(x_k, \hat{x}_k^*)^p \frac{p(y_k^{(i)}|x_k, u)p(x_k)}{p(y_k^{(i)}|u)} dx_k \\ &\approx \arg\min_{\hat{x}_k^*} \sum_{j=1}^{N} \epsilon_p^{(c)}(x_k^{(j)}, \hat{x}_k^*)^p p(y_k^{(i)}|x_k^{(j)}, u). \end{aligned} \tag{48}$$

Needless to say, computing (48) can be extremely difficult due to the need of searching for the optimal $\hat{x}_k^*$ in the state space.

If the number of objects is fixed, and we choose $p = 2$ and $c = \infty$, we may explore the relation with the MMSE estimate as described in Section IV-B. For instance, for two targets, we can use the following particle-based algorithm (proposed by Svensson et al. [15]):

1) Set test $= 0$ and compute $\hat{x}_k(y_k^{(i)})$ according to (43)
2) While test $= 0$, do
   a) Set test $= 1$
   b) For $j = 1, \dots, N$
      i) If $\|\hat{x}_k(y_k^{(i)}) - x_k^{(j)}\|^2 > \|\hat{x}_k(y_k^{(i)}) - \chi x_k^{(j)}\|^2$ (where $\chi x_k^{(j)}$ corresponds to $x_k^{(j)}$ with the states of both objects permuted) set test $= 0$ and make $x_k^{(j)} = \chi x_k^{(j)}$ (for this algorithm only)
   c) Compute $\hat{x}_k(y_k^{(i)})$ according to (43)

We have empirically verified that, in many cases, this algorithm converges to the same OSPA-based estimate obtained using (48), but its theoretical convergence properties have not been yet studied in detail.

VII. Conclusions

In this memorandum, we derived approximations, based on particle filters, for open loop (OLC) and open loop feedback (OLFC) sensor management, which can be useful for both practical problems and empirical studies on sensor management criteria.

Our work may also be a good starting point for those seeking solutions for closed loop (CLC) sensor management or for reduction of computational cost of long-term OLC and OFLC.

VIII. Acknowledgements

The research leading to these results has received funding from the EU's Seventh Framework Programme under grant agreement n° 238710. The research has been carried out in the MC IMPULSE project: https://mcimpulse.isy.liu.se.

This research has been also supported by the Netherlands Organisation for Scientific Research (NWO) under the Casimir program, contract 018.003.004. Under this grant Yvo Boers holds a part-time position at the Department of Applied Mathematics at the University of Twente.

The authors would also like to thank Hans Driessen (Thales Nederland B.V.) for contributing to the memorandum.

References

[1] J. L. Williams, J. W. Fisher, and A. S. Willsky, "Approximate dynamic programming for communication-constrained sensor network management," *IEEE Trans. Signal Process.*, vol. 55, no. 8, pp. 4300–4311, 2007.

[2] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Process. Mag.*, vol. 19, no. 2, pp. 61–72, 2002.

[3] A. Doucet, B. Vo, C. Andrieu, and M. D. Thomas, "Particle filtering for multi-target tracking and sensor management," in *Proc. 5th International Conference on Information Fusion*, Annapolis, MD, Jul. 7–11, 2002, pp. 474–481.

[4] C. Kreucher, K. Kastella, and A. O. Hero, "Multi-target sensor management using alpha-divergence measures," in *Lecture Notes in Computer Science, Proc. of the 2nd International Conference on Information Processing in Sensor Networks*, no. 2634, 2003, pp. 209–222.

[5] C. Kreucher, A. O. Hero, and K. Kastella, "A comparison of task driven and information driven sensor management for target tracking," in *Proc. IEEE 44th Conference on Decision and Control*, Dec. 12–15, 2005, pp. 4004–4009.

[6] A. Bagchi, *Optimal Control of Stochastic Systems*. Upper Saddle River, NJ: Prentice-Hall, 1993.

[7] C. Kreucher and A. O. Hero, "Non-myopic approaches to scheduling agile sensors for multitarget detection, tracking, and identification," in *Proc. IEEE Conference on Acoustics, Speech, and Signal Processing*, vol. 5, Philadelphia, PA, Mar. 18–23, 2005, pp. 885–888.

[8] A. Kuwertz, M. F. Huber, and F. Sawo, "Multi-step sensor management for localizing movable sources of spatially distributed phenomena," in *Proc. 13th International Conference on Information Fusion*, Edinburgh, UK, Jul. 26–29, 2010.

[9] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Belmont, MA: Athena Scientific, 2000.

[10] T. Hanselmann, M. Morelande, B. Moran, and P. Sarunic, "Sensor scheduling for multiple target tracking and detection using passive measurements," in *Proc. 11th International Conference on Information Fusion*, Cologne, Germany, Jun. 30, Jul. 1–3, 2008, pp. 1528–1535.

[11] E. H. Aoki, A. Bagchi, P. Mandal, and Y. Boers, "A theoretical look at information-driven sensor management criterions," in *(submitted to) The 14th International Conference of Information Fusion*, Chicago, IL, Jul. 5–8, 2011.

[12] D. Schuhmacher, B.-T. Vo, , and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, 2008.

[13] M. Guerriero, L. Svensson, D. Svensson, and P. Willett, "Shooting two birds with two bullets: how to find Minimum Mean OSPA estimates," in *Proc. 13th International Conference on Information Fusion*, Edinburgh, UK, Jul. 26–29, 2010.

[14] C. Kreucher, A. O. Hero, K. Kastella, and D. Chang, "Efficient methods of non-myopic sensor management for multitarget tracking," in *Proc. IEEE 43th Conference on Decision and Control*, Atlantis, Bahamas, Dec. 14–17, 2004, pp. 722–727.

[15] D. Svensson, L. Svensson, M. Guerriero, and P. Willett, "On the calculation of minimum mean OSPA estimates," Department of Signals and Systems, Chalmers University of Technology, Tech. Rep. R004/2011, 2011.