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

A Mean-field Control-oriented Approach to Particle Filtering

Tao Yang, Prashant G. Mehta, Sean P. Meyn

***Abstract*— A new formulation of the particle filter for nonlinear filtering is presented, based on concepts from optimal control, and from the mean-field game theory framework of Huang et. al. [8]. The optimal control is chosen so that the posterior distribution of a particle matches as closely as possible the posterior distribution of the true state, given the observations. In the infinite-N limit, the empirical distribution of ensemble particles converges to the posterior distribution of an individual particle.**

The cost function in this control problem is the Kullback-Leibler (K-L) divergence between the actual posterior, and the posterior of any particle. The optimal control input is characterized by a certain Euler-Lagrange (E-L) equation.

A numerical algorithm is introduced and implemented in two general examples: A linear SDE with partial linear observations, and a nonlinear oscillator perturbed by white noise, with partial nonlinear observations.

I. INTRODUCTION

We consider a scalar filtering problem:

$$\mathrm{d}X_t = a(X_t)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t, \tag{1}$$

$$\mathrm{d}Z_t = h(X_t)\,\mathrm{d}t + \sigma_W\,\mathrm{d}W_t, \tag{2}$$

where $X_t \in \mathbb{R}$ is the state at time t, $Z_t \in \mathbb{R}$ is the observation process, $a(\cdot)$, $h(\cdot)$ are C^1 functions, and $\{B_t\}$, $\{W_t\}$ are mutually independent standard Wiener processes.

The objective of the filtering problem is to obtain the posterior distribution p^* of X_t given the history $\mathscr{Z}_t := \sigma(Z_s : s \le t)$. If $a(\cdot)$, $h(\cdot)$ are linear functions, then the solution is given by the well-known Kalman filter.

For nonlinear systems, the particle filter is a simulation-based algorithm to approximate the filtering task [7], [6]. The key step is the construction of N stochastic processes $\{X_t^i\}_{i=1}^N$. For each time t the empirical distribution formed by the "particle population" is used to approximate the conditional distribution. Recall that this is defined for any measurable set $A \subset \mathbb{R}$ by,

$$p^{(N)}(A,t) = \frac{1}{N}\sum_{i=1}^{N} 1\{X_t^i \in A\} \tag{3}$$

A common approach in particle filtering is called the sequential importance sampling, where particles are generated according to their importance weight at every time stage [6], [2]. By choosing the sampling mechanism properly, particle filtering can approximately propagate the posterior distribution, with the accuracy improving as N increases [5]. Several modifications and improvements of these and other related algorithms have been suggested when dynamics are linear and/or noise is Gaussian; cf., [3], [1], [9].

The objective of this paper is to introduce an alternative approach to the construction of a particle filter for (1)-(2) inspired from the mean-field optimal control techniques; cf., [8], [11]. In this approach, the model for the i^{th} particle is defined by a controlled system,

$$\mathrm{d}X_t^i = a(X_t^i)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i + \mathrm{d}U_t^i, \tag{4}$$

where $X_t^i \in \mathbb{R}$ is the state for the i^{th} particle at time t, U_t^i is its control input, and $\{B_t^i\}$ are mutually independent standard Wiener processes. We assume the initial conditions $\{X_0^i\}_{i=1}^N$ are i.i.d., and drawn from initial distribution $p^*(x,0)$ of X_0.

With a specific formulation for (4), the resulting particle trajectories $\{X_t^i\}_{i=1}^N$ (obtained under optimal control input U_t^i) define a particle filter, in the sense that the empirical distribution $p^{(N)}(\cdot,t)$ approximates the posterior distribution $p^*(X_t|Z^t)$ for large N.

We cast the synthesis of control input as an optimization problem, with the Kullback-Leibler metric serving as the cost function. The optimal control input is obtained via solution of the associated Euler-Lagrange (E-L) boundary value problem (BVP).

The outline of this paper is as follows. We begin with a discussion of the continuous-discrete filtering problem: the equation for dynamics is given by (1), but the observations are made only at discrete times. In section II, we formulate the optimization problem for this case and describe the E-L BVP. For the linear Gaussian case, we obtain an explicit solution of the BVP.

In section III, we consider the continuous-continuous filtering problem (for (1)-(2)) as a limiting case of the continuous-discrete problem. For the linear Gaussian case, the solution is explicitly obtained, and shown to be consistent with the Kalman filter.

In section IV, we describe the results of numerical experiments. A numerical algorithm is introduced and implemented in two general examples: A linear SDE with partial linear observations, and a nonlinear oscillator perturbed by white noise, with partial nonlinear observations.

Apart from the serving the pedagogical purpose, the scalar case is especially relevant to filtering in oscillator models. These models (also considered in our earlier mean-field control paper [11]) provide one of the main motivations for the present work. Extension to the multi-state case is planned for the journal version of this paper.

Financial support from the NSF grant EECS-0925534 and the AFOSR grant FA9550-09-1-0190 is gratefully acknowledged.

T. Yang and P. G. Mehta are with the Coordinated Science Laboratory and the Department of Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign (UIUC) taoyang1@illinois.edu; mehtapg@illinois.edu

S. P. Meyn is with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory at UIUC meyn@illinois.edu

II. Continuous-discrete time filtering

We consider the continuous-discrete time filtering problem. The equation for dynamics is given by (1), and the observations are made only at discrete times $\{t_n\}$:

$$Z_{t_n} = h(X_{t_n}) + W_{t_n}, \tag{5}$$

where $\{W_{t_n}\}$ is i.i.d and drawn from $\boldsymbol{N}(0,\sigma_W^2)$.

The particle model in this case is a hybrid dynamical system: For $t \in [t_{n-1}, t_n)$, the i^{th} particle evolves according to the stochastic differential equation,

$$\mathrm{d}X_t^i = a(X_t^i)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i, \quad t_{n-1} \le t < t_n, \tag{6}$$

where the initial condition $X_{t_{n-1}}^i$ is given. At time $t = t_n$ there is a potential jump that is determined by the input $U_{t_n}^i$:

$$X_{t_n}^i = X_{t_n^-}^i + U_{t_n}^i, \tag{7}$$

where $X_{t_n^-}^i$ denotes the right limit of $\{X_t^i : t_{n-1} \le t < t_n\}$. The specification (7) defines the initial condition for the process on the next interval $[t_n, t_{n+1})$.

The filtering problem is to construct a control law that defines $\{U_{t_n}^i : n \ge 1\}$ such that, for each $n \ge 1$, the resulting empirical distribution $p^{(N)}$ approximates the posterior distribution of X_{t_n} given the history $\mathscr{Z}_n := \sigma(Z_{t_k} : k \le n)$. To solve this problem we first define "belief maps" that propagate the conditional distributions of X and X^i.

A. Belief maps

For each n we let $p_n^*(x)$ denote the probability density function (pdf) for the conditional distribution of X_{t_n} given $\mathscr{Z}_n$, and we let $p_n^{*-}(x)$ denote the conditional distribution of X_{t_n} given $\mathscr{Z}_{n-1}$. Similarly, we let $p_n(x)$ denote the pdf for the conditional distribution of $X_{t_n}^i$ given $\mathscr{Z}_n$, and we let $p_n^-(x)$ denote the conditional distribution of $X_{t_n}^i$ given $\mathscr{Z}_{n-1}$. We will see that these densities evolve according to deterministic equations of the form,

$$p_n^* = \mathscr{P}^*(p_{n-1}^*, Z_{t_n}), \qquad p_n = \mathscr{P}(p_{n-1}, Z_{t_n}). \tag{8}$$

The mappings $\mathscr{P}^*$ and $\mathscr{P}$ can be decomposed into two parts. The first part is identical for each of these mappings: The transformation that takes p_{n-1} to p_n^- coincides with the mapping from p_{n-1}^* to p_n^{*-}. In each case it is defined by the kolmogorov's forward equation associated with the diffusion on $[t_{n-1}, t_n)$.

Consider now how $p_n^{*-}(x)$ is mapped to $p_n^*(x)$ to define $\mathscr{P}^*$. Given the observation Z_{t_n} made at time $t = t_n$, the pdf for the actual state is updated using Bayes rule:

$$p_n^*(s) = \frac{p_n^{*-}(s) \cdot p_{Z|X}(Z_{t_n}|s)}{p_Z(Z_{t_n})}, \quad s \in \mathbb{R}, \tag{9}$$

where p_Z denotes the pdf for Z_{t_n}, and $p_{Z|X}(\cdot \mid s)$ denotes the conditional distribution of Z_{t_n} given $X_{t_n} = s$. Applying (5) gives,

$$p_{Z|X}(Z_{t_n} \mid s) = \frac{1}{\sqrt{2\pi\sigma_W^2}} \exp\left(-\frac{(Z_{t_n} - h(s))^2}{2\sigma_W^2}\right).$$

Combining (9) with the forward equation defines $\mathscr{P}^*$.

To complete the definition of $\mathscr{P}$ we impose further structure on the input. At time $t = t_n$, we seek a control input $U_{t_n}^i$ that is an *admissible function* of $X_{t_n^-}^i$, and $\{Z_{t_k} : k \le n\}$.

Definition 1 (Admissible function): A function $u : \mathbb{R} \to \mathbb{R}$ is said to be *admissible* if u is twice continuously differentiable and

$$\lim_{x \to \pm\infty} u(x) p_n^-(x) = 0.$$

The space of admissible functions is denoted as C_b^2. ∎

We suppress the dependency on the observations, writing $U_{t_n}^i = u(x)$ when $X_{t_n^-}^i = x$. We further assume that $1 + u'(x)$ is non-zero for all x. In this case we can write,

$$p_n(s) = \frac{p_n^-(x)}{|1 + u'(x)|} \tag{10}$$

where $s = x + u(x)$ and $u'(x) = \frac{\mathrm{d}}{\mathrm{d}x} u(x)$.

B. Variational Problem

Our goal is to choose the function u so that the mapping (10) approximates the mapping (9). More specifically, given the pdf p_{n-1} we have already defined the mapping $\mathscr{P}$ so that $p_n = \mathscr{P}(p_{n-1}, Z_{t_n})$. We denote $\hat{p}_n^* = \mathscr{P}^*(p_{n-1}, Z_{t_n})$, and choose $u(x)$ so that these pdfs are as close as possible. We approach this goal through the formulation of an optimization problem with respect to the Kullback-Leibler (KL) divergence metric. That is, at time $t = t_n$, the function, denoted as u_n, is the solution to the following optimization problem,

$$u_n(x) = \arg\min_u \mathrm{KL}\left(p_n \| \hat{p}_n^*\right). \tag{11}$$

Based on the definitions, for any u the KL divergence can be expressed,

$$\mathrm{KL}(p_n \| \hat{p}_n^*) = -\int_{\mathbb{R}} p_n^-(x) \Big\{ \ln|1 + u'(x)| + \ln(p_n^-(x + u(x)) p_{Z|X}(Z_{t_n}|x + u(x))) \Big\} \mathrm{d}x + C, \tag{12}$$

where $C = \int_{\mathbb{R}} p_n^-(x) \ln(p_n^-(x) p_Z(Z_{t_n}))\,\mathrm{d}x$ is a constant that does not depend on u; cf., Appendix A for the calculation.

The solution to (11) is described in the following proposition, whose proof appears in Appendix B.

Proposition 1: Suppose that the admissible function $u : \mathbb{R} \to \mathbb{R}$ is a minimizer for the optimization problem (11). Then it is a solution of the following Euler-Lagrange (E-L) BVP:

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{p_n^-(x)}{|1 + u'(x)|}\right) = p_n^-(x) \frac{\partial}{\partial u}\left(\ln(p_n^-(x + u) p_{Z|X}(Z_{t_n}|x + u))\right), \tag{13}$$

with boundary condition $\lim_{x \to \pm\infty} u(x) p_n^-(x) = 0$. ∎

We refer to the minimizer as the *optimal control function*.

C. Forward equation

We now provide a complete description of the forward equation that is intended to approximate the continuous-discrete time particle filter (6)-(7) for large N.

At measurement times $t = t_n$, the optimal control function is obtained as a solution to the E-L BVP,

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{p_n^-(x)}{|1+u'(x)|}\right) = p_n^-(x)\frac{\partial}{\partial u}\left(\ln(p_n^-(x+u)p_{Z|X}(Z_{t_n}|x+u))\right), \tag{14}$$

with boundary condition $\lim_{x\to\pm\infty} u(x)p_n^-(x) = 0$. We denote the solution of (14) as $u_n(x)$.

We let $p(x,t)$ denote the conditional distribution of the particle X_t^i given $\mathscr{Z}_t$. The evolution of the density $p(x,t)$ is given by the forward evolution equation:

$$\frac{\partial p(x,t)}{\partial t} + \frac{\partial}{\partial x}\left\{\left(a(x) + \sum_n u_n(x)\delta(t-t_n)\right)p(x,t)\right\} = \frac{\sigma_B^2}{2}\frac{\partial^2 p(x,t)}{\partial x^2}, \tag{15}$$

where $\delta(t-t_n)$ is the Dirac delta function at time $t = t_n$. The two equations (14)-(15) are coupled because $p_n^-(x) = p(x,t_n^-)$ by construction.

D. Example: Linear Gaussian case

In this section, we assume dynamics and observation equations are both linear:

$$\mathrm{d}X_t = a\,X_t\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t, \tag{16}$$

$$Z_{t_n} = h\,X_{t_n} + W_{t_n}, \tag{17}$$

where a, h are now real numbers. We assume that the initial distribution $p^*(x,0)$ is Gaussian with mean μ_0 and variance Σ_0.

For the linear Gaussian case, the E-L equation (13) can be solved in closed form. The proof appears in Appendix C.

Lemma 1: Suppose $p(x,t)$ is assumed to be Gaussian with mean μ_t and variance Σ_t. Then the solution of E-L BVP (13) is given by:

$$u_n(x) = k_n x + c_n, \tag{18}$$

where

$$k_n = \sqrt{\frac{\sigma_W^2}{\sigma_W^2 + \Sigma_{t_n^-}h^2}} - 1, \tag{19}$$

$$c_n = -\left(\frac{h^2\Sigma_{t_n^-}}{h^2\Sigma_{t_n^-}+\sigma_W^2} + k_n\right)\mu_{t_n^-} + \frac{h\Sigma_{t_n^-}}{h^2\Sigma_{t_n^-}+\sigma_W^2}Z_{t_n}. \tag{20}$$

■

The optimal control function $u_n(x)$ yields the following hybrid dynamical system model for the particle filter:

$$t \in [t_{n-1}, t_n): \quad \mathrm{d}X_t^i = a\,X_t^i\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i;$$
$$t = t_n: \quad X_{t_n}^i = X_{t_n^-}^i + k_n\,X_{t_n^-}^i + c_n, \tag{21}$$

where k_n and c_n are as in (19)-(20).

Finally, with the control function (19)-(20) obtained by E-L BVP, the forward equation (15) reduces to the following recursive equations for only the mean and the variance:

$$t \in [t_{n-1}, t_n): \begin{cases} \mathrm{d}\mu_t/\mathrm{d}t = a\mu_t, \\ \mathrm{d}\Sigma_t/\mathrm{d}t = 2a\Sigma_t + \sigma_B^2; \end{cases} \tag{22}$$

$$t = t_n: \begin{cases} \mu_{t_n} = (\sigma_W^2\mu_{t_n^-} + h\Sigma_{t_n^-}Z_{t_n})/\Sigma_{t_n^-}^+, \\ \Sigma_{t_n} = \sigma_W^2\Sigma_{t_n^-}/\Sigma_{t_n^-}^+; \end{cases} \tag{23}$$

where we define $\Sigma_{t_n^-}^+ := h^2\Sigma_{t_n^-} + \sigma_W^2$.

One can also readily verify that the solution of the BVP is consistent with the Kalman filtering solution. These calculations for the discrete time case appear in the Appendix D.

The implementation of the particle filter appears to suffer from the same drawback as importance sampling: We must compute the object that we wish to simulate. In practice, the conditional mean $\mu_{t_n^-}$ and the variance $\Sigma_{t_n^-}$ appearing in the recursions above will be estimated from the ensemble $\{X_{t_n^-}^i\}_{i=1}^N$:

$$\mu_{t_n^-} \approx \bar{\mu}_{t_n^-}^{(N)} := \frac{1}{N}\sum X_{t_n^-}^i,$$
$$\Sigma_{t_n^-} \approx \bar{\Sigma}_{t_n^-}^{(N)} := \frac{1}{N-1}\sum (X_{t_n^-}^i - \bar{\mu}_{t_n^-}^{(N)})^2.$$

Further discussion is contained at the end of Section III-A.

III. Continuous-time filtering

We now return to the filtering problem (1)-(2) introduced in Section I. The solution for the continuous time problem is obtained by considering an observation time sequence $\{t_n\}$ where the time between observations, $t_{n+1} - t_n$, goes to zero.

In the following, we describe the continuous-time filter: We assume the initial conditions $\{X_0^i\}_{i=1}^N$ are i.i.d., and drawn from initial distribution $p^*(x,0)$ of X_0. The dynamics of the i^{th} particle are defined by a controlled system:

$$\mathrm{d}X_t^i = a(X_t^i)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i + v(X_t^i,t)\,\mathrm{d}I_t^i + \frac{1}{2}\sigma_W^2 v(X_t^i,t)v'(X_t^i,t)\,\mathrm{d}t \tag{24}$$

in which $v'(x,t) = \frac{\partial v}{\partial x}(x,t)$, where $\{B_t^i\}$ are mutually independent standard Wiener processes, and I^i is intended to mirror the *innovations process* that appears in the nonlinear filter,

$$\mathrm{d}I_t^i =: \mathrm{d}Z_t - \frac{1}{2}(h(X_t^i) + \hat{h})\,\mathrm{d}t \tag{25}$$

where $\hat{h} = \frac{1}{N}\sum_{i=1}^N h(X_t^i)$.

The gain function $v(x,t)$ is the solution to the Euler-Lagrange boundary value problem (E-L BVP):

$$-\frac{\partial}{\partial x}\left(\frac{1}{p(x,t)}\frac{\partial}{\partial x}\{p(x,t)v(x,t)\}\right) = \frac{1}{\sigma_W^2}h'(x), \tag{26}$$

with boundary conditions $\lim_{x\to\pm\infty} p(x,t)v(x,t) = 0$, where $h'(x) = \frac{\mathrm{d}}{\mathrm{d}x}h(x)$, and $p(x,t)$ denotes the conditional distribution of X_t^i given $\mathscr{Z}_t$.

The filter (24) represents the continuous-time counterpart of the continuous-discrete time filter (6)-(7). Likewise, the E-L BVP (26) is the continuous-time counterpart of the E-L BVP (13) derived earlier for the continuous-discrete time problem. The details of derivation of the continuous-time filter are omitted here on account of space.

The evolution of the conditional distribution $p(x,t)$ is now described by the Kolmogorov's forward equation for the controlled system (24): Setting

$$u(x,t) = v(x,t)\left(-\frac{1}{2}(h(x)+\hat{h}) + \frac{1}{2}\sigma_W^2 v'(x,t)\right),$$

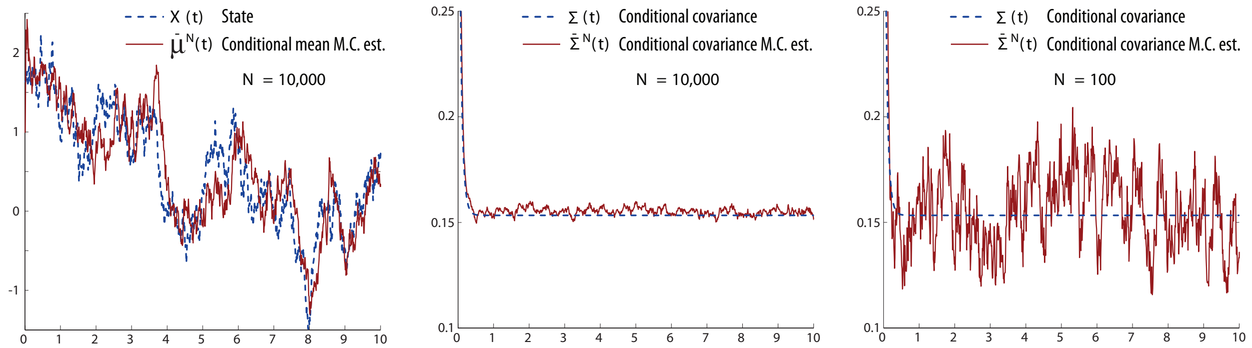

Fig. 1. (a) Comparison of the true state $\{X_t\}$ and the conditional mean $\{\bar{\mu}_t^{(N)}\}$. (b) and (c) Plots of estimated conditional covariance with $N = 10{,}000$ and $N = 100$ particles, respectively. For comparison, the true conditional covariance obtained using Kalman filtering equations is also shown.

the forward equation is given by

$$\begin{aligned} \mathrm{d}p = \mathscr{L}^\dagger p \, \mathrm{d}t - \frac{\partial}{\partial x}(vp)\, \mathrm{d}Z_t \\ - \frac{\partial}{\partial x}(up)\, \mathrm{d}t + \sigma_W^2 \frac{1}{2}\frac{\partial^2}{\partial x^2}\left(pv^2\right) \mathrm{d}t, \end{aligned} \tag{27}$$

where $\mathscr{L}^\dagger p = -\frac{\partial(pa)}{\partial x} + \frac{\sigma_B^2}{2}\frac{\partial^2 p}{\partial x^2}$.

We refer to particle system described by (24) as the *feedback particle filter*. In the following, we describe the feedback particle filter for the linear Gaussian case.

A. Example: Linear Gaussian case

Consider the scalar linear model in continuous time defined by,

$$\mathrm{d}X_t = a\, X_t \, \mathrm{d}t + \sigma_B \, \mathrm{d}B_t, \tag{28}$$

$$\mathrm{d}Z_t = h\, X_t \, \mathrm{d}t + \sigma_W \, \mathrm{d}W_t, \tag{29}$$

where a, h are real numbers. We assume that the initial distribution $p^*(x,0)$ is Gaussian with mean μ_0 and variance Σ_0.

The following lemma provides the solution of the optimal control function $v(x,t)$ in the linear Gaussian case.

Lemma 2: Consider the linear observation equation (29). Suppose $p(x,t) = \frac{1}{\sqrt{2\pi\Sigma_t}}\exp(-\frac{(x-\mu_t)^2}{2\Sigma_t})$ is assumed to be Gaussian with mean μ_t and variance Σ_t. Then the solution of E-L BVP (26) is given by:

$$v(x,t) = \frac{h\Sigma_t}{\sigma_W^2} \tag{30}$$

■

The formula (30) is verified by direct substitution in the ODE (26) where the distribution p is Gaussian and $h'(x) = h$ which is a constant.

The optimal control yields the following form for the particle filter in this linear Gaussian model:

$$\mathrm{d}X_t^i = a\, X_t^i \, \mathrm{d}t + \sigma_B \, \mathrm{d}B_t^i + \frac{h\Sigma_t}{\sigma_W^2}\left(\mathrm{d}Z_t - h\frac{X_t^i + \mu_t}{2}\,\mathrm{d}t\right). \tag{31}$$

Now we show that $p = p^*$ in this case. That is, the conditional distributions of X and X^i coincide, and are defined by the well-known dynamic equations that characterize the mean and the variance of the continuous-time Kalman filter.

Theorem 1: Consider the linear Gaussian filtering problem defined by the state-observation equations (28,29). In this case the posterior distributions of X and X^i are Gaussian, whose conditional mean and covariance are given by the respective stochastic differential equation and ordinary differential equation,

$$\mathrm{d}\mu_t = a\mu_t \, \mathrm{d}t + \frac{h\Sigma_t}{\sigma_W^2}\Big(\mathrm{d}Z_t - h\mu_t \, \mathrm{d}t\Big) \tag{32}$$

$$\frac{\mathrm{d}\Sigma_t}{\mathrm{d}t} = 2a\Sigma_t + \sigma_B^2 - \frac{h^2\Sigma_t^2}{\sigma_W^2} \tag{33}$$

■

The proof of Theorem 1 appears in Appendix E.

Notice that particle system (31) is not practical since it requires computation of the conditional mean and variance $\{\mu_t, \Sigma_t\}$. If we are to compute these quantities, then there is no reason to run a particle filter!

In practice $\{\mu_t, \Sigma_t\}$ are approximated as sample means and sample covariances from the ensemble $\{X_t^i\}_{i=1}^N$:

$$\begin{aligned} \mu_t \approx \bar{\mu}_t^{(N)} &:= \frac{1}{N}\sum_{i=1}^N X_t^i, \\ \Sigma_t \approx \bar{\Sigma}_t^{(N)} &:= \frac{1}{N-1}\sum_{i=1}^N (X_t^i - \bar{\mu}_t^{(N)})^2. \end{aligned} \tag{34}$$

The resulting equation (31) for the i^{th} particle is given by

$$\mathrm{d}X_t^i = a\, X_t^i \, \mathrm{d}t + \sigma_B \, \mathrm{d}B_t^i + \frac{h\bar{\Sigma}_t^{(N)}}{\sigma_W^2}\left(\mathrm{d}Z_t - h\frac{X_t^i + \bar{\mu}_t^{(N)}}{2}\,\mathrm{d}t\right). \tag{35}$$

It is very similar to the mean-field "synchronization-type" control laws and oblivious equilibria constructions as in [8], [11], [10]. For large N, the model (31) represents the mean-field approximation of (35).

For any t and any set $A \in \mathscr{B}(\mathbb{R})$, we define the empirical distribution, $p^{(N)}(A,t) =: \frac{1}{N}\sum_{i=1}^N 1\{X_t^i \in A\}$. Using Theorem 1, the empirical distribution $p^{(N)}$ approximates p^*, in the sense that:

$$\lim_{N\to\infty} p^{(N)}(A,t) = \int_A p^*(x,t)\, \mathrm{d}x.$$

IV. NUMERICS

In this section, we provide numerical verification of the particle filter for a linear Gaussian system (section IV-A) and a nonlinear oscillator system (section IV-B).

A. Linear Gaussian case

Consider the following system:

$$dX_t = -0.5\, X_t\, dt + 1\, dB_t, \tag{36}$$

$$dZ_t = 3\, X_t\, dt + 0.5\, dW_t, \tag{37}$$

where $\{B_t\},\{W_t\}$ are mutually independent standard Wiener process.

The particle filter comprises of N particles where the dynamic of the ith particle is given by:

$$dX_t^i = -0.5\, X_t^i\, dt + 1\, dB_t^i + dU_t^i, \tag{38}$$

where

$$dU_t^i = \frac{3\,\bar{\Sigma}_t^{(N)}}{0.5^2}\Big[dZ_t - 3\frac{X_t^i + \bar{\mu}_t^{(N)}}{2}\, dt\Big],$$

and $\{B_t^i\}$ are mutually independent standard Wiener process. We initialize the particle system by drawing initial conditions $\{X_0^i\}_{i=1}^N$ from the distribution $\boldsymbol{N}(1,1)$. In the simulation discussed next, the mean $\bar{\mu}_t^{(N)}$ and the variance $\bar{\Sigma}_t^{(N)}$ are obtained from the ensemble $\{X_t^i\}_{i=1}^N$ according to (34).

Figure 1 summarizes some of the results of the numerical experiments: Part (a) depicts a sample path of the state $\{X_t\}$ and the mean $\{\bar{\mu}_t^{(N)}\}$ obtained using a particle filter with $N = 10,000$ particles. Part (b) provides a comparison between the estimated variance $\bar{\Sigma}_t^{(N)}$ and the true error variance Σ_t that one would obtain by using the Kalman filtering equations. The accuracy of the results is sensitive to the number of particles. For example, part (c) of the figure provides a comparison of the variance with $N = 100$ particles.

B. Nonlinear oscillator case

We consider the filtering problem for a nonlinear oscillator:

$$d\theta_t = \omega\, dt + \sigma_B\, dB_t \mod 2\pi, \tag{39}$$

$$dZ_t = h(\theta_t)\, dt + \sigma_W\, dW_t, \tag{40}$$

where ω is the frequency, $h(\theta) = \frac{1+\cos\theta}{2}$, and $\{B_t\}$ and $\{W_t\}$ are mutually independent standard Wiener process. For numerical simulations, we pick $\omega = 1$ and the standard deviation parameters $\sigma_B = 0.5$ and $\sigma_W = 0.4$. We consider oscillator models because of their significance to applications including neuroscience; cf., [11].

The feedback particle filter is given by:

$$d\theta_t^i = \omega\, dt + \sigma_B\, dB_t^i + v(\theta_t^i,t)\Big[dZ_t - \frac{1}{2}(h(\theta_t^i) + \hat{h})\, dt\Big] + \frac{1}{2}\sigma_W^2 v(\theta_t^i,t) v'(\theta_t^i,t)\, dt \mod 2\pi, \quad i = 1,\dots,N. \tag{41}$$

where $\hat{h} = \frac{1}{N}\sum_{j=1}^N h(\theta_t^j)$, and the function $v(\theta,t)$ is obtained via the solution of the E-L equation:

$$-\frac{\partial}{\partial\theta}\left(\frac{1}{p(\theta,t)}\frac{\partial}{\partial\theta}\{p(\theta,t)v(\theta,t)\}\right) = -\frac{\sin\theta}{2\sigma_W^2}. \tag{42}$$

Although the equation (42) can be solved numerically to obtain the optimal control function $v(\theta,t)$, here we investigate a solution based on perturbation method. Suppose, at some time t, $p(\theta,t) = \frac{1}{2\pi} =: p_0$, the uniform density. In this case, the E-L equation is given by:

$$\partial_{\theta\theta} v = \frac{\sin\theta}{2\sigma_W^2}.$$

A straightforward calculation shows that the solution in this case is given by

$$v(\theta,t) = -\frac{\sin\theta}{2\sigma_W^2} =: v_0(\theta). \tag{43}$$

To obtain the solution of the E-L equation (42), we assume that the density $p(\theta,t)$ is a small harmonic perturbation of the uniform density. In particular, we express $p(\theta,t)$ as:

$$p(\theta,t) = p_0 + \varepsilon\tilde{p}(\theta,t), \tag{44}$$

where ε is a small perturbation parameter. Since $p(\theta,t)$ is a density, $\int_0^{2\pi}\tilde{p}(\theta,t)\, d\theta = 0$.

We are interested in obtaining a solution of the form:

$$v(\theta,t) = v_0(\theta) + \varepsilon\tilde{v}(\theta,t). \tag{45}$$

On substituting the ansatz (44) and (45) in (42), and retaining only $O(\varepsilon)$ term, we obtain the following linearized equation:

$$\partial_{\theta\theta}\tilde{v} = -2\pi\partial_\theta[(\partial_\theta\tilde{p})v_0]. \tag{46}$$

The linearized E-L equation (46) can be solved easily by considering a Fourier series expansion of $\varepsilon\tilde{p}(\theta,t)$:

$$\varepsilon\tilde{p}(\theta,t) = P_c(t)\cos\theta + P_s(t)\sin\theta + \text{h.o.h}, \tag{47}$$

where "h.o.h" denotes the terms due to higher order harmonics. The Fourier coefficients are given by,

$$P_c(t) = \frac{1}{\pi}\int_0^{2\pi} p(\theta,t)\cos\theta\, d\theta,\; P_s(t) = \frac{1}{\pi}\int_0^{2\pi} p(\theta,t)\sin\theta\, d\theta.$$

For a harmonic perturbation, the solution of the linearized E-L equation (46) is given by:

$$\varepsilon\tilde{v}(\theta,t) = \frac{\pi}{4\sigma_W^2}\left(P_c(t)\sin 2\theta - P_s(t)\cos 2\theta\right) =: v_1(\theta; P_c(t), P_s(t)) \tag{48}$$

For "h.o.h" terms in the Fourier series expansion (47) of the density in $p(\theta,t)$, the linearized E-L equation (46) can be solved in a similar manner. In numerical simulation provided here, we ignore the higher order harmonics, and use a control input as summarized in the following proposition:

Proposition 2: Consider the E-L equation (42) where the density $p(\theta,t)$ is assumed to be a small harmonic perturbation of the uniform density $\frac{1}{2\pi}$. As $\varepsilon \to 0$, the optimal control function is given by the following asymptotic formula:

$$v(\theta,t) = v_0(\theta) + v_1(\theta; P_c(t), P_s(t)), \tag{49}$$

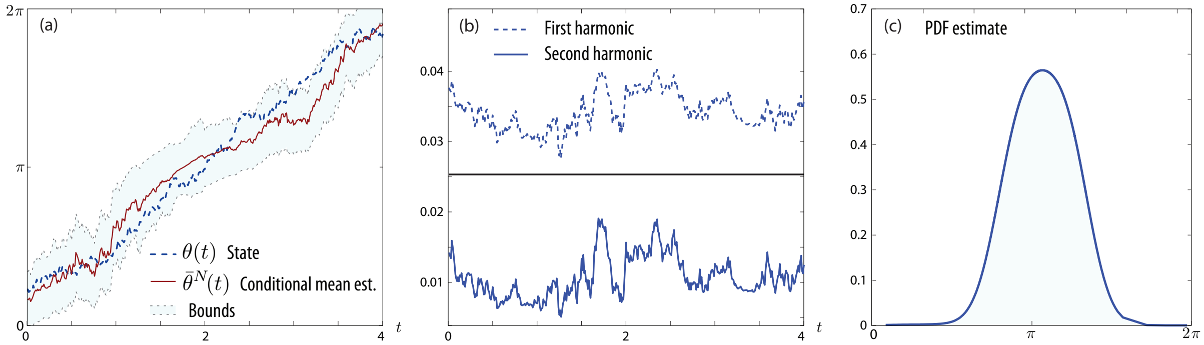

Fig. 2. Summary of the numerical experiments with the nonlinear oscillator filter: (a) Comparison of the true state $\{\theta_t\}$ and the conditional mean $\{\bar{\theta}_t^N\}$. (b) The mean-squared estimate of the first and second harmonics of the density $p(\theta,t)$ and (c) a plot of the typical density.

where $P_c(t), P_s(t)$ denote the harmonic coefficients of density $p(\theta,t)$. For large N, these are approximated by using the formulae:

$$\bar{P}_c^N(t) = \frac{1}{\pi N}\sum_{j=1}^{N}\cos\theta_j(t), \quad \bar{P}_s^N(t) = \frac{1}{\pi N}\sum_{j=1}^{N}\sin\theta_j(t). \quad (50)$$

■

We next discuss the result of numerical experiments. The particle filter model is given by (41) with gain function $v(\theta_t^i,t)$, obtained using formula (49). The number of particles $N = 10{,}000$ and their initial condition $\{\theta_0^i\}_{i=1}^N$ was sampled from a uniform distribution on circle $[0,2\pi]$.

Figure 2 summarizes some of the results of the numerical simulation. For illustration purposes, we depict only a single cycle from a time-window after transients due to initial condition have converged. Part (a) of the figure compares the sample path of the actual state $\{\theta_t\}$ (as a dashed line) with the estimated mean $\{\bar{\theta}_t^{(N)}\}$ (as a solid line). The shaded area indicates $\pm$ one standard deviation bounds. Part (b) of the figure provides a comparison of the magnitude of the first and the second harmonics (as dashed and solid lines, respectively) of the density $p(\theta,t)$. The density at any time instant during the time-window is approximately harmonic (see also part (c) where the density at one typical time instant is shown).

Note that at each time instant t, the estimated mean, the bounds and the density $p(\theta,t)$ shown here are all approximated from the ensemble $\{\theta_t^i\}_{i=1}^N$. For the sake of illustration, we have used a Gaussian mixture approximation to construct a smooth approximation of the density.

V. CONCLUSION AND FUTURE WORK

We have shown how a version of the particle filter can be derived using optimal control techniques. A significant advantage of the control-oriented formulation is that it provides self-correcting feedback mechanism to stabilize the particles around the common posterior $p(\cdot)$ for $\{X^i\}$.

There are several connections to be made to both the mean-field control literature as well as the filter stability literature:

(i) The background to this paper is the literature on mean-field games (e.g., [8], [10], [11]). The message here is that one can potentially improve particle filters by using (a small amount of) global information regarding particles. In the linear Gaussian case, this information is the mean and the variance of the ensemble. In the nonlinear oscillator case, this information is the first Fourier harmonic.

(ii) The research on ergodic properties of filters (e.g., [4]) is potentially relevant to the resolution of the following inter-related questions: Does the posterior approximation match the true posterior in the infinite-N limit? What are the stability properties of the algorithm in this limit? And what is the nature of performance bounds for large but finite values of N? Although some of these issues are addressed here for the linear Gaussian case, a resolution of these questions for the nonlinear case is a subject of ongoing research.

APPENDIX

A. Calculation of KL divergence

Recall the definition of K-L divergence for densities,

$$\mathrm{KL}\left(p_n \| \hat{p}_n^*\right) = \int_{\mathbb{R}} p_n(s)\ln\left(\frac{p_n(s)}{\hat{p}_n^*(s)}\right)\mathrm{d}s.$$

We make a co-ordinate transformation $s = x + u(x)$ and use (10) to express the K-L divergence as:

$$\mathrm{KL}\left(p_n \| \hat{p}_n^*\right) = \int_{\mathbb{R}} \frac{p_n^-(x)}{|1+u'(x)|}\ln\Big(\frac{p_n^-(x)}{|1+u'(x)|\,\hat{p}_n^*(x+u(x)|Z_{t_n})}\Big)|1+u'(x)|\,\mathrm{d}x$$

The expression for K-L divergence given in (12) follows on using (9).

B. Solution of the optimization problem

Denote:

$$\mathscr{L}(x,u,u') = -p_n^-(x)\Big(\ln|1+u'(x)| + \ln\big(p_n^-(x+u)\,p_{Z|X}(Z_{t_n}|x+u)\big)\Big). \quad (51)$$

The optimization problem (11) is a calculus of variation problem:

$$\min_u \int \mathscr{L}(x,u,u')\,\mathrm{d}x.$$

The minimizer is obtained via the analysis of first variation given by the well-known Euler-Lagrange equation:

$$\frac{\partial \mathscr{L}}{\partial u} = \frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{\partial \mathscr{L}}{\partial u'}\right),$$

Explicitly substituting the expression (51) for $\mathscr{L}$, we obtain (13).

C. Proof of Lemma 1

Denote $\Sigma_{n^-} := \Sigma_{t_n^-}, \mu_{n^-} := \mu_{t_n^-}$. Under the linear Gaussian assumption, we have:

$$p_n^-(x) = \frac{1}{\sqrt{2\pi\Sigma_{n^-}}}\exp[-\frac{(x-\mu_{n^-})^2}{2\Sigma_{n^-}}], \tag{52}$$

$$p_{Z|X}(Z_{t_n}|x+u_n) = \frac{1}{\sqrt{2\pi\sigma_W^2}}\exp[-\frac{(Z_{t_n}-h(x+u_n))^2}{2\sigma_W^2}]. \tag{53}$$

We consider a linear ansatz for the solution $u_n(x)$:

$$u_n(x) = k_n x + c_n. \tag{54}$$

So $u_n' = k_n$. Substituting (52) and (53) into (13) and using ansatz (54), we obtain:

$$\frac{1}{(1+k_n)\Sigma_{n^-}}x - \frac{1}{(1+k_n)\Sigma_{n^-}}\mu_{n^-} = (\frac{1}{\Sigma_{n^-}} + \frac{h^2}{\sigma_W^2})(1+k_n)x + (\frac{1}{\Sigma_{n^-}} + \frac{h^2}{\sigma_W^2})c_n - \frac{h}{\sigma_W^2}Z_{t_n} - \frac{\mu_{n^-}}{\Sigma_{n^-}}.$$

This gives:

$$\frac{1}{(1+k_n)\Sigma_{n^-}} = (\frac{1}{\Sigma_{n^-}} + \frac{h^2}{\sigma_W^2})(1+k_n),$$

$$(\frac{1}{\Sigma_{n^-}} + \frac{h^2}{\sigma_W^2})c_n = (\frac{1}{\Sigma_{n^-}} - \frac{1}{(1+k_n)\Sigma_{n^-}})\mu_{n^-} + \frac{h}{\sigma_W^2}Z_{t_n},$$

whose solution is given by (19) and (20).

D. Calculations for the discrete-time case

In this section, we consider the discrete time case:

$$X_n = a\,X_{n-1} + B_n, \tag{55}$$

$$Z_n = h\,X_n + W_n, \tag{56}$$

where a, h are real numbers, and B_n, W_n are mutually independent i.i.d noise with distribution $N(0,\sigma_B^2)$ and $N(0,\sigma_W^2)$, respectively.

The discrete-time filter is a special case of the continuous-discrete filter discussed in Section II. One can view (55) as the discrete time system obtained by integrating the continuous time SDE (16) forward in time from t_{n-1} to t_n.

Using Lemma 1, the discrete-time particle filter is given by:

$$t \in [t_n - 1, t_n): \quad X_{n^-}^i = a\,X_{n-1}^i + B_n^i; \tag{57}$$

$$t = t_n: \quad X_n^i = X_{n^-}^i + k_n X_{n^-}^i + c_n, \tag{58}$$

where $B_n^i \sim N(0,\sigma_B^2)$ and k_n, c_n are defined in (19) and (20). The mean $\mu_{t_n^-}$ and variance $\Sigma_{t_n^-}$ is approximated from the ensemble $\{X_{n^-}^i\}_{i=1}^N$ in practice. In the infinite-N limit, one obtains the recursive equations for mean and variance associated with the discrete-time Kalman filter:

$$\mu_{n+1} = a\,\mu_n + K_{n+1}(Z_{n+1} - h\,a\,\mu_n), \tag{59}$$

$$\Sigma_{n+1} = (1 - h\,K_{n+1})(a^2\,\Sigma_n + \sigma_B^2). \tag{60}$$

where $K_{n+1} = \dfrac{h(a^2\Sigma_n + \sigma_B^2)}{h^2(a^2\Sigma_n + \sigma_B^2) + \sigma_W^2}$ is the Kalman gain.

E. Proof of Theorem 1

The Gaussian density is given by:

$$p(x,t) = \frac{1}{\sqrt{2\pi\Sigma_t}}\exp(-\frac{(x-\mu_t)^2}{2\Sigma_t}), \tag{61}$$

The density (61) is a function of the stochastic process μ_t. Using Itô's formula,

$$\mathrm{d}p(x,t) = \frac{\partial p}{\partial \mu}\,\mathrm{d}\mu_t + \frac{\partial p}{\partial \Sigma}\,\mathrm{d}\Sigma_t + \frac{1}{2}\frac{\partial^2 p}{\partial \mu^2}\,\mathrm{d}\mu_t^2,$$

where $\frac{\partial p}{\partial \mu} = \frac{x-\mu_t}{\Sigma_t}p$, $\frac{\partial p}{\partial \Sigma} = \frac{1}{2\Sigma_t}\left(\frac{(x-\mu_t)^2}{\Sigma_t} - 1\right)p$, and $\frac{\partial^2 p}{\partial \mu^2} = \frac{1}{\Sigma_t}\left(\frac{(x-\mu_t)^2}{\Sigma_t} - 1\right)p$. Substituting these into the forward equation (27), we obtain a quadratic equation $Ax^2 + Bx = 0$, where

$$A = \mathrm{d}\Sigma_t - \left(2a\Sigma_t + \sigma_B^2 - \frac{h^2\Sigma_t^2}{\sigma_W^2}\right)\mathrm{d}t,$$

$$B = \mathrm{d}\mu_t - \left(a\mu_t\,\mathrm{d}t + \frac{h\Sigma_t}{\sigma_W^2}(\mathrm{d}Z_t - h\mu_t\,\mathrm{d}t)\right).$$

This leads to the model (32) and (33). ■

REFERENCES

[1] C. Andrieu and A. Doucet. Particle filtering for partially observed Gaussian state space models. *J. R. Statist. Soc. B*, 64:827–836, 2000.

[2] A. Budhiraja, L. Chen, and C. Lee. A survey of numerical methods for nonlinear filtering problems. *Physica D: Nonlinear Phenomena*, 230(1-2):27 – 36, 2007.

[3] R. Chen and J. S. Liu. Mixture Kalman filters. *J. R. Statist. Soc. B*, 62:493–508, 2000.

[4] P. Chigansky, R. Liptser, and R. Van Handel. Intrinsic methods in filter stability. In *Handbook of Nonlinear Filtering*. Oxford University Press, 2008. To appear.

[5] D. Crisan and A. Doucet. A survey of convergence results on particle filtering methods for practitioners. *IEEE Trans. Signal Process.*, 50(3):736–746, 2002.

[6] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte-Carlo Methods in Practice*. Springer-Verlag, April 2001.

[7] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F Radar and Signal Processing*, 140(2):107–113, 1993.

[8] M. Huang, P. E. Caines, and R. P. Malhame. Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε-Nash equilibria. *IEEE Trans. Automat. Control*, 52(9):1560–1571, 2007.

[9] G. Terejanu, P. Singla, T. Singh, and P. D. Scott. Uncertainty propagation for nonlinear dynamical systems using Gaussian mixture models. *AIAA Journal of Guidance, Control and Dynamics*, 31(6):1623–1633, Nov. 2008.

[10] G. Y. Weintraub, L. Benkard, and B. Van Roy. Oblivious equilibrium: A mean field approximation for large-scale dynamic games. In *Advances in Neural Information Processing Systems*, volume 18. MIT Press, 2006.

[11] H. Yin, P. G. Mehta, S. P. Meyn, and U. V. Shanbhag. Synchronization of coupled oscillators is a game. *In American Control Conference*, pages 1783–1790, June 2010.