2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Feedback Particle Filter with Mean-field Coupling

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. 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: This is achieved by introducing a cost function, defined by 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, and is shown to admit an innovation error-based feedback structure. For diffusions with continuous observations, the value of the optimal control solution is ideal: The two posteriors match exactly, provided they are initialized with identical priors. The resulting control system is called the *feedback particle filter*.

An algorithm is introduced and implemented in two numerical examples. A numerical comparison of the feedback particle filter with the bootstrap particle filter is provided.

I. INTRODUCTION

We consider a scalar filtering problem:

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

$$\mathrm{d}Z_t = h(X_t)\,\mathrm{d}t + \sigma_W\,\mathrm{d}W_t, \tag{1b}$$

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 estimate 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, the solution is given by the finite-dimensional Kalman filter. The theory of nonlinear filtering is described in the classic monograph [7]. The filter is infinite dimensional since it defines the evolution, in the space of probability measures, of $\{p^*(\cdot,t) : t \ge 0\}$.

The article [2] surveys numerical methods to approximate the nonlinear filter. One approach described in this survey is particle filtering.

The particle filter is a simulation-based algorithm to approximate the filtering task [3]. The key step is the construction of N stochastic processes $\{X_t^i : 1 \le i \le N\}$: The value $X_t^i \in \mathbb{R}$ is the state for the i^{th} particle at time t. For each time t, the empirical distribution formed by, the "particle population" is used to approximate the posterior distribution. A common approach in particle filtering is called *sequential importance sampling*, where particles are generated according to their importance weight at every time step [1], [3].

The objective of this paper is to introduce an alternative feedback control-based approach to the construction of a particle filter for (1a)-(1b). The main result of this paper is to derive an explicit formula for the optimal control input. The optimally controlled dynamics of the i^{th} particle have the following gain feedback form,

$$\begin{aligned}\mathrm{d}X_t^i &= a(X_t^i)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i \\ &\quad + \mathsf{K}(X_t^i,t)\,\mathrm{d}I_t^i + \frac{1}{2}\sigma_W^2 \mathsf{K}(X_t^i,t)\mathsf{K}'(X_t^i,t)\,\mathrm{d}t\end{aligned} \tag{2}$$

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

$$\mathrm{d}I_t^i := \mathrm{d}Z_t - \frac{1}{2}(h(X_t^i) + \hat{h}_t)\,\mathrm{d}t \tag{3}$$

where $\hat{h}_t := \mathsf{E}[h(X_t^i)|\mathscr{Z}_t]$. In a numerical implementation, we approximate $\hat{h}_t \approx \frac{1}{N}\sum_{i=1}^N h(X_t^i) =: \hat{h}_t^{(N)}$.

The gain function K is shown to be the solution to the Euler-Lagrange boundary value problem (E-L BVP):

$$-\frac{\partial}{\partial x}\left(\frac{1}{p(x,t)}\frac{\partial}{\partial x}\{p(x,t)\mathsf{K}(x,t)\}\right) = \frac{1}{\sigma_W^2}h'(x), \tag{4}$$

with boundary conditions $\lim_{x\to\pm\infty} p(x,t)\mathsf{K}(x,t) = 0$, where p denotes the conditional distribution of X_t^i given $\mathscr{Z}_t$ and $h'(x) = \frac{\mathrm{d}}{\mathrm{d}x}h(x)$. Note that the gain function needs to be obtained for each value of time t.

The controlled system (2)-(4) is called the *feedback particle filter*.

The contributions of this paper are as follows:

- **Consistency.** We show that the feedback particle filter (2) is consistent with nonlinear filter in the following sense: Suppose at time 0, $p(x,0) = p^*(x,0)$ and the gain function $\mathsf{K}(x,t)$ is obtained as the solution to (4). Then, for all $t > 0$,

$$p(x,t) = p^*(x,t).$$

This implies that if the initial conditions $\{X_0^i\}_{i=1}^N$ are drawn from initial distribution $p^*(x,0)$ of X_0, then, as $N \to \infty$, the empirical distribution of the particle system approximates the posterior distribution p^*.

- **Algorithms.** We propose algorithms for synthesis of the gain function $\mathsf{K}(x,t)$. If $a(\cdot)$ and $h(\cdot)$ are linear and the density p is Gaussian, then the gain function admits a closed-form solution in terms of variance alone. The variance is approximated empirically as a sample covariance.

In the nonlinear case, we approximate the density as a sum of Gaussian and provide exact and approximate formulae

Financial support from NSF grants EECS-0925534 and CCF-08-30776, and the AFOSR grant FA9550-09-1-0190 is gratefully acknowledged.

T. Yang, P. Mehta and S. Meyn are with Coordinated Science Laboratory at the University of Illinois at Urbana-Champaign taoyang1@illinois.edu; mehtapg@illinois.edu; meyn@illinois.edu

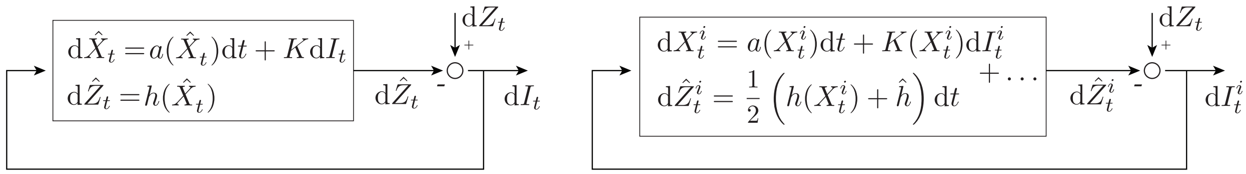

(a): Kalman filter (b): Feedback particle filter

Fig. 1. Innovation error-based feedback structure for (a) Kalman filter and (b) nonlinear feedback particle filter.

for the solution of the gain function. With the sum of Gaussian approximation, each Gaussian models a sub-cluster of particles $\{X^i_t\}$.

In recent decades, there have been many important advances in importance sampling based approaches for particle filtering; cf., [1], [3], [2]. The crucial distinction here is that there is no resampling of particles. We believe that the introduction of control in the feedback particle filter has several useful features/advantages:

Innovation error. The innovation error-based feedback structure is a key feature of the feedback particle filter (2). The innovation error is now based on the average value of the prediction $h(X^i_t)$ of the i^{th}-particle and the prediction $\hat{h}_t$ due to the entire population.

The feedback particle filter thus provides for a generalization of the Kalman filter to nonlinear systems, where the innovation error-based feedback structure of the control is preserved (see Fig. 1). For the linear case, the optimal gain function is the Kalman gain. For the nonlinear case, the Kalman gain is replaced by a nonlinear function of the state.

Variance reduction. We believe that the feedback can help reduce the high variance that is sometimes observed in the usual particle filter. Numerical results in Sec V support this claim — See Fig. 3 for a comparison of the feedback particle filter and the bootstrap filter for the linear filtering problem.

Applications. Bayesian inference is an important paradigm used to model functions of certain neural circuits in brain [4]. Compared to techniques that rely on importance sampling, a feedback particle filter may provide a more neurobiologically plausible model to implement filtering and inference functions [8].

The biggest limitation of our approach is the need to solve the BVP at each time-step, that additionally requires one to approximate the density. We are encouraged however by the extensive set of tools in feedback control: After all, one rarely needs to solve the HJB equations in closed-form to obtain a reasonable feedback control law. Moreover, there are many approaches in nonlinear and adaptive control to both approximate control laws as well as learn/adapt these in online fashion. For feedback particle filter, this is a subject of continuing research.

The outline of this paper is as follows. In Sec. II, we begin by reviewing the results of our earlier paper [8] that introduced the continuous-discrete time filtering problem. The nonlinear filter is introduced in III, and algorithms discussed in IV. We conclude with a discussion of some numerical results in Sec. V.

II. Preliminaries

In this section we briefly summarize the main results of our earlier paper [8] that dealt with the continuous-discrete time filtering problem.

For the continuous-discrete time filtering problem, the equation for dynamics is given by (1a), 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 $N(0, \sigma^2_W)$. We denote $\mathscr{Z}^{\blacksquare}_n := \sigma\{Z_{t_k} : k \le n\}$.

The particle model in this case is a hybrid dynamical system: At time t_n, assuming that $X^i_{t_{n-1}}$ is given, the i^{th} particle evolves on $[t_{n-1}, t_n)$ according to the stochastic differential equation

$$\mathrm{d}X^i_t = a(X^i_t)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B^i_t, \quad t_{n-1} \le t < t_n. \tag{6}$$

At the end of this time-horizon, there is a potential jump that is determined by the control input $U^i_{t_n}$:

$$X^i_{t_n} = X^i_{t_n^-} + U^i_{t_n}, \tag{7}$$

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

A. Belief maps

For each n we denote:

1) p^*_n and p^{*-}_n: The conditional distribution of X_{t_n} given $\mathscr{Z}^{\blacksquare}_n$ and $\mathscr{Z}^{\blacksquare}_{n-1}$, respectively.
2) p_n and p^-_n: The conditional distribution of $X^i_{t_n}$ given $\mathscr{Z}^{\blacksquare}_n$ and $\mathscr{Z}^{\blacksquare}_{n-1}$, respectively.

These densities evolve according to recursions 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: 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 forward equation associated with the diffusion on $[t_{n-1}, t_n)$.

The second part of the mapping is different: The transformation that takes p^{*-}_n to p^*_n is obtained from Bayes' rule. The transformation that takes p^-_n to p_n depends upon the choice of control $U^i_{t_n}$ in (7).

In [8], we seek a control input $U^i_{t_n}$ that is *admissible*.

Definition 1 (Admissible Input): The control sequence $\{U^i_{t_n} : n \ge 0\}$ is *admissible* if there is a sequence of maps $\{u_n(x; z^n_0)\}$ such that $U^i_{t_n} = u_n(X^i_{t_n}, Z_{t_0}, \ldots, Z_{t_n})$, and,

(i) $\mathsf{E}[|U^i_{t_n}|] < \infty$, and with probability one,

$$\lim_{x\to\pm\infty} u_n(x, Z_{t_0}, \ldots, Z_{t_n}) p_n^-(x) = 0.$$

(ii) u_n is twice continuously differentiable as a function of x.

(iii) $1 + u_n'(x)$ is non-zero for all x, where $u'(x) = \frac{\mathrm{d}}{\mathrm{d}x} u(x)$.

We suppress the dependency on the observations, writing $U^i_{t_n} = u(x)$ when $X^i_{t_n^-} = x$.

B. Variational Problem

Our goal is to choose an admissible input u so that the mapping $\mathscr{P}$ approximates the mapping $\mathscr{P}^*$ in (8). In particular, we assume $p_n^- = p_n^{*-}$ and seek

$$u_n(x) = \arg\min_u \mathrm{KL}\left(p_n \| p_n^*\right), \tag{9}$$

where $\mathrm{KL}(\cdot)$ denote the Kullback-Leibler (KL) divergence between the two distributions. The KL divergence can be expressed as

$$\begin{aligned}\mathrm{KL}(p_n \| p_n^*) = C - \int_{\mathbb{R}} p_n^-(x) \Big\{ \ln|1 + u'(x)| \\ + \ln(p_n^-(x + u(x)) p_{Z|X}(Z_{t_n} | x + u(x))) \Big\} \mathrm{d}x\end{aligned} \tag{10}$$

where $p_{Z|X}(Z_{t_n}|\cdot) = \frac{1}{\sqrt{2\pi\sigma_W^2}} \exp\left(-\frac{(Z_{t_n} - h(\cdot))^2}{2\sigma_W^2}\right)$, and $C = \int_{\mathbb{R}} p_n^-(x) \ln(p_n^-(x) p_Z(Z_{t_n}))\, \mathrm{d}x$ is a constant that does not depend on u; see [8] for the calculation.

The solution to (9) is described in the following:

Proposition 1 (Proposition 1 in [8]): Suppose that the admissible function u is a minimizer for the optimization problem (9). Then it is a solution of the following Euler-Lagrange (E-L) BVP:

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{p_n^-(x)}{|1 + u'(x)|}\right) = p_n^-(x) \frac{\partial}{\partial u} \left(\ln(p_n^-(x+u) p_{Z|X}(Z_{t_n}|x+u)) \right), \tag{11}$$

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

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

III. Continuous-time filtering

Consider now the continuous time filtering problem (1a, 1b) introduced in Section I.

We denote as $p^*(x,t)$ the conditional distribution of X_t given $\mathscr{Z}_t = \sigma(Z_s : s \le t)$. The evolution of $p^*(x,t)$ is described by the Kushner-Stratonovich (K-S) equation:

$$\mathrm{d}p^* = \mathscr{L}^\dagger p^* \, \mathrm{d}t + \frac{1}{\sigma_W^2}(h - \hat{h}_t)(\mathrm{d}Z_t - \hat{h}_t \, \mathrm{d}t) p^*, \tag{12}$$

where $\hat{h}_t = \int h(x) p^*(x,t)\, \mathrm{d}x$ and $\mathscr{L}^\dagger p^* = -\frac{\partial(p^* a)}{\partial x} + \frac{\sigma_B^2}{2} \frac{\partial^2 p^*}{\partial x^2}$.

A. Belief state dynamics & control architecture

The model for the particle filter is given by,

$$\mathrm{d}X^i_t = a(X^i_t) + \sigma_B \, \mathrm{d}B^i_t + \underbrace{u(X^i_t, t)\, \mathrm{d}t + \mathsf{K}(X^i_t, t)\, \mathrm{d}Z_t}_{\mathrm{d}U^i_t}, \tag{13}$$

where $X^i_t \in \mathbb{R}$ is the state for the i^{th} particle at time t, and $\{B^i_t\}$ are mutually independent standard Wiener processes. We assume the initial conditions $\{X^i_0\}_{i=1}^N$ are i.i.d., independent of $\{B^i_t\}$, and drawn from the initial distribution $p^*(x,0)$ of X_0. Both $\{B^i_t\}$ and $\{X^i_0\}$ are also assumed to be independent of X_t, Z_t.

As in Sec II, we impose admissibility requirements on the control input U^i_t in (13):

Definition 2 (Admissible Input): The control input U^i_t is *admissible* if the random variables $u(x,t)$ and $\mathsf{K}(x,t)$ are $\mathscr{Z}_t = \sigma(Z_s : s \le t)$ measurable for each t. Moreover, each t,

(i) $\mathsf{E}[|u(X^i_t,t)| + |\mathsf{K}(X^i_t,t)|^2] < \infty$, and with probability one,

$$\lim_{x\to\pm\infty} u(x,t) p(x,t) = 0, \tag{14a}$$

$$\lim_{x\to\pm\infty} \mathsf{K}(x,t) p(x,t) = 0. \tag{14b}$$

where p is the posterior distribution of X^i_t given $\mathscr{Z}_t$.

(ii) $u : \mathbb{R}^2 \to \mathbb{R}$, $\mathsf{K} : \mathbb{R}^2 \to \mathbb{R}$ are twice continuously differentiable in their first arguments.

■

Recall that are two types of conditional distributions of interest in our analysis:

1) $p(x,t)$: Defines the conditional dist. of X^i_t given $\mathscr{Z}_t$.

2) $p^*(x,t)$: Defines the conditional dist. of X_t given $\mathscr{Z}_t$.

The functions $\{u(x,t), \mathsf{K}(x,t)\}$ represent the continuous time-counterparts of the optimal control function $u_n(x)$ (see (9)). We say that these functions are *optimal* if $p \equiv p^*$. That is, given $p^*(\cdot,0) = p(\cdot,0)$, our goal is to choose $\{u, \mathsf{K}\}$ in the feedback particle filter so that the evolution equations of these conditional distributions coincide (see (12) and (15)).

The evolution equation for the belief state is described in the next result. Its proof appears in Appendix A.

Proposition 2: Consider the process X^i_t that evolves according to the particle filter model (13). The conditional distribution of X^i_t given the filtration $\mathscr{Z}_t$, $p(x,t)$, satisfies the forward equation

$$\begin{aligned}\mathrm{d}p = \mathscr{L}^\dagger p \, \mathrm{d}t - \frac{\partial}{\partial x}(\mathsf{K}p)\, \mathrm{d}Z_t \\ - \frac{\partial}{\partial x}(up)\, \mathrm{d}t + \sigma_W^2 \frac{1}{2} \frac{\partial^2}{\partial x^2}\left(p\mathsf{K}^2\right) \mathrm{d}t.\end{aligned} \tag{15}$$

■

B. Consistency with the nonlinear filter

The main result of this section is the construction of an optimal pair $\{u, \mathsf{K}\}$ under the following general assumption:

Assumption A1 The conditional distributions (p^*, p) are C^2, with $p^*(x,t) > 0$ and $p(x,t) > 0$, for all $x \in \mathbb{R}$, $t > 0$.

■

We henceforth choose $\{u, \mathsf{K}\}$ as the solution to a certain E-L BVP based on p: The function K as the solution to

$$-\frac{\partial}{\partial x}\left(\frac{1}{p(x,t)} \frac{\partial}{\partial x}\{p(x,t)\mathsf{K}(x,t)\}\right) = \frac{1}{\sigma_W^2} h'(x), \tag{16}$$

with boundary condition (14b). The function $u(\cdot,t) : \mathbb{R} \to \mathbb{R}$ is obtained as:

$$u(x,t) = \mathsf{K}(x,t)\left(-\frac{1}{2}(h(x) + \hat{h}_t) + \frac{1}{2}\sigma_W^2 \mathsf{K}'(x,t)\right), \tag{17}$$

where $\hat{h}_t = \int h(x)p(x,t)\,\mathrm{d}x$. We assume moreover that the boundary conditions given in (14a) also hold, so that $\{u, \mathsf{K}\}$ is admissible. The BVP is motivated by the continuous-time limit of (11), obtained on letting $t_{n+1} - t_n$ go to zero; the calculations appear in Appendix B.

Existence and uniqueness of $\{u, \mathsf{K}\}$ is obtained in the following proposition — Its proof is given in Appendix B.

Proposition 3: Consider the BVP (16), subject to Assumption A1. Then,

1) There exists a unique solution K, subject to the boundary condition (14b).
2) The solution satisfies $\mathsf{K}(x,t) \geq 0$ for all x, t, provided $h'(x) \geq 0$ for all x. ∎

The following theorem shows that the two evolution equations (12) and (15) are identical. The proof appears in Appendix C.

Theorem 1: Consider the two evolution equations for p and p^*, defined according to the solution of the forward equation (15) and the K-S equation (12), respectively. Suppose that the control functions $u(x,t)$ and $\mathsf{K}(x,t)$ are obtained according to (16) and (17), respectively. Then, provided $p(x,0) = p^*(x,0)$, we have for all $t \geq 0$,

$$p(x,t) = p^*(x,t)$$

∎

C. Linear Gaussian case

We provide here a special case for linear system:

$$\mathrm{d}X_t = \alpha\, X_t\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t, \tag{18a}$$

$$\mathrm{d}Z_t = \gamma\, X_t\,\mathrm{d}t + \sigma_W\,\mathrm{d}W_t, \tag{18b}$$

where α, γ are real numbers. We assume 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 functions $u(x,t), \mathsf{K}(x,t)$ in the linear Gaussian case.

Lemma 1: Consider the linear observation equation (18b). 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 (4) is given by:

$$\mathsf{K}(x,t) = \frac{\gamma\Sigma_t}{\sigma_W^2}, \quad u(x,t) = -\frac{\gamma^2\Sigma_t}{2\sigma_W^2}(x+\mu_t) \tag{19}$$

∎

The formulae (19) are verified by direct substitution in the ODE (4) where the distribution p is Gaussian.

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

$$dX_t^i = \alpha\, X_t^i\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i + \frac{\gamma\Sigma_t}{\sigma_W^2}\left(\mathrm{d}Z_t - \gamma\frac{X_t^i + \mu_t}{2}\,\mathrm{d}t\right). \tag{20}$$

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 2: Consider the linear Gaussian filtering problem defined by the state-observation equations (18a,18b). In this case the posterior distributions of X and X^i are Gaussian, whose conditional mean and covariance are given by the respective SDE and the ODE,

$$\mathrm{d}\mu_t = \alpha\mu_t\,\mathrm{d}t + \frac{\gamma\Sigma_t}{\sigma_W^2}\Big(\mathrm{d}Z_t - \gamma\mu_t\,\mathrm{d}t\Big)$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\Sigma_t = 2\alpha\Sigma_t + \sigma_B^2 - \frac{\gamma^2\Sigma_t^2}{\sigma_W^2}$$

∎

The result is verified by substituting $p(x,t) = \frac{1}{\sqrt{2\pi\Sigma_t}}\exp(-\frac{(x-\mu_t)^2}{2\Sigma_t})$ in the forward equation (15). The details are omitted on account of space, and because the result is a special case of Theorem 1.

Notice that particle system (20) is not practical since it requires computation of the conditional mean and variance $\{\mu_t, \Sigma_t\}$. In practice $\{\mu_t, \Sigma_t\}$ are approximated as sample means and sample covariances from the ensemble $\{X_t^i\}_{i=1}^N$:

$$\mu_t \approx \mu_t^{(N)} := \frac{1}{N}\sum_{i=1}^N X_t^i,$$

$$\Sigma_t \approx \Sigma_t^{(N)} := \frac{1}{N-1}\sum_{i=1}^N (X_t^i - \mu_t^{(N)})^2.$$

The resulting equation (20) for the i^{th} particle is given by

$$\mathrm{d}X_t^i = \alpha\, X_t^i\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t^i + \frac{\gamma\Sigma_t^{(N)}}{\sigma_W^2}\left(\mathrm{d}Z_t - \gamma\frac{X_t^i + \mu_t^{(N)}}{2}\,\mathrm{d}t\right). \tag{21}$$

It is very similar to the mean-field "synchronization-type" control laws and oblivious equilibria constructions as in [6], [9]. As $N \to \infty$, the empirical distribution of the particle system approximates the posterior distribution $p^*(x,t)$ (by Theorem 2).

IV. Synthesis of the Gain Function $\mathsf{K}(x,t)$

Implementation of the nonlinear filter (2) requires solution of the E-L BVP (4) to obtain the gain function $\mathsf{K}(x,t)$ for each fixed t. If $p(x,t)$ is known, the linear BVP admits a closed-form solution (43) – the main issue thus is the approximation of the distribution $p(x,t)$.

In this section, we consider the following approximation:

Assumption A2 For each fixed t, the distribution $p(x,t)$ is a sum of Gaussian:

$$p(x,t) = \sum_{j=1}^m \lambda_t^j q_t^j(x), \tag{22}$$

where $q_t^j(x) = q(x;\mu_t^j,\Sigma_t^j) = \frac{1}{\sqrt{2\pi\Sigma_t^j}}\exp(-\frac{(x-\mu_t^j)^2}{2\Sigma_t^j})$, $\lambda_t^j > 0$, $\sum\lambda_t^j = 1$. We assume an ordering so $\mu_t^j < \mu_t^k$ for $j < k$. ∎

The approximation is motivated by the numerical algorithm: At each discrete time-step t, we have particle states $\{X_t^i\}_{i=1}^N$. We identify m-clusters each of which is assumed to be localized in $\mathbb{R}$. We approximate the j^{th}-cluster with a Gaussian pdf with weight $\lambda_t^j \in (0,1)$, empirical mean μ_t^j and variance Σ_t^j.

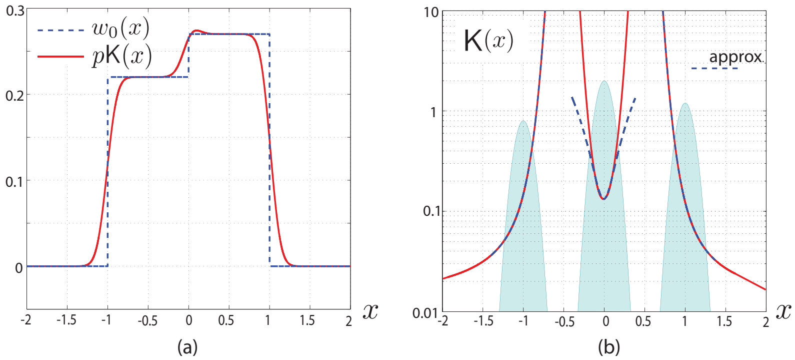

TABLE I

j	λ^j	μ^j
1	0.2	-1
2	0.5	0
3	0.3	1

TABLE II: Taylor series coefficients

v_0	$\Sigma^j + \sqrt{2\pi\Sigma^j}\left(\frac{1}{2}(\hat{h}-\mu^j) + \sum_{k<j}\frac{\lambda^k}{\lambda^j}(\hat{h}-\mu^k)\right)$
v_1	$(\hat{h}-\mu^j)$
v_2	$\sqrt{\frac{2\pi}{\Sigma^j}}\left(\frac{1}{2}(\hat{h}-\mu^j) + \sum_{k<j}\frac{\lambda^k}{\lambda^j}(\hat{h}-\mu^k)\right)$

(c)

Fig. 2. (a) Comparison of the solutions (23) and (25). (b) Comparison of the exact (solid) and approximate (dashed) solutions. In the background, the three Gaussians are also shown. (c) For both parts (a) and (b) $p(x,t)$ is a sum of $(m=3)$ Gaussian densities with weights λ^j and means μ^j tabulated in Table I. Table II tabulates the formulae for the Taylor series coefficients in (26).

For the ease of presentation, we also assume $h(x) = x$ in the observation model (1b):

Assumption A3 We assume the observation model (1b) with $h(x) \equiv x$.

This assumption is not critical. Other modalities can also be considered as discussed in Remark 1 below.

The following proposition provides a closed-form solution of the E-L BVP with $p(x,t)$ of the form (22):

Proposition 4: Consider the BVP (4) with $h(x) = x$. Suppose $p(x,t)$ is of the form (22). Then the solution is given by

$$p\mathsf{K}(x,t) = \frac{1}{\sigma_W^2}\Big(\sum_{j=1}^{m}\lambda_t^j(\hat{h}_t - h(\mu_t^j))Q_t^j(x) + \sum_{j=1}^{m} h'(\mu_t^j)\lambda_t^j\Sigma_t^j q_t^j(x)\Big), \tag{23}$$

where $Q_t^j(x) = \int_{-\infty}^{x} q_t^j(y)\,\mathrm{d}y$. ■

The proof is omitted – it is a straightforward verification by direct substitution of the solution in the ODE (4).

Remark 1: For general h, the expression in (23) represents an approximate solution of the E-L BVP in the asymptotic limit that $\Sigma_t^j \to 0$ for $j = 1,\ldots,m$.

This is seen by considering a BVP

$$-\frac{\partial}{\partial x}\left(\frac{1}{p(x,t)}\frac{\partial w}{\partial x}\right) = \frac{1}{\sigma_W^2}h'(x), \tag{24}$$

with boundary condition $\lim_{x\to\pm\infty} w(x) = 0$. Consider also a limiting distribution $p^0(x,t) = \sum_{j=1}^{m}\lambda_t^j\delta(x-\mu_t^j)$. With $p(x,t) = p^0(x,t)$ in (24), the weak solution is given by the staircase function:

$$w(x) = w^0(x) := \frac{1}{\sigma_W^2}\sum_{j=1}^{m}\lambda_t^j(\hat{h}_t - h(\mu_t^j))H(x-\mu_t^j), \tag{25}$$

where $H(\cdot)$ is the Heaviside function.

For small values of Σ_t^j, the solution given by (23) is a small perturbation of $w^0(x)$ in (25); see Fig. 2(a) which depicts the two solutions for $h(x) = x$ and (λ_t^j, μ_t^j) tabulated in Table I and $\Sigma_t^j = 0.01$.

In the case of general h, the approximate nature of solution (23) follows by considering a perturbation argument in the asymptotic limit $\Sigma_t^j \to 0$. ■

Now given $w(x,t) = p\mathsf{K}(x,t)$, the gain function is obtained as $\mathsf{K}(x,t) = w(x,t)/p(x,t)$. Since dividing by p may be a problem, we also provide asymptotic formulae for values of x near the j^{th} empirical mean $(x \approx \mu_t^j)$:

1) For $2 \le j \le m-1$, the gain function may be approximated by using a Taylor series approximation: For $x \approx \mu_t^j$,

$$\mathsf{K}(x,t) = \frac{1}{\sigma_W^2}\left(v_0 + v_1(x-\mu_t^j) + v_2(x-\mu_t^j)^2\right), \tag{26}$$

where formulae for the coefficients appear in Table II.

2) For $i = 1$ or m, the Taylor series does not yield a good approximation and one can use the following formula:

$$\text{For } x \approx \mu_t^1 : \mathsf{K}(x,t) = \frac{1}{\sigma_W^2}\left(\Sigma_t^1 + (\hat{h}_t - \mu_t^1)\frac{Q_t^1}{q_t^1}(x,t)\right), \tag{27a}$$

$$\text{For } x \approx \mu_t^m : \mathsf{K}(x,t) = \frac{1}{\sigma_W^2}\left(\Sigma_t^m + (\hat{h}_t - \mu_t^m)\frac{Q_t^m - 1}{q_t^m}(x,t)\right). \tag{27b}$$

Using (27b), $\lim_{x\to-\infty}\mathsf{K}(x,t) = \frac{\Sigma_t^1}{\sigma_W^2}$, $\lim_{x\to\infty}\mathsf{K}(x,t) = \frac{\Sigma_t^m}{\sigma_W^2}$. So, the asymptotic value of the gain is consistent with the formula for the gain obtained in the linear case (see (19)).

Figure 2 (b) depicts a comparison between the exact numerical solution K and the approximate formulae; parameter values are the same as in Table I. We obtain the approximate solution for $x \in [\mu_t^j - 4\sqrt{\Sigma_t^j}, \mu_t^j + 4\sqrt{\Sigma_t^j}]$ – for $j = 2$ using (26), and for $j = 1$ and 3 using (27a) and (27b), respectively.

Remark 2: Depending upon the problem as well as the available computational resources, one may choose different approximate structures for the gain function. For example, using the Taylor series approximation, one possibility is to chose $\mathsf{K}(x,t) = \frac{1}{\sigma_W^2}v_0$, a constant for the j^{th} cluster. A better choice, however, may be to pick the constant to be the mean value taken over ± 1 standard deviation, i.e., $\mathsf{K}(x,t) = \frac{1}{\sigma_W^2}(v_0 + \frac{1}{3}v_2\Sigma_t^j)$.

A. Algorithm

At time t, we assume m clusters with empirical mean μ_t^j and variance Σ_t^j, ordered such that $\mu_t^j < \mu_t^k$ for $j < k$.

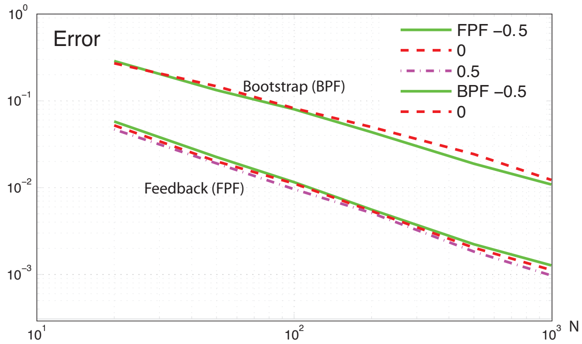

Fig. 3. Comparison of the *mse* using feedback particle filter and the bootstrap particle filter.

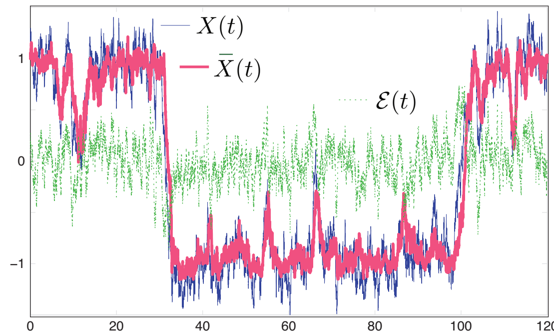

Fig. 4. Comparison of the true state $X(t)$ and the conditional mean $\bar{X}(t)$ by using feedback particle filter. The error $\mathscr{E}(t) = X(t) - \bar{X}(t)$ remains small even during a transition of the state.

We assign the i^{th}-particle to the j^{th}-cluster based on proximity ($j = \arg\min_k |X_t^i - \mu_t^k|$) and evaluate the gain function $\mathsf{K}(X_t^i, t)$ according to (27a,27b) or (26).

Now, one problem with directly implementing the particle filter model (2) is that there possibly are multiple time-scales in the problem: $\mathsf{K}(x,t)$ may become very large for particles that lie between two clusters (see Fig. 2(b)). So, $\mathrm{d}t$ would need to be chosen extremely small to avoid numerical instabilities. This is impractical.

To help deal with this issue, we propose to break up the trajectory into the following two phases:

(i) *Control phase.* If $\mathsf{K}(X_t^i,t) < C\mathsf{K}(\mu_t^i,t)$ then we use the particle filter model (2) to obtain $\mathrm{d}X_t^i$. The constant $C > 1$ is selected based on the available $\mathrm{d}t$.

(ii) *Flight phase.* If $\mathsf{K}(X_t^i,t) \geq C\mathsf{K}(\mu_t^i,t)$ then we set $X_t^i = \mu_t^j$ or $X_t^i = \mu_t^{j+1}$ depending upon the sign of the innovation term $\mathrm{d}I_t^i$.

Note that the flight phase implements motion during the fast time-scale: It allows the particle to escape the cluster without the need to make $\mathrm{d}t$ arbitrarily small.

V. Numerics

A. Linear case

We first provide a comparison between the feedback particle filter and the bootstrap particle filter for the linear problem (18a, 18b).

For the linear filtering problem, the optimal solution is given by the Kalman filter. We use this solution to define the relative mean-squared error:

$$mse = \frac{1}{T}\int_0^T \left(\frac{\Sigma_t^{(N)} - \Sigma_t}{\Sigma_t}\right)^2 \mathrm{d}t, \tag{28}$$

where Σ_t is the error covariance using the Kalman filter, and $\Sigma_t^{(N)}$ is its approximation using the particle filter.

Figure 3 depicts a comparison between *mse* obtained using the feedback particle filter (21) and the bootstrap filter. The latter implementation is based on an algorithm taken from Ch. 9 of [1]. For simulation purposes, we used a range of values of $\alpha \in [-1,1]$, $\gamma = 3$, $\sigma_B = 1$, $\sigma_W = 0.5$, $\mathrm{d}t = 0.01$, and $T = 50$. The plot in Fig. 3 is generated using simulations with $N = 20, 50, 100, 200, 500, 1000$ particles.

We refer the reader to our earlier paper [8] for additional simulation plots for the pdf and estimates obtained using the linear feedback particle filter (21).

B. Nonlinear example

We consider

$$\mathrm{d}X_t = X_t(1 - X_t^2)\,\mathrm{d}t + \sigma_B\,\mathrm{d}B_t, \tag{29a}$$
$$\mathrm{d}Z_t = X_t\,\mathrm{d}t + \sigma_W\,\mathrm{d}W_t, \tag{29b}$$

where $\sigma_B = 0.4$, $\sigma_W = 0.2$. Without noise, the ODE (29a) has two stable equilibria at ± 1. With noise, the state of the SDE "transitions" between these two "equilibria" (see Fig. 4).

Figure 4 depicts the simulation results obtained using the nonlinear feedback particle filter (2). The implementation is based on the algorithm presented in Sec IV-A with $\mathrm{d}t = 0.01$. The control gain $\mathsf{K}(x,t)$ is chosen via the constant approximation discussed in Remark 2. We initialize the simulation with $m = 2$ clusters at ± 1. After a brief period of transients, these clusters merge into a single cluster, which adequately tracks the true state including the transition events.

We refer the reader to our earlier paper [8] for another numerical example that uses the nonlinear feedback particle filter (2) for filtering of a nonlinear oscillator system.

Appendix

A. Derivation of the Forward Equation

We denote the filtration $\mathscr{B}_t = \sigma(X_0^i, B_s^i : s \leq t)$, and we recall that $\mathscr{Z}_t = \sigma(Z_s : s \leq t)$ for $t \geq 0$. These two filterations are *independent* by construction.

On denoting $\tilde{a}(x,t) = a(x) + u(x,t)$, the particle evolution (13) is expressed,

$$X_t^i = X_0^i + \int_0^t \tilde{a}(X_s^i, s)\,\mathrm{d}s + \int_0^t \mathsf{K}(X_s^i, s)\,\mathrm{d}Z(s) + \sigma_B B_t^i. \tag{30}$$

By assumption on Lipschitz continuity of $\tilde{a}$ and K, there exists a unique solution that is adapted to the larger filtration $\mathscr{B}_t \vee \mathscr{Z}_t = \sigma(X_0^i, B_s^i, Z_s : s \leq t)$. In fact, there is a functional F_t such that

$$X_t^i = F_t(X_0^i, B_t^i, Z^t), \tag{31}$$

where $Z^t := \{Z_s : 0 \leq s \leq t\}$ denotes the trajectory.

The conditional distribution of X_t^i given $\mathscr{Z}_t = \sigma(Z_s : s \leq t)$ was introduced in Sec. II-A: Its density is denoted $p(x,t)$, defined by any bounded and measurable function $f\colon \mathbb{R}\to\mathbb{R}$ via,

$$\mathsf{E}(f(X_t^i)\mid \mathscr{Z}_t)=\int_{\mathbb{R}} p(x,t)f(x)\,\mathrm{d}x =: \langle p_t, f\rangle.$$

We begin by stating a lemma that is the key to proving Proposition 2.

Lemma 2: Suppose that f is an $\mathscr{B}_t\vee\mathscr{Z}_t$-adapted process satisfying $\mathsf{E}\int_0^t |f(s)|^2\,\mathrm{d}s<\infty$. Then,

$$\mathsf{E}\Big[\int_0^t f(s)\,\mathrm{d}s|\mathscr{Z}_t\Big]=\int_0^t \mathsf{E}[f(s)|\mathscr{Z}_s]\,\mathrm{d}s, \tag{32}$$

$$\mathsf{E}\Big[\int_0^t f(s)\,\mathrm{d}Z_s|\mathscr{Z}_t\Big]=\int_0^t \mathsf{E}[f(s)|\mathscr{Z}_s]\,\mathrm{d}Z_s. \tag{33}$$

■

We next provide a proof of the Proposition 2 and follow it up with the proof of the Lemma 2.

Proof of Proposition 2 Applying Itô's formula to equation (13) gives, for any smooth and bounded function f,

$$\mathrm{d}f(X_t^i)=\mathscr{L}f(X_t^i)\,\mathrm{d}t+\mathsf{K}(X_t^i,t)\frac{\partial f}{\partial x}(X_t^i)\,\mathrm{d}Z_t+\sigma_B\frac{\partial f}{\partial x}(X_t^i)\,\mathrm{d}B_t^i,$$

where $\mathscr{L}f := (a+u)\frac{\partial f}{\partial x}+\frac{1}{2}(\sigma_W^2\mathsf{K}^2+\sigma_B^2)\frac{\partial^2 f}{\partial x^2}$. Therefore,

$$f(X_t^i)=f(X_0^i)+\int_0^t \mathscr{L}f(X_s^i)\,\mathrm{d}s+\int_0^t \mathsf{K}(X_s^i,s)\frac{\partial f}{\partial x}(X_s^i)\,\mathrm{d}Z_s + \sigma_B\int_0^t \frac{\partial f}{\partial x}(X_s^i)\,\mathrm{d}B_s^i.$$

Taking conditional expectations on both sides,

$$\langle p_t,f\rangle=\mathsf{E}(f(X_0^i)\mid\mathscr{Z}_t)+\mathsf{E}\Big[\int_0^t\mathscr{L}f(X_s^i)\,\mathrm{d}s|\mathscr{Z}_t\Big] +\mathsf{E}\Big[\int_0^t \mathsf{K}(X_s^i,s)\frac{\partial f}{\partial x}(X_s^i)\,\mathrm{d}Z_s\mid\mathscr{Z}_t\Big] +\sigma_B\mathsf{E}\Big[\int_0^t\frac{\partial f}{\partial x}(X_s^i)\,\mathrm{d}B_s^i\mid\mathscr{Z}_t\Big]$$

On applying Lemma 2, and the fact that B_t^i is a Wiener process, we conclude that

$$\langle p_t,f\rangle=\langle p_0,f\rangle+\int_0^t\langle p_s,Lf\rangle\,\mathrm{d}s+\int_0^t\langle p_s,\mathsf{K}\frac{\partial f}{\partial x}\rangle\,\mathrm{d}Z_s.$$

The forward equation (15) follows using integration by parts.

■

We now provide a proof of Lemma 2.

Proof of Lemma 2 The key is the functional form (31) of the solution X_t^i: It says that apart from the past values of Z, the solution depends only upon initial condition X_0^i and Wiener process B_t^i that are both independent of Z.

First we suppose that f is simple, i.e.,

$$f(s)=\sum_{i=1}^k F_i 1_{(a_i,b_i]}(s),$$

where $(a_i,b_i]$ are disjoint intervals of $[0,t]$ and F_i is measurable with respect to $\mathscr{B}_{a_i}\vee\mathscr{Z}_{a_i}$. For general f satisfying the assumptions of the lemma, the result will then follow via an application of the dominated convergence theorem.

Once we restrict to simple functions, the essence of the proof is to establish the identity,

$$\mathsf{E}(F_i|\mathscr{Z}_t)=\mathsf{E}(F_i|\mathscr{Z}_{a_i}). \tag{34}$$

Under the measurability assumption we can write $F_i=\phi(\zeta,\xi)$, where $\zeta\in\mathscr{Z}_{a_i}$, $\xi\in\mathscr{B}_{a_i}$ are random variables, and ϕ is a real-valued function. The random variable ξ is independent of $\mathscr{Z}_t$, so that

$$\mathsf{E}(F_i|\mathscr{Z}_t)=\mathsf{E}(\bar\phi(\zeta)|\mathscr{Z}_t),$$

with $\bar\phi(\,\cdot\,)=\mathsf{E}[\phi(\,\cdot\,,\xi)]$. Using the fact that $\zeta\in\mathscr{Z}_{a_i}\subset\mathscr{Z}_t$ we obtain (34):

$$\mathsf{E}(F_i|\mathscr{Z}_t)=\mathsf{E}(\bar\phi(\zeta)|\mathscr{Z}_t)=\mathsf{E}(F_i|\mathscr{Z}_{a_i}).$$

The desired results follow easily from (34): To obtain (32),

$$\begin{aligned}\mathsf{E}(\int_0^t f(s)\,\mathrm{d}s\mid\mathscr{Z}_t)&=\sum_{i=1}^k\mathsf{E}(F_i(b_i-a_i)\mid\mathscr{Z}_t)\\&=\sum_{i=1}^k\mathsf{E}(F_i\mid\mathscr{Z}_{a_i})(b_i-a_i)\\&=\int_0^t\mathsf{E}(f(s)\mid\mathscr{Z}_s)\,\mathrm{d}s.\end{aligned}$$

The proof of (33) is similar:

$$\begin{aligned}\mathsf{E}\Big[\int_0^t f(s)\,\mathrm{d}Z_s\mid\mathscr{Z}_t\Big]&=\sum_{i=1}^k\mathsf{E}[F_i(Z_{b_i}-Z_{a_i})\mid\mathscr{Z}_t]\\&=\sum_{i=1}^k\mathsf{E}[F_i\mid\mathscr{Z}_{a_i}](Z_{b_i}-Z_{a_i})\\&=\int_0^t\mathsf{E}[f(s)\mid\mathscr{Z}_s]\,\mathrm{d}Z_s,\end{aligned}$$

where the second equality uses the fact that Z is adapted to $\mathscr{Z}_t$, and $a_i<b_i\le t$ for each i. ■

B. Euler-Lagrange BVP

In this section we describe, formally, the continuous-time limit of the discrete-time E-L BVP (11).

In the continuous-time case, the control is of the form:

$$U_t^i=u(X_t^i,t)\,\mathrm{d}t+\mathsf{K}(X_t^i,t)\,\mathrm{d}Z_t. \tag{35}$$

Substituting this in the E-L BVP (11) for the continuous-discrete time case, we arrive at the formal equation:

$$\frac{\partial}{\partial x}\left(\frac{p(x,t)}{1+u'\,\mathrm{d}t+\mathsf{K}'\,\mathrm{d}Z_t}\right)=p(x,t)\frac{\partial}{\partial\tilde u}\Big(\ln p(x+u\,\mathrm{d}t+\mathsf{K}\,\mathrm{d}Z_t,t) + \ln p_{\mathrm{dZ}|\mathrm{X}}(\,\mathrm{d}Z_t\mid x+u\,\mathrm{d}t+\mathsf{K}\,\mathrm{d}Z_t)\Big), \tag{36}$$

where $p_{\mathrm{dZ}|\mathrm{X}}(\mathrm{d}Z_t|\cdot)=\frac{1}{\sqrt{2\pi\sigma_W^2}}\exp\Big(-\frac{(\mathrm{d}Z_t-h(\cdot))^2}{2\sigma_W^2}\Big)$ and $\tilde u = u\,\mathrm{d}t+\mathsf{K}\,\mathrm{d}Z_t$.

For notational ease, we use primes to denote partial derivatives with respect to x: p is used to denote $p(x,t)$, $p':=\frac{\partial p}{\partial x}(x,t)$, $p'':=\frac{\partial^2 p}{\partial x^2}(x,t)$, $u':=\frac{\partial u}{\partial x}(x,t)$, $\mathsf{K}':=\frac{\partial\mathsf{K}}{\partial x}(x,t)$ etc. Note that the time t is fixed.

A sketch of calculations to obtain (16) and (17) starting from (36) appears in the following three steps:

Step 1: The three terms in (36) are simplified as:

$$\frac{\partial}{\partial x}\left(\frac{p}{1+u'\,\mathrm{d}t+\mathsf{K}'\,\mathrm{d}Z_t}\right)=p'-f_1\,\mathrm{d}t-(p'\mathsf{K}'+p\mathsf{K}'')\,\mathrm{d}Z_t$$

$$p\frac{\partial}{\partial\tilde{u}}\ln p(x+u\,\mathrm{d}t+\mathsf{K}\,\mathrm{d}Z_t)=p'+f_2\,\mathrm{d}t+(p''\mathsf{K}-\frac{p'^2\mathsf{K}}{p})\,\mathrm{d}Z_t$$

$$p\frac{\partial}{\partial\tilde{u}}\ln p_{\mathrm{d}Z|X}(\,\mathrm{d}Z_t|x+u\,\mathrm{d}t+\mathsf{K}\,\mathrm{d}Z_t)=\frac{p}{\sigma_W^2}[h'\,\mathrm{d}Z_t+(h''\mathsf{K}-hh')\,\mathrm{d}t]$$

where we have used Itô's rules $\mathrm{d}Z_t^2=\sigma_W^2\,\mathrm{d}t$, $\mathrm{d}Z_t\,\mathrm{d}t=0$ etc., and the functions

$$f_1=(p'u'+pu'')-\sigma_W^2(p'\mathsf{K}'^2+2p\mathsf{K}'\mathsf{K}''),$$
$$f_2=(p''u-\frac{p'^2u}{p})+\sigma_W^2\mathsf{K}^2\left(\frac{1}{2}p'''-\frac{3p'p''}{2p}+\frac{p'^3}{p^2}\right).$$

Collecting terms in $O(\,\mathrm{d}Z_t)$ and $O(\,\mathrm{d}t)$, after some simplification, leads to the following ODEs:

$$\mathscr{E}(\mathsf{K})=\frac{1}{\sigma_W^2}h'(x) \tag{37}$$

$$\mathscr{E}(u)=-\frac{1}{\sigma_W^2}h(x)h'(x)+h''(x)\mathsf{K}+\sigma_W^2G(x,t) \tag{38}$$

where $\mathscr{E}(\mathsf{K})=-\frac{\partial}{\partial x}\left(\frac{1}{p(x,t)}\frac{\partial}{\partial x}\{p(x,t)\mathsf{K}(x,t)\}\right)$, and $G=-2\mathsf{K}'\mathsf{K}''-(\mathsf{K}')^2(\ln p)'+\frac{1}{2}\mathsf{K}^2(\ln p)'''$.

Step 2. Suppose (u,K) are admissible solutions of the E-L BVP (37)-(38). Then it is claimed that

$$-(p\mathsf{K})'=\frac{h-\hat{h}_t}{\sigma_W^2}p \tag{39}$$

$$-(pu)'=-\frac{(h-\hat{h}_t)\hat{h}_t}{\sigma_W^2}p-\frac{1}{2}\sigma_W^2(p\mathsf{K}^2)''. \tag{40}$$

Recall that admissible here means

$$\lim_{x\to\pm\infty}p(x,t)u(x,t)=0,\quad \lim_{x\to\pm\infty}p(x,t)\mathsf{K}(x,t)=0. \tag{41}$$

To show (39), integrate (37) once to obtain

$$-(p\mathsf{K})'=\frac{1}{\sigma_W^2}hp+Cp,$$

where the constant of integration $C=-\frac{\hat{h}_t}{\sigma_W^2}$ is obtained by integrating once again between $-\infty$ to ∞ and using the boundary conditions for K (41). This gives (39).

To show (40), we denote its right hand side as $\mathscr{R}$ and claim

$$\left(\frac{\mathscr{R}}{p}\right)'=-\frac{hh'}{\sigma_W^2}+h''\mathsf{K}+\sigma_W^2G. \tag{42}$$

The equation (40) then follows by using the ODE (38) together with the boundary conditions for u (41). The verification of the claim involves a straightforward calculation, where we use (37) to obtain expressions for h' and K''. The details of this calculation are omitted on account of space.

Step 3. The E-L equation for K is given by (37) which is the same as (16). The proof of (17) involves a short calculation starting from (40), which is simplified to the form (17) by using (39).

Proof of Proposition 3. Consider the ODE (16). It is a linear ODE whose unique solution is given by

$$\mathsf{K}(x,t)=\frac{1}{p(x,t)}\left(C_1+C_2\int_{-\infty}^{x}p(y,t)\,\mathrm{d}y-\frac{1}{\sigma_W^2}\int_{-\infty}^{x}h(y)p(y,t)\,\mathrm{d}y\right), \tag{43}$$

where the constant of integrations $C_1=0$ and $C_2=\frac{\hat{h}_t}{\sigma_W^2}$ because of the boundary conditions for K. Part 2 is an easy consequence of the minimum principle for elliptic PDEs [5].

C. Consistency with the nonlinear filter

Proof of Theorem 1 It is only necessary to show that with this choice of $\{u,\mathsf{K}\}$, we have $\mathrm{d}p(x,t)=\mathrm{d}p^*(x,t)$, for all x and t, in the sense that they are defined by identical stochastic differential equations. Recall $\mathrm{d}p^*$ is defined according to the K-S equation (12), and $\mathrm{d}p$ according to the forward equation (15).

If K solves the E-L BVP (16) then using (43),

$$\frac{\partial}{\partial x}(p\mathsf{K})=-\frac{1}{\sigma_W^2}(h-\hat{h}_t)p. \tag{44}$$

On multiplying both sides of (17) by $-p$, we have

$$\begin{aligned}-up&=\frac{1}{2}(h-\hat{h}_t)p\mathsf{K}-\frac{1}{2}\sigma_W^2(p\mathsf{K})\frac{\partial\mathsf{K}}{\partial x}+\hat{h}_tp\mathsf{K}\\&=-\frac{1}{2}\sigma_W^2\frac{\partial(p\mathsf{K})}{\partial x}\mathsf{K}-\frac{1}{2}\sigma_W^2(p\mathsf{K})\frac{\partial\mathsf{K}}{\partial x}+\hat{h}_tp\mathsf{K}\\&=-\frac{1}{2}\sigma_W^2\frac{\partial}{\partial x}(p\mathsf{K}^2)+\hat{h}_tp\mathsf{K},\end{aligned}$$

where we have used (44) to obtain the second equality. Differentiating once with respect to x and using (44),

$$-\frac{\partial}{\partial x}(up)+\frac{1}{2}\sigma_W^2\frac{\partial^2}{\partial x^2}(p\mathsf{K}^2)=-\frac{\hat{h}_t}{\sigma_W^2}(h-\hat{h}_t)p. \tag{45}$$

Using (44)-(45) in the forward equation (15), we have

$$\mathrm{d}p=\mathscr{L}^\dagger p+\frac{1}{\sigma_W^2}(h-\hat{h}_t)(\,\mathrm{d}Z_t-\hat{h}_t\,\mathrm{d}t)p.$$

This is precisely the SDE (12), as desired.

REFERENCES

[1] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*. Springer, Cambridge, Mass, 2010.

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

[4] K. Doya, S. Ishii, A. Pouget, and R. P. N. Rao. *Bayesian Brain*. Comput. Neurosci. MIT Press, Cambridge, MA, 2007.

[5] L. C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 1998.

[6] 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.

[7] Gopinath Kallianpur. *Stochastic filtering theory*, volume 13 of *Applications of Mathematics*. Springer-Verlag, New York, 1980.

[8] T. Yang, P. G. Mehta, and S. P. Meyn. A mean-field control-oriented approach to particle filtering. In *Proc. of American Control Conference*, pages 2037–2043, June 2011.

[9] 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.