Multivariable Feedback Particle Filter

Tao Yang, Richard S. Laugesen, Prashant G. Mehta, Sean P. Meyn

Abstract— In recent work it is shown that importance sampling can be avoided in the particle filter through an innovation structure inspired by traditional nonlinear filtering combined with Mean-Field Game formalisms [9], [19]. The resulting *feedback particle filter* (FPF) offers significant variance improvements; in particular, the algorithm can be applied to systems that are not stable. The filter comes with an up-front computational cost to obtain the filter gain. This paper describes new representations and algorithms to compute the gain in the general multivariable setting. The main contributions are,

(i) Theory surrounding the FPF is improved: Consistency is established in the multivariate setting, as well as well-posedness of the associated PDE to obtain the filter gain.

(ii) The gain can be expressed as the gradient of a function, which is precisely the solution to Poisson's equation for a related MCMC diffusion (the Smoluchowski equation). This provides a bridge to MCMC as well as to approximate optimal filtering approaches such as TD-learning, which can in turn be used to approximate the gain.

(iii) Motivated by a weak formulation of Poisson's equation, a Galerkin finite-element algorithm is proposed for approximation of the gain. Its performance is illustrated in numerical experiments.

I. INTRODUCTION

In a recent work, we introduced a new feedback control-based formulation of the particle filter for the nonlinear filtering problem [18],[17]. The resulting filter is referred to as the *feedback particle filter*. In [18],[17], the filter was described for the scalar case, where the signal and the observation processes are both real-valued. The aim of this paper is to generalize the scalar results of our earlier papers to the multivariable filtering problem:

$$dX_t = a(X_t)\,dt + dB_t, \qquad (1a)$$

$$dZ_t = h(X_t)\,dt + dW_t, \qquad (1b)$$

where $X_t \in \mathbb{R}^d$ is the state at time t, $Z_t \in \mathbb{R}^m$ is the observation process, $a(\cdot)$, $h(\cdot)$ are C^1 functions, and $\{B_t\}$, $\{W_t\}$ are mutually independent Wiener processes of appropriate dimension. The covariance matrix of the observation noise $\{W_t\}$ is assumed to be positive definite. The function h is a column vector whose j-th coordinate is denoted as h_j (i.e., $h = (h_1, h_2, \ldots, h_m)^T$). For notational ease, the process noise $\{B_t\}$ is assumed to be a standard Wiener process. By scaling,

we may assume without loss of generality that the covariance matrices associated with $\{B_t\}$, $\{W_t\}$ are identity matrices.

The objective of the filtering problem is to estimate the posterior distribution of X_t given the history $\mathscr{Z}_t := \sigma(Z_s : s \leq t)$. The posterior is denoted by p^*, so that for any measurable set $A \subset \mathbb{R}^d$,

$$\int_{x \in A} p^*(x,t)\,dx = \mathsf{P}\{X_t \in A \mid \mathscr{Z}_t\}.$$

The filter is infinite-dimensional since it defines the evolution, in the space of probability measures, of $\{p^*(\cdot, t) : t \geq 0\}$. If $a(\cdot)$, $h(\cdot)$ are linear functions, the solution is given by the finite-dimensional Kalman filter. The article [3] 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 [6]. The key step is the construction of N stochastic processes $\{X_t^i : 1 \leq i \leq N\}$: The value $X_t^i \in \mathbb{R}^d$ 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. Recall that this is defined for any measurable set $A \subset \mathbb{R}^d$ by,

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

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], [6].

In our earlier papers [18],[17], an alternative feedback control-based approach to the construction of a particle filter was introduced; see also [4], [14], [5], [12], [16] for related approaches. The resulting particle filter, referred to as the feedback particle filter, was described for the scalar filtering problem (where $d = m = 1$). The main result of this paper is to describe the feedback particle filter for the multivariable filtering problem (1a)-(1b):

The particle filter is a controlled system. The dynamics of the i^{th} particle have the following gain feedback form,

$$\begin{aligned} dX_t^i = {} &a(X_t^i)\,dt + dB_t^i \\ &+ \mathsf{K}(X_t^i,t)\,dI_t^i + \Omega(X_t^i,t)\,dt, \end{aligned} \qquad (2)$$

where $\{B_t^i\}$ are mutually independent standard Wiener processes, I^i is similar to the *innovation process* that appears in the nonlinear filter,

$$dI_t^i := dZ_t - \frac{1}{2}(h(X_t^i) + \hat{h})\,dt, \qquad (3)$$

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

The gain function K is obtained as a solution to an Euler-Lagrange boundary value problem (E-L BVP): For $j = 1, 2, \ldots, m$, the function ϕ_j is a solution to the second-order differential equation,

$$\nabla \cdot (p(x,t)\nabla\phi_j(x,t)) = -(h_j(x) - \hat{h}_j)p(x,t),$$
$$\int \phi_j(x,t)p(x,t)\, dx = 0, \tag{4}$$

where p denotes the conditional distribution of X_t^i given $\mathscr{Z}_t$. In terms of these solutions, the gain function is given by

$$[\mathsf{K}]_{lj} = \frac{\partial \phi_j}{\partial x_l}. \tag{5}$$

Note that the gain function needs to be obtained for each value of time t.

Finally, $\Omega = (\Omega_1, \Omega_2, \ldots, \Omega_d)^T$ is the Wong-Zakai correction term:

$$\Omega_l(x,t) := \frac{1}{2}\sum_{k=1}^{d}\sum_{s=1}^{m} \mathsf{K}_{ks}(x,t)\frac{\partial \mathsf{K}_{ls}}{\partial x_k}(x,t). \tag{6}$$

The controlled system (2)-(6) is called the multivariable feedback particle filter.

The contributions of this paper are as follows:

• **Consistency.** The feedback particle filter (2) is consistent with the nonlinear filter, given consistent initializations $p(\cdot, 0) = p^*(\cdot, 0)$. Consequently, if the initial conditions $\{X_0^i\}_{i=1}^N$ are drawn from the initial distribution $p^*(\cdot, 0)$ of X_0, then, as $N \to \infty$, the empirical distribution of the particle system approximates the posterior distribution $p^*(\cdot, t)$ for each t.

• **Well-posedness.** A weak formulation of (4) is introduced, and used to prove an existence-uniqueness result for ϕ_j in a suitable function space. Certain apriori bounds are derived for the gain function to show that the resulting control input in (2) is admissible (That is, the filter (2) is well-posed in the Itô sense).

• **Numerical algorithms.** Based on the weak formulation, a Galerkin finite-element algorithm is proposed for approximation of the gain function $\mathsf{K}(x,t)$. The algorithm is completely adapted to data (That is, it does not require an explicit approximation of $p(x,t)$ or computation of derivatives). Certain closed-form expressions for gain function are derived in certain special cases. The conclusions are illustrated with numerical examples.

• **Characterization of the feedback gain.** The *Smoluchowski equation* models a d-dimensional gradient flow with "noise":

$$d\Phi_t = -\nabla\mathscr{G}(\Phi_t)\, dt + \sqrt{2}\, d\xi_t, \tag{7}$$

where ξ is a standard Wiener process. It is regarded as the original MCMC algorithm: Under general conditions it is ergodic, with (unnormalized) stationary distribution $e^{-\mathscr{G}}$. The BVP (4) can be expressed as an instance of Poisson's equation for this diffusion,

$$\mathscr{D}\phi_j = -(h_j - \hat{h}_j), \quad 1 \le j \le m, \tag{8}$$

where $\mathscr{D}$ is the differential generator for the Smoluchowski equation, with potential $\mathscr{G} = -\log p$. Subject to growth conditions on h and p, this implies the mean-integral representation for the vector-valued function,

$$\phi_j(x) = \int_0^\infty \mathsf{E}[h_j(\Phi_t) - \hat{h}_j \mid \Phi_0 = x]\, dt. \tag{9}$$

This representation also suggests an alternate proof of well-posedness and construction of numerical algorithms; cf., [8]. This will be the subject of future work.

The outline of the remainder of this paper is as follows. The nonlinear filter is introduced and shown to be consistent in Sec II. The weak formulation of the BVP appears in Sec III where well-posedness results are also derived. Algorithms are discussed in Sec IV and a numerical example in Sec V.

II. MULTIVARIABLE FEEDBACK PARTICLE FILTER

Consider the continuous time filtering problem (1a, 1b) introduced in Sec 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:

$$dp^* = \mathscr{L}^\dagger p^*\, dt + (h - \hat{h})^T(dZ_t - \hat{h}\, dt)p^*, \tag{10}$$

where $\hat{h} = \int h(x)p^*(x,t)\, dx$ and $\mathscr{L}^\dagger p^* = -\nabla \cdot (pa) + \frac{1}{2}\Delta p$, where Δ denotes the Laplacian in $\mathbb{R}^d$.

A. Belief state dynamics & control architecture

The model for the particle filter is given by,

$$dX_t^i = a(X_t^i)\, dt + dB_t^i + \underbrace{u(X_t^i, t)\, dt + \mathsf{K}(X_t^i, t)\, dZ_t}_{dU_t^i}, \tag{11}$$

where $X_t^i \in \mathbb{R}^d$ is the state for the i^{th} particle at time t, 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., independent of $\{B_t^i\}$, and drawn from the initial distribution $p^*(x,0)$ of X_0. Both $\{B_t^i\}$ and $\{X_0^i\}$ are also assumed to be independent of X_t, Z_t. Note that the gain function $\mathsf{K}(x,t)$ is a $d \times m$ matrix and $u(x,t) \in \mathbb{R}^d$.

We impose admissibility requirements on the control input U_t^i in (11):

Definition 1 (Admissible Input): The control input U_t^i 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. And for each t, $\mathsf{E}[|u|] := \mathsf{E}[\sum_l |u_l(X_t^i, t)|] < \infty$ and $\mathsf{E}[|\mathsf{K}|^2] := \mathsf{E}[\sum_{lj} |\mathsf{K}_{lj}(X_t^i, t)|^2] < \infty$. ∎

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

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

The functions $\{u(x,t), \mathsf{K}(x,t)\}$ are said to be *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 (10) and (12)).
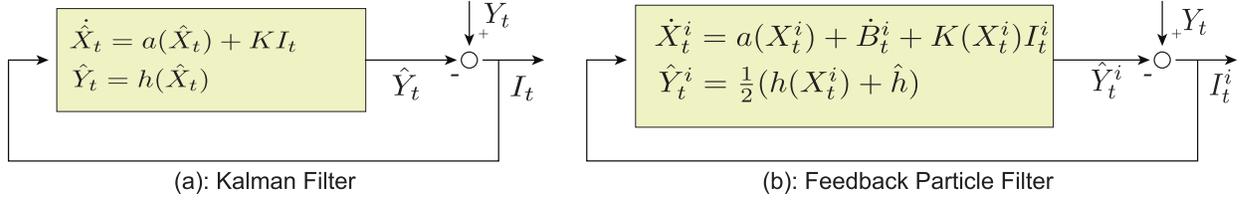
Fig. 1. Innovation error-based feedback structure for (a) Kalman filter and (b) nonlinear feedback particle filter (see Remark (1)).

The evolution equation for the belief state is described in the next result. The proof is identical to the proof in the scalar case (see Proposition 2 in [17]). It is omitted here.

Proposition 1: Consider the process X_t^i that evolves according to the particle filter model (11). The conditional distribution of X_t^i given the filtration $\mathscr{Z}_t$, $p(x,t)$, satisfies the forward equation

$$dp = \mathscr{L}^\dagger p\, dt - \nabla \cdot (p\mathsf{K})\, dZ_t$$
$$- \nabla \cdot (pu)\, dt + \frac{1}{2} \sum_{l,k=1}^{d} \frac{\partial^2}{\partial x_l \partial x_k} \left(p[\mathsf{K}\mathsf{K}^T]_{lk} \right) dt. \tag{12}$$
∎

B. General Form of the Feedback Particle Filter

The general form of the feedback particle filter is obtained by choosing $\{u, \mathsf{K}\}$ as the solution to a certain E-L BVP based on p. The function K is a solution to

$$\nabla \cdot (p\mathsf{K}) = -(h - \hat{h})^T p, \tag{13}$$

and the function u is obtained as

$$u(x,t) = -\frac{1}{2}\mathsf{K}(x,t)\left(h(x) + \hat{h}\right) + \Omega(x,t). \tag{14}$$

The reader is referred to our earlier paper [17] for additional justification regarding these choices.

Remark 1: Substituting (13)-(14) into (11) gives the feedback particle filter model (2)-(3) in Sec I.

In the Stratonovich form, the filter admits a simpler representation,

$$dX_t^i = a(X_t^i)\, dt + dB_t^i + \mathsf{K}(X^i,t) \circ \left(dZ_t - \frac{1}{2}(h(X_t^i) + \hat{h})\, dt \right).$$

Given that the Stratonovich form provides a mathematical interpretation of the (formal) ODE model [15, Section 3.3 of the SDE text by Øksendal], we also obtain the (formal) ODE model of the filter. Denoting $Y_t \doteq \frac{dZ_t}{dt}$ and white noise process $\dot{B}_t^i \doteq \frac{dB_t^i}{dt}$, the ODE model of the filter is given by,

$$\frac{dX_t^i}{dt} = a(X_t^i) + \dot{B}_t^i + \mathsf{K}(X^i,t) \cdot \left(Y_t - \frac{1}{2}(h(X_t^i) + \hat{h}) \right).$$

The feedback particle filter thus provides 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, it is shown in Sec III-D that the gain function is the Kalman gain. For the nonlinear case, the Kalman gain is replaced by a nonlinear function of the state. ∎

If one further *assumes* that the control input U_t^i is admissible, a short calculation shows that the feedback particle filter is consistent with the choice of $\{u, \mathsf{K}\}$ given by (13)-(14). This calculation appears in Appendix A.

C. Consistency with the Nonlinear Filter

To establish admissibility of the input U_t^i requires additional assumptions on the density p and function h:

(i) **Assumption A1** The probability density $p(x,t)$ is of the form $p(x,t) = e^{-\mathscr{G}(x,t)}$, where $\mathscr{G}(x,t)$ is a twice continuously differentiable function with

$$|\nabla \mathscr{G}|^2 - 2\Delta \mathscr{G} \to \infty \quad \text{as } |x| \to \infty. \tag{15}$$

(ii) **Assumption A2** The function h satisfies,

$$\int |h(x)|^2 p(x,t)\, dx < \infty,$$

where $|h(x)|^2 := \sum_j |h_j(x)|^2$. For admissibility of u, our arguments require additional assumptions:

(iii) **Assumption A3** The second derivatives of $\mathscr{G}(x,t)$ with respect to x are uniformly bounded at each t, i.e., $|\frac{\partial^2 \mathscr{G}}{\partial x_j \partial x_k}(x,t)| \leq c_2(t)$ for all $x \in \mathbb{R}^d$, $t > 0$.

(iv) **Assumption A4** The first (weak) derivatives of h satisfy

$$\int |\nabla h(x)|^2 p(x,t)\, dx < \infty,$$

where $|\nabla h(x)|^2 := \sum_{jk} |\frac{\partial h_j}{\partial x_k}(x)|^2$. ∎

Under these assumptions, it is shown in Theorem 2 that the gradient-form (4) of the E-L BVP (13) is uniquely obtained to give ϕ and thence K.

The admissibility of the resulting control input is established in Corollary 1. Theorem 2 and Corollary 1 are stated and proved in Sec III.

The following theorem then shows that the two evolution equations (10) and (12) are identical. The proof appears in Appendix A.

Theorem 1: Consider the two evolution equations for p and p^*, defined according to the solution of the forward equation (12) and the K-S equation (10), respectively. Suppose that the gain function $\mathsf{K}(x,t)$ is obtained according to (4)-(5). Then, provided $p(\cdot,0) = p^*(\cdot,0)$, we have for all $t \geq 0$,

$$p(\cdot,t) = p^*(\cdot,t).$$
∎

III. EXISTENCE, UNIQUENESS AND ADMISSIBILITY

The aim of this section is to introduce a particular *gradient-form solution* of the BVP (13). The gradient-form solution is obtained in terms of m real-valued functions $\{\phi_1(\cdot,t), \phi_2(\cdot,t), \ldots, \phi_m(\cdot,t)\}$. For $j = 1, 2, \ldots, m$, the function ϕ_j is a solution to,

$$\nabla \cdot (p(x,t)\nabla\phi_j(x,t)) = -(h_j(x) - \hat{h}_j)p(x,t),$$
$$\int \phi_j(x,t)p(x,t)\,\mathrm{d}x = 0. \tag{16}$$

The normalization $\int \phi_j(x,t)p(x,t)\,\mathrm{d}x = 0$ is for convenience: If ϕ_j^o is an solution to the differential equation (16), we obtain the desired normalization on subtracting its mean.

In terms of these solutions, the gain function is given by,

$$\mathsf{K}_{lj}(x,t) = \frac{\partial \phi_j}{\partial x_l}(x,t), \quad x \in \mathbb{R}^d. \tag{17}$$

It is straightforward to verify that K thus defined is a particular solution of the BVP (13).

A. Poisson's Equation Interpretation

The differential equation (16) is solved for each t to give the m functions $\{\phi_j(\cdot,t) : 1 \leq j \leq m\}$. On dividing each side of this equation by p, elementary calculus leads to the equivalent equation (8), with generator $\mathscr{D}$ defined for C^2 functions f via,

$$\mathscr{D}f = -\nabla\mathscr{G} \cdot \nabla f + \Delta f$$

and with $\mathscr{G}(\cdot) = -\log p(\cdot,t)$. This is the differential generator for the Smoluchowski equation (7). It is shown in [10] that this diffusion is exponentially ergodic under mild conditions on $\mathscr{G}$. Consequently, $\mathsf{E}[h_j(\Phi_t) - \hat{h}_j \mid \Phi_0 = x]$ converges to zero exponentially fast, subject to growth conditions on h_j, and from this we can conclude that (9) is well defined, and provides a solution to Poisson's equation (8) [8].

Poisson's equation can be regarded as the *value function* that arises in average-cost optimal control, and this is the object of interest in the approximation techniques used in TD-learning for average-cost optimal control [13]. The integral representation (9) suggests approximation techniques based on approximate models for the diffusion Φ.

The remainder of this section is devoted to showing existence and uniqueness of the solution of (16), and admissibility of the resulting control input, obtained using gain function defined by (17).

B. Weak Formulation

Further analysis of this problem requires introduction of Hilbert spaces: $L^2(\mathbb{R}^d; p)$ is used to denote the Hilbert space of functions on $\mathbb{R}^d$ that are square-integrable with respect to density $p(\cdot,t)$ (for a fixed time t); $H^k(\mathbb{R}^d; p)$ is used to denote the Hilbert space of functions whose first k-derivatives (defined in the weak sense) are in $L^2(\mathbb{R}^d; p)$. Denote

$$H_0^1(\mathbb{R}^d; p) := \left\{ \phi \in H^1(\mathbb{R}^d; p) \,\Big|\, \int \phi(x)p(x)\,\mathrm{d}x = 0 \right\}.$$

A function $\phi_j \in H_0^1(\mathbb{R}^d; p)$ is said to be a weak solution of the BVP (16) if

$$\int \nabla\phi_j(x,t) \cdot \nabla\psi(x)p(x,t)\,\mathrm{d}x = \int (h_j(x) - \hat{h}_j)\psi(x)p(x,t)\,\mathrm{d}x, \tag{18}$$

for all $\psi \in H^1(\mathbb{R}^d; p)$.

Denoting $\mathsf{E}[\cdot] := \int \cdot p(x,t)\,\mathrm{d}x$, the weak form of the BVP (16) can also be expressed as

$$\mathsf{E}[\nabla\phi_j \cdot \nabla\psi] = \mathsf{E}[(h_j - \hat{h}_j)\psi], \quad \forall\, \psi \in H^1(\mathbb{R}^d; p). \tag{19}$$

This representation is useful for the numerical algorithm described in Sec IV.

C. Main Results

The existence-uniqueness result for the BVP (16) is described next — Its proof is given in Appendix B.

Theorem 2: Under Assumptions A1-A2, the BVP (16) possesses a unique weak solution $\phi_j \in H_0^1(\mathbb{R}^d; p)$, satisfying

$$\int |\nabla\phi_j(x)|^2 p(x,t)\,\mathrm{d}x \leq \frac{1}{\lambda} \int |h_j(x) - \hat{h}_j|^2 p(x,t)\,\mathrm{d}x. \tag{20}$$

If in addition Assumptions A3-A4 hold, then $\phi_j \in H^2(\mathbb{R}^d; p)$ with

$$\int \left|(D^2\phi_j)(\nabla\phi_j)\right| p(x,t)\,\mathrm{d}x \leq C(\lambda; p) \int |\nabla h_j|^2 p(x,t)\,\mathrm{d}x, \tag{21}$$

where λ is (spectral gap) constant (see Appendix B) and $C(\lambda; p) = \frac{1}{\lambda^{3/2}} \left(\frac{\|D^2(\log p)\|_{L^\infty}}{\lambda} + 1 \right)^{1/2}$. ∎

The apriori bounds (20)-(21) are used to show that the control input for the feedback particle filter is admissible. The proof is omitted on account of space.

Corollary 1: Suppose ϕ_j is the weak solution of BVP (16) as described in Theorem 2. The gain function K is obtained using (17) and u is given by (14). Then

$$\mathsf{E}[|\mathsf{K}|^2] \leq \frac{1}{\lambda} \sum_{j=1}^m \int |h_j(x)|^2 p(x,t)\,\mathrm{d}x,$$

$$\mathsf{E}[|u|] \leq \left(\frac{1}{\lambda} + C(\lambda; p) \right) \sum_{j=1}^m \int \left(|h_j(x)|^2 + |\nabla h_j|^2 \right) p(x,t)\,\mathrm{d}x,$$

where $C(\lambda; p)$ is given in Theorem 2. That is, the resulting control input in (17) is admissible. ∎

D. Linear Gaussian case

Consider the linear system,

$$\mathrm{d}X_t = A\,X_t\,\mathrm{d}t + \mathrm{d}B_t \tag{22a}$$
$$\mathrm{d}Z_t = H\,X_t\,\mathrm{d}t + \mathrm{d}W_t \tag{22b}$$

where A is an $d \times d$ matrix, and H is an $m \times d$ matrix. The initial distribution $p^*(x,0)$ is Gaussian with mean vector μ_0 and covariance matrix Σ_0.

The following proposition shows that the Kalman gain is a gradient-form solution of the multivariable BVP (13):

Proposition 2: Consider the d-dimensional linear system (22a)-(22b). Suppose $p(x,t)$ is assumed to be Gaussian: $p(x,t) = \frac{1}{(2\pi)^{\frac{d}{2}}|\Sigma_t|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x - \mu_t)^T \Sigma_t^{-1}(x - \mu_t)\right)$, where $x =$

$(x_1, x_2, ..., x_d)^T$, $\mu_t = (\mu_{1t}, \mu_{2t}, \ldots, \mu_{dt})^T$ is the mean, Σ_t is the covariance matrix, and $|\Sigma_t| > 0$ denotes the determinant. A solution of the BVP (16) is given by,

$$\phi_j(x,t) = \sum_{k=1}^{d} [\Sigma_t H^T]_{kj} (x_k - \mu_{kt}). \tag{23}$$

Using (17), $\mathsf{K}(x,t) = \Sigma_t H^T$ (the Kalman gain) is the gradient form solution of (13). ∎

The formula (23) is verified by direct substitution in the BVP (16) where the distribution p is multivariable Gaussian.

The gain function yields the following form for the particle filter in this linear Gaussian model:

$$dX_t^i = A\,X_t^i\,dt + dB_t^i + \Sigma_t H^T \left(dZ_t - H \frac{X_t^i + \mu_t}{2}\,dt \right). \tag{24}$$

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 3: Consider the linear Gaussian filtering problem defined by the state-observation equations (22a)-(22b). In this case the posterior distributions of X_t and X_t^i are Gaussian, whose conditional mean and covariance are given by the respective SDE and the ODE,

$$d\mu_t = A\mu_t\,dt + \Sigma_t H^T \left(dZ_t - H\mu_t\,dt \right)$$

$$\frac{d}{dt}\Sigma_t = A\Sigma_t + \Sigma_t A^T + I - \Sigma_t H^T H \Sigma_t$$

∎

The result is verified by substituting $p(x,t) = (2\pi)^{-\frac{d}{2}} |\Sigma_t|^{-\frac{1}{2}} \exp\left[-\frac{1}{2}(x - \mu_t)^T \Sigma_t^{-1}(x - \mu_t) \right]$ in the forward equation (12). The details are omitted on account of space, and because the result is a special case of Theorem 1.

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 (24) for the i^{th} particle is given by

$$dX_t^i = A\,X_t^i\,dt + dB_t^i + \Sigma_t^{(N)} H^T \left(dZ_t - H \frac{X_t^i + \mu_t^{(N)}}{2}\,dt \right).$$

As $N \to \infty$, the empirical distribution of the particle system approximates the posterior distribution $p^*(x,t)$ (by Theorem 3).

IV. FINITE-ELEMENT ALGORITHM

In this section, a Galerkin finite-element algorithm is described to construct an approximate solution of (18). Since there are m uncoupled BVPs, without loss of generality, we assume scalar-valued observation in this section, with $m = 1$, so that $\mathsf{K} = \nabla\phi$. The time t is fixed. The explicit dependence on time is suppressed for notational ease (That is, $p(x,t)$ is denoted as $p(x)$, $\phi(x,t)$ as $\phi(x)$ etc.).

A. Galerkin Approximation

Using (19), the gain function $\mathsf{K} = \nabla\phi$ is a weak solution if

$$\mathsf{E}[\mathsf{K} \cdot \nabla\psi] = \mathsf{E}[(h - \hat{h})\psi], \quad \forall\,\psi \in H^1(\mathbb{R}^d; p). \tag{25}$$

The gain function is approximated as,

$$\mathsf{K} = \sum_{l=1}^{L} \kappa_l \chi_l(x),$$

where $\{\chi_l(x)\}_{l=1}^L$ are *basis functions*. For each $l = 1, \ldots, L$, $\chi_l(x)$ is a gradient function; That is, $\chi_l(x) = \nabla\zeta_l(x)$ for some function $\zeta_l(x) \in H_0^1(\mathbb{R}^d; p)$.

The *test functions* are denoted as $\{\psi_k(x)\}_{k=1}^L$ and $S := \text{span}\{\psi_1(x), \psi_2(x), \ldots, \psi_L(x)\} \subset H^1(\mathbb{R}^d; p)$.

The finite-dimensional approximation of the BVP (25) is to choose constants $\{\kappa_l\}_{l=1}^L$ such that

$$\sum_{l=1}^{L} \kappa_l \mathsf{E}[\chi_l \cdot \nabla\psi] = \mathsf{E}[(h - \hat{h})\psi], \quad \forall\,\psi \in S. \tag{26}$$

Denoting $[A]_{kl} = \mathsf{E}[\chi_l \cdot \nabla\psi_k]$, $b_k = \mathsf{E}[(h - \hat{h})\psi_k]$, $\kappa = (\kappa_1, \kappa_2, \ldots, \kappa_L)^T$, the finite-dimensional approximation (26) is expressed as a linear matrix equation:

$$A\kappa = b.$$

The matrix A and vector b are easily approximated by using only the particles:

$$[A]_{kl} = \mathsf{E}[\chi_l \cdot \nabla\psi_k] \approx \frac{1}{N} \sum_{i=1}^{N} \chi_l(X_t^i) \cdot \nabla\psi_k(X_t^i), \tag{27}$$

$$b_k = \mathsf{E}[(h - \hat{h})\psi_k] \approx \frac{1}{N} \sum_{i=1}^{N} (h(X_t^i) - \hat{h})\psi_k(X_t^i), \tag{28}$$

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

B. Example 1: Constant Gain Approximation

Suppose $\chi_l = e_l$, the canonical coordinate vector with value 1 for the l^{th} coordinate and zero otherwise. The test functions are the coordinate functions $\psi_k(x) = x_k$ for $k = 1, 2, \ldots, d$. Denoting $\psi(x) = (\psi_1, \psi_2, \ldots, \psi_d)^T = x$,

$$\kappa = \mathsf{E}[\mathsf{K}] = \mathsf{E}[(h - \hat{h})\psi] = \int (h(x) - \hat{h})\psi(x)p(x)\,dx$$

$$\approx \frac{1}{N} \sum_{i=1}^{N} (h(X_t^i) - \hat{h})X_t^i. \tag{29}$$

This formula yields the *constant-gain approximation* of the gain function.

C. Example 2: Single-state Case

Consider a scalar example, where the density is a sum of Gaussian,

$$p(x) \approx \sum_{j=1}^{3} \lambda^j q^j(x),$$

where $q^j(x) = q(x; \mu^j, \Sigma^j) = \frac{1}{\sqrt{2\pi\Sigma^j}} \exp(-\frac{(x - \mu^j)^2}{2\Sigma^j})$, $\lambda^j > 0$, $\sum \lambda^j = 1$. The parameter values for $\lambda^j, \mu^j, \Sigma^j$ are tabulated in Table I.
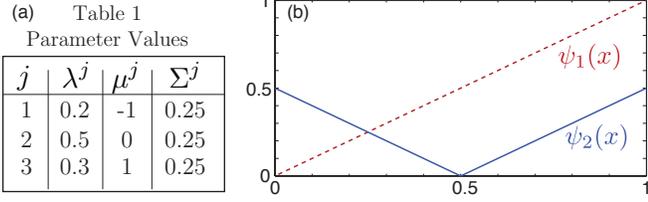
Fig. 2. (a) Parameter values and (b) (ψ_1, ψ_2) in the Example.

In the scalar case, a direct numerical solution (DNS) of the gain function is obtained by numerically approximating the integral

$$\mathsf{K}(x) = -\frac{1}{p(x)} \int_{-\infty}^{x} (h(y) - \hat{h}) p(y) \, \mathrm{d}y.$$

The DNS solution is used to provide comparisons with the approximate Galerkin solutions.

The Galerkin approximation of the gain function is constructed on an interval domain $D \subset \mathbb{R}$. The domain is a union of finitely many non-intersecting intervals $D_l = [a_{l-1}, a_l)$, where $a_0 < a_1 < \ldots < a_L$.

Define for $l = 1, 2, \ldots, L$ and $k = 1, 2, \ldots, L$:

Basis functions: $\chi_l(x) = 1_{D_l}(x)$,

Test functions: $\psi_k(x) = |x - a_k|$.

Figure 2 depicts the test functions $\{\psi_1(x), \psi_2(x)\}$ for $D = [0, 1]$ and $a_0 = 0$, $a_1 = \frac{1}{2}$ and $a_2 = 0$. The basis functions are the indicator functions on $[0, \frac{1}{2})$ and $[\frac{1}{2}, 1)$.

Figure 3 depicts a comparison of the DNS solution and the Galerkin solution for $h(x) = x^2$, $D = [-2, 2]$ and $L = 1, 5, 15$. For a given L, the basis and test functions are constructed for a uniform partition of the domain (That is, $a_l = -2 + \frac{l}{L}4$). The Galerkin solution is obtained using $N = 1000$ particles that are sampled from the distribution p. The particles are used to compute matrix A and vector b, using formulae (27) and (28), respectively. Since the analytical form of p is known, these matrices can also be assembled by using the integrals:

$$[A]_{kl} = \int \chi_l(x) \cdot \nabla \psi_k(x) p(x) \, \mathrm{d}x, \quad (30)$$

$$b_k = \int (h(x) - \hat{h}) \psi_k(x) p(x) \, \mathrm{d}x. \quad (31)$$

The figure also depicts the Galerkin solution based on the integral evaluation of the matrix A and vector b.

For $L = 15$, the matrix A was found to be singular for the particle-based implementation. This is because there are no particles in D_{15}. In this case, the Galerkin solution is obtained using only the integral formulae (30)-(31). These formulae are exact while the particle-based formulae (27) and (28) are approximations. In the other two cases ($L = 1$ and $L = 5$), the particle-based solution provides a good approximation.

V. NUMERICS

Consider a target tracking problem with two bearing-only sensors [2]. A single target moves in a two-dimensional
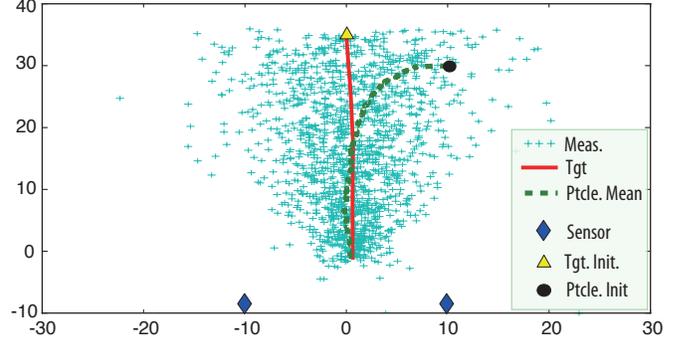

Fig. 4. Simulation results: Comparison of the true target trajectory with the estimate obtained using FPF.

(2d) plane according to the standard white-noise acceleration model:

$$\mathrm{d}X_t = AX_t \, \mathrm{d}t + \Gamma \, \mathrm{d}B_t,$$

where $X := (X_1, V_1, X_2, V_2)^T \in \mathbb{R}^4$, (X_1, X_2) denotes the position and (V_1, V_2) denotes the velocity. The matrices,

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \Gamma = \sigma_B \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix},$$

and B_t is a standard 2d Wiener process.

The observation model is given by,

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

where W_t is a standard 2d Wiener process, $h = (h_1, h_2)^T$ and

$$h_j(x_1, v_1, x_2, v_2) = \arctan\left(\frac{x_2 - x_2^{(\text{sen } j)}}{x_1 - x_1^{(\text{sen } j)}}\right), \quad j = 1, 2,$$

where $(x_1^{(\text{sen } j)}, x_2^{(\text{sen } j)})$ denote the position of sensor j.

Figure 4 depicts a sample path obtained for a typical numerical experiment. The sensor and target locations are depicted together with an estimate (conditional mean) that is approximated using a feedback particle filter. The background depicts the ensemble of observations that were made over the simulation run. Each point in the ensemble is obtained by using the process of triangulation based on two (noisy) angle measurements. The simulation parameters are: The initial position of the target is depicted, the initial velocity was chosen as $(0.2, -5)$ and $\sigma_B = 0.1$; The two sensor positions are depicted and $\sigma_W = 0.017$; The particle filter comprised of $N = 200$ particles whose initial position was chosen from a Gaussian distribution whose mean is depicted. The gain function was obtained using the constant gain approximation in (29). The simulation results show that the filter can adequately track the target.

APPENDIX

A. Proof of Theorem 1

It is only necessary to show that with the choice of $\{u, \mathsf{K}\}$ given by (13)-(14), we have $\mathrm{d}p(x,t) = \mathrm{d}p^*(x,t)$, for all x and
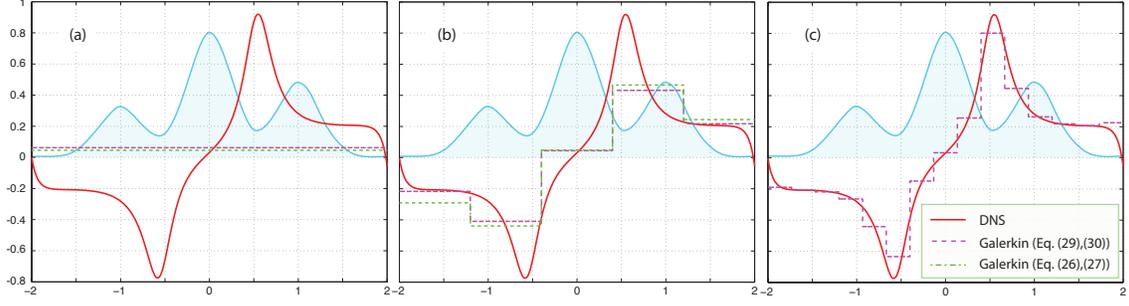
Fig. 3. Comparison of the DNS and the Galerkin approximations of the gain function for $h(x) = x^2$ and: (a) $L = 1$, (b) $L = 5$ and (c) $L = 15$. The density is depicted as the shaded curve in the background.

t, in the sense that they are defined by identical stochastic differential equations. Recall dp^* is defined according to the K-S equation (10), and dp according to the forward equation (12).

Recall that the gain function K is a solution of the following BVP:

$$\nabla \cdot (pK) = -p(h - \hat{h})^T . \tag{32}$$

On multiplying both sides of (14) by $-p$, we obtain

$$-up = \frac{1}{2}p(h - \hat{h})K - \Omega p + pK\hat{h}$$
$$= -\frac{1}{2}K[\nabla \cdot (pK)]^T - \Omega p + pK\hat{h} \tag{33}$$

where (32) is used to obtain the second equality. Denoting $E := \frac{1}{2}K[\nabla \cdot (pK)]^T$, a direct calculation shows that

$$E_l + \Omega_l p = \sum_{k=1}^{d} \frac{\partial}{\partial x_k}\left(p[KK^T]_{lk}\right).$$

Substituting this in (33), on taking the divergence of both sides, we obtain

$$-\nabla \cdot (pu) + \frac{1}{2}\sum_{l,k=1}^{d} \frac{\partial^2}{\partial x_l \partial x_k}\left(p[KK^T]_{lk}\right) = \nabla \cdot (pK)\hat{h}. \tag{34}$$

Using (32) and (34) in the forward equation (12),

$$dp = \mathscr{L}^\dagger p + (h - \hat{h})^T(dZ_t - \hat{h}dt)p.$$

This is precisely the SDE (10), as desired. ∎

B. Proof of Theorem 2

We omit the subscript "j" in this proof, writing just ϕ and h. We also suppress explicit dependence on time t, writing $p(x)$ instead of $p(x,t)$ and $\phi(x)$ instead of $\phi(x,t)$.

Under Assumption A1, p is known to admit a spectral gap (or Poincaré inequality) with constant $\lambda \geq c_1$ [11]: That is, for all functions $\phi \in H_0^1(\mathbb{R}^d; p)$,

$$\int |\phi(x)|^2 p(x)\,dx \leq \frac{1}{\lambda}\int |\nabla \phi(x)|^2 p(x)\,dx. \tag{35}$$

Consider now the inner product

$$< \phi, \psi >:= \int \nabla \phi(x) \cdot \nabla \psi(x) p(x)\,dx.$$

On account of (35), the norm defined by using the inner product $< \cdot, \cdot >$ is equivalent to the standard norm in $H^1(\mathbb{R}^d; p)$.

(i) Consider the BVP in its weak form (18). The integral on the righthand-side is a bounded linear functional on $\psi \in H_0^1$, since

$$|\int (h(x) - \hat{h})\psi(x)p(x)\,dx|^2$$
$$\leq \int |h(x) - \hat{h}|^2 p(x)\,dx \int |\psi(x)|^2 p(x)\,dx$$
$$\leq (\text{const.}) \int |\nabla \psi(x)|^2 p(x)\,dx,$$

where (35) is used to obtain the second inequality. It follows from the Riesz representation theorem that there exists a unique $\phi \in H_0^1$ such that

$$< \phi, \psi >= \int (h(x) - \hat{h})\psi(x)p(x)\,dx,$$

for all $\psi \in H_0^1(\mathbb{R}^d; p)$. Thus ϕ is a weak solution of the BVP, satisfying (18).

(ii) Suppose ϕ is a weak solution. Using $\psi = \phi$ in (18),

$$\int |\nabla \phi|^2 p(x)\,dx = \int (h(x) - \hat{h})\phi(x)p(x)\,dx$$
$$\leq \left(\int |h(x) - \hat{h}|^2 p(x)\,dx\right)^{\frac{1}{2}}\left(\int |\phi(x)|^2 p(x)\,dx\right)^{\frac{1}{2}}$$
$$\leq \left(\int |h(x) - \hat{h}|^2 p(x)\,dx\right)^{\frac{1}{2}}\left(\frac{1}{\lambda}\int |\nabla \phi(x)|^2 p(x)\,dx\right)^{\frac{1}{2}}$$

by (35). The estimate (20) follows.

(iii) For the final estimate (21), we need:

Lemma 1: Under Assumptions A1-A4, the weak solution ϕ of the BVP (18) belongs to $H^2(\mathbb{R}^d; p)$, with

$$\int |D^2\phi|^2 p\,dx \leq \int \nabla \phi \cdot Gp\,dx \tag{36}$$

where the vector function $G \in L^2(\mathbb{R}^d; p)$ is defined by

$$G = D^2(\log p)\nabla \phi + \nabla h$$

and where $|D^2\phi|^2 = \sum_{j,k}\left(\frac{\partial^2 \phi}{\partial x_j \partial x_k}\right)^2$.

Proof: First note that each entry of the Hessian matrix $D^2(\log p)$ is bounded, by Assumption A3, and that $\nabla h \in L^2(\mathbb{R}^d; p)$ by Assumption A4. Hence $G \in L^2(\mathbb{R}; p)$.

Next, elliptic regularity theory [7, Section 6.3 of the PDE text by Evans] applied to the weak solution $\phi \in$

$H^1(\mathbb{R}^d; p)$ says that $\phi \in H^3_{loc}(\mathbb{R}^d)$. Hence the partial differential equation holds pointwise:

$$-\nabla \cdot (p\nabla\phi) = (h - \hat{h})p. \qquad (37)$$

We differentiate with respect to x_k to obtain:

$$-\nabla \cdot (p\nabla\frac{\partial\phi}{\partial x_k}) - \nabla(\frac{\partial \log p}{\partial x_k}) \cdot (p\nabla\phi) - \frac{\partial \log p}{\partial x_k}\nabla \cdot (p\nabla\phi)$$
$$= \frac{\partial h}{\partial x_k}p + (h - \hat{h})\frac{\partial \log p}{\partial x_k}p.$$

The final terms on the left and right sides cancel, by equation (37). Thus the preceding formula becomes

$$-\nabla \cdot (p\nabla\frac{\partial\phi}{\partial x_k}) = G_k p, \qquad (38)$$

where G_k denotes the kth component of $G(x)$.

Let $\beta(x) \geq 0$ be a smooth, compactly supported "bump" function, meaning $\beta(x)$ is radially decreasing with $\beta(0) = 1$. Let $s > 0$ and multiply (38) by $\beta(sx)^2\frac{\partial\phi}{\partial x_k}$. Integrate by parts on the left side (noting the boundary terms vanish because β has compact support) to obtain

$$\int \nabla[\beta(sx)^2\frac{\partial\phi}{\partial x_k}] \cdot (\nabla\frac{\partial\phi}{\partial x_k})p\,dx = \int \beta(sx)^2\frac{\partial\phi}{\partial x_k}G_k p\,dx.$$

The right hand side RHS $\to \int \frac{\partial\phi}{\partial x_k}G_k p\,dx$ by dominated convergence as $s \to 0$, since $\beta(0) = 1$. The left side,

$$\text{LHS} = \int \beta(sx)^2|\nabla\frac{\partial\phi}{\partial x_k}|^2 p\,dx$$
$$+ 2s\int \frac{\partial\phi}{\partial x_k}\beta(sx)(\nabla\beta)(sx) \cdot (\nabla\frac{\partial\phi}{\partial x_k})p\,dx.$$

Clearly the second term is bounded by

$$2s\|\nabla\beta\|_{L^\infty(\mathbb{R}^d)}\int |\frac{\partial\phi}{\partial x_k}|\beta(sx)|\nabla\frac{\partial\phi}{\partial x_k}|p\,dx$$
$$\leq s\|\nabla\beta\|_{L^\infty(\mathbb{R}^d)}\int \left[(\frac{\partial\phi}{\partial x_k})^2 + \beta(sx)^2|\nabla\frac{\partial\phi}{\partial x_k}|^2\right]p\,dx$$

and so

$$(1 - s\|\nabla\beta\|_{L^\infty(\mathbb{R}^d)})\int \beta(sx)^2|\nabla\frac{\partial\phi}{\partial x_k}|^2 p\,dx$$
$$-s\|\nabla\beta\|_{L^\infty(\mathbb{R}^d)}\int (\frac{\partial\phi}{\partial x_k})^2\,dx \leq \text{LHS}.$$

Letting $s \to 0$ in LHS and RHS, and recalling that $\beta(x)$ is radially decreasing, we conclude from the monotone convergence theorem that

$$\int |\nabla\frac{\partial\phi}{\partial x_k}|^2 p\,dx \leq \int \frac{\partial\phi}{\partial x_k}G_k p\,dx.$$

Summing over k completes the proof of the lemma. ∎
Next we prove (21). We will use several times that

$$\int |\nabla\phi|^2 p\,dx \leq \frac{1}{\lambda}\int |h - \hat{h}|^2 p\,dx \leq \frac{1}{\lambda^2}\int |\nabla h|^2 p\,dx, \qquad (39)$$

by (20) followed by (35) applied to the function $h - \hat{h} \in H^1_0(\mathbb{R}^d; p)$.

We have

$$\int |(D^2\phi)(\nabla\phi)|p\,dx \leq \int |D^2\phi||\nabla\phi|p\,dx$$
$$\leq (\frac{1}{\lambda^2}\int |\nabla h|^2 p\,dx)^{3/4}(\int |G|^2 p\,dx)^{1/4}$$

by Lemma 1, Cauchy–Schwarz, and (39). The definition of G, the L^2-triangle inequality and (39) show that

$$(\int |G|^2 p\,dx)^{1/2}$$
$$\leq \|D^2(\log p)\|_{L^\infty}(\int |\nabla\phi|^2 p\,dx)^{1/2} + (\int |\nabla h|^2 p\,dx)^{1/2}$$
$$\leq \left(\frac{\|D^2(\log p)\|_{L^\infty}}{\lambda} + 1\right)(\int |\nabla h|^2 p\,dx)^{1/2}.$$

The estimate (21) now follows.

REFERENCES

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

[2] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., New York, NY, USA, 2002.

[3] 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.

[4] D. Crisan and J. Xiong. Approximate McKean-Vlasov representations for a class of SPDEs. *Stochastics: An International Journal of Probability and Stochastic Processes*, pages 1–16, 2009.

[5] F. Daum and J. Huang. Generalized particle flow for nonlinear filters. *In Proc. of SPIE*, 7698, 2010.

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

[7] L. C. Evans. *Partial differential equations*. American Mathematical Society, Providence, RI, 1998.

[8] P. W. Glynn and S. P. Meyn. A Liapounov bound for solutions of the Poisson equation. *Ann. Probab.*, 24(2):916–931, 1996.

[9] 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.

[10] W. Huisinga, S. Meyn, and C. Schütte. Phase transitions and metastability in Markovian and molecular systems. *Ann. Appl. Probab.*, 14(1):419–458, 2004.

[11] J. Johnsen. On the spectral properties of Witten-Laplacians, their range projections and Brascamp-Lieb's inequality. *Integral Equations Operator Theory*, 36(3):288–324, 2000.

[12] R. Ma and T. P. Coleman. Generalizing the posterior matching scheme to higher dimensions via optimal transportation. In *Allerton Conference on Communication, Control, and Computing*, pages 96–102, 2011.

[13] S. P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, Cambridge, 2007. Pre-pub. edition available online.

[14] S. K. Mitter and N. J. Newton. A variational approach to nonlinear estimation. *SIAM J. Cont. Opt.*, 42(5):1813–1833, 2003.

[15] B. K. Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, Berlin, 2003.

[16] S. Pequito, A. P. Aguiar, B. Sinopoli, and P. A. Gome. Nonlinear estimation using mean field games. In *NetGCOOP 2011 : International conference on Network Games, Control and Opt.*, 2011.

[17] T. Yang, P. G. Mehta, and S. P. Meyn. Feedback particle filter with mean-field coupling. *In Proc. of IEEE Conference on Decision and Control*, pages 7909–7916, December 2011.

[18] 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.

[19] 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.