Particle Filters for State Estimation of Jump Markov Linear Systems

Arnaud Doucet, Neil J. Gordon, and Vikram Krishnamurthy, *Senior Member, IEEE*

Abstract—**Jump Markov linear systems (JMLS) are linear systems whose parameters evolve with time according to a finite state Markov chain. In this paper, our aim is to recursively compute optimal state estimates for this class of systems. We present efficient simulation-based algorithms called particle filters to solve the optimal filtering problem as well as the optimal fixed-lag smoothing problem. Our algorithms combine sequential importance sampling, a selection scheme, and Markov chain Monte Carlo methods. They use several variance reduction methods to make the most of the statistical structure of JMLS.**

Computer simulations are carried out to evaluate the performance of the proposed algorithms. The problems of on-line deconvolution of impulsive processes and of tracking a maneuvering target are considered. It is shown that our algorithms outperform the current methods.

Index Terms—**Filtering theory, Monte Carlo methods, state estimation, switching systems.**

NOMENCLATURE

n_z	dimension of an arbitrary vector z.
$t \in \{1, 2, \cdots\}$	discrete time.
k	iteration number of the various iterative algorithms.
For $p < q$	$\mathbf{z}_{p:q} \triangleq (z_p, z_{p+1}, \cdots, z_q)$.
$\Phi(m, \Sigma)$	$\lvert 2\pi\Sigma \rvert^{-1/2} \exp(-(1/2)m^{\mathsf{T}}\Sigma^{-1/2}m)$.
$\mathcal{N}(m, \Sigma)$	Gaussian distribution of mean m and covariance Σ.
$\mathcal{U}[0, 1]$	uniform distribution on $[0, 1]$.
$z \sim p(z)$	z distributed according to $p(z)$.
$z \lvert y \sim p(z)$	conditional upon y, z distributed according to $p(z)$.
I_n	identity matrix of dimensions $n \times n$.
A^{T}	transpose matrix.

I. INTRODUCTION

THE FIELD of applied statistics has been revolutionized during the past ten years due to the development of several remarkable stochastic sampling algorithms that are collectively termed *Markov chain Monte Carlo* (MCMC) methods [33].

A. Doucet is with the Signal Processing Group, Department of Engineering, University of Cambridge, Cambridge, U.K. (e-mail: ad2@eng.cam.ac.uk).

N. J. Gordon is with the Signal Processing and Imagery Department, Defense Evaluation and Research Agency, Malvern, Worcestershire, U.K. (e-mail: N.Gordon@signal.dera.gov.uk).

V. Krishnamurthy is with the Department of Electrical and Electronic Engineering, the University of Melbourne, Parkville, Australia (e-mail: v.krishnamurthy@ee.mu.oz.au).

Publisher Item Identifier S 1053-587X(01)01419-2.

MCMC methods are simulation-based algorithms that have led to powerful numerical methods for computation of likelihoods, posterior distributions, and estimates derived from them. Most of the development in MCMC methods so far has focused on off-line algorithms that operate on a fixed batch of data. The aim of this paper is to propose and analyze recursive (on-line) simulation-based algorithms. These algorithms combine sequential importance sampling and MCMC algorithms. Motivated by several applications in signal processing outlined below, we focus on deriving recursive algorithms for optimal state estimation of jump Markov linear systems (JMLS)—which is a well-known NP-hard problem.

Let $r_t, t = 1, 2, \cdots$ denote a discrete time Markov chain with known transition probabilities. A jump Markov linear system can be modeled as

$$x_{t+1} = A(r_{t+1})x_t + B(r_{t+1})v_{t+1} + F(r_{t+1})u_{t+1} \quad (1)$$

$$y_t = C(r_t)x_t + D(r_t)\varepsilon_t + G(r_t)u_t \quad (2)$$

where u_t denotes a known exogenous input, and v_t and ε_t denote independent white Gaussian noise sequences. A jump Markov linear system can be viewed as a linear system whose parameters $(A(r_t), B(r_t), C(r_t), D(r_t), F(r_t), G(r_t))$ evolve with time according to a finite state Markov chain r_t. Neither the continuous-state process x_t nor the finite state process r_t are observed—instead, we observe the noisy measurement process y_t.

Jump Markov linear systems are widely used in several fields of signal processing including seismic signal processing [30], digital communications such as interference suppression in CDMA spread spectrum systems [26], target tracking [4], [29], and de-interleaving of pulse trains [31]. They can be viewed as a generalization of the hidden Markov model (HMM) (which consists of a Markov chain r_t observed in white noise) to correlated noise.

Under assumptions detailed later on, it is well known that exact computation of the conditional mean filtered or smoothed state estimates of x_t and r_t involves a prohibitive computational cost exponential in the (growing) number of observations [37]. This is unlike the standard HMM for which conditional mean state estimates can be computed with linear complexity in the number of observations via the HMM filter. Recently, efficient batch (off-line) deterministic and stochastic iterative algorithms have been proposed to compute fixed-interval smoothed conditional mean and maximum *a posteriori* (MAP) state estimates of x_t and r_t; see [10], [11], and [26]. However, in most real-world applications, one wants to compute state estimates of x_t and r_t

recursively (on-line). The aim of this paper is to present *simulation-based recursive* filtering algorithms for computing conditional mean estimates of the states x_t and r_t given the observation history, namely, $\mathbb{E}(x_t | \mathbf{y}_{1:t})$ and $\mathbb{E}(r_t | \mathbf{y}_{1:t})$. Simulation-based algorithms for computing fixed-lag state estimates $\mathbb{E}(x_t | \mathbf{y}_{1:t+L})$ and $\mathbb{E}(r_t | \mathbf{y}_{1:t+L})$ are also presented—the fixed lag L is a fixed positive integer.

Due to the prohibitive computational cost required to compute fixed-lag and filtered state estimates of x_t and r_t, it is necessary to consider in practice suboptimal estimation algorithms. A variety of algorithms has already been proposed in the literature to solve these estimation problems [4], [18], [36], [37]. Most of these algorithms are based on deterministic finite Gaussian mixture approximations like the popular Interacting multiple model (IMM) or the generalized pseudo-Bayes (GPB) algorithms [4]. These methods are computationally cheap, but they can fail in difficult situations.

Another possible suboptimal strategy is to compute a fixed grid approximation to the filtered state density. This involves approximating the continuous-valued process x_t by a finite state process $\overline{x}_t$ with fixed states $q_1, q_2, \cdots, q_N$ (say). The *fixed grid* point values $q_1, q_2, \cdots, q_N$ are called "particles." The filtered state density and, hence, filtered state estimates can then be computed easily at these grid points according to Bayes' rule. The values of the filtered state density at these grid points are called the "weights" of the particles. However, such a fixed grid approximation suffers from the curse of dimensionality—the approximation error depends on the state dimension of the underlying jump Markov linear system. For example, for an uniform fixed grid, the approximation error behaves as $O(N^{-1/D})$, where D denotes the state space dimension and N the number of grid points (particles).

In this paper, we present Monte Carlo (MC) particle filters for computing the conditional mean estimates. These particle filters can be viewed as a *randomized adaptive grid* approximation. As will be shown later, the particles (values of the grid) evolve randomly in time according to a simulation-based rule. The weights of the particles are updated according to Bayes' rule. The most striking advantage of these MC particle filters is that the convergence rate toward zero of the approximation error is independent of the state dimension. That is, the randomization implicit in the MC particle filter gets around the curse of dimensionality [8].

Taking advantage of the increase of computational power and the availability of parallel computers, several authors have recently proposed such MC particle methods [3], [16], [21], [22]. It has been shown that these methods outperform the standard suboptimal methods. In this paper, we propose improved simulation-based approximations of the optimal filter and smoother with novel variance reduction methods; the filtering and fixed-lag smoothing distributions of interest are approximated by a Gaussian mixture of a large number, say, N, of components that evolve stochastically over time and are driven by the observations.

MC particle methods to solve optimal estimation problems were introduced in automatic control at the end of the 1960s by Handschin and Mayne [17] . Interesting developments were then subsequently proposed in the 1970s [1], [35]. Most likely because of the primitive computers available at the time, these

papers were overlooked and forgotten. In the beginning of the 1990s, the great increase in computational power allowed for the rebirth of this field. In 1993, Gordon *et al.* [15] proposed an algorithm (the bootstrap filter) that introduced a selection step that statistically multiplies and/or discards particles at each time. This key step led to the first operational particle filter. Following this seminal paper, particle filters have stimulated great interest in the engineering and statistical literature; see [12] for a summary of the state of the art. With these filters, complex nonlinear non-Gaussian estimation problems can be solved efficiently in an on-line manner. Moreover, they are much easier to implement than classical numerical methods.

The bootstrap filter is a simple algorithm that can be easily applied to JMLS. However, in its standard form, it does not use all the salient structure of this model. We propose here simulation-based algorithms that make use of this structure and include efficient variance reduction techniques. Our algorithms have a computational complexity of $O(N)$ at each time step and can be easily implemented on parallel computers. The filtering algorithm is shown to be more efficient than the most recent computational methods; see [3], [6], [16], [21], and [22]. Finally, we show how this filtering algorithm can be combined with MCMC methods [33] to obtain an efficient fixed-lag smoothing algorithm. Given the importance of JMLS, earlier papers have already partly developed similar ideas [1], [6], [35]. We discuss in detail these issues in Section V.

This paper is organized as follows. Section II presents the signal model and estimation objectives in a formal way. In Section III, an original simulation-based method is proposed to solve the optimal filtering problem. We detail the different steps of this method. In Section IV, after having shown that a direct application of the previous methods to fixed-lag smoothing is inefficient, we propose an original method based on the introduction of MCMC algorithms. In Section V, a discussion of the previous work on related problems and on the algorithms developed here is given. In Section VI, we demonstrate the performance of the proposed algorithms via computer simulations for deconvolution of impulsive processes and tracking of a maneuvering target. In Appendix A, the backward information filter equations are recalled. Finally, the proofs of some propositions are grouped in Appendix B. Detailed proofs of propositions can be found in [14].

II. PROBLEM FORMULATION

A. Signal Model

Let r_t denote a discrete-time, time-homogeneous, s-state, first-order Markov chain with transition probabilities $p_{m,n} \triangleq \Pr\{r_{t+1} = n | r_t = m\}$ for any $m, n \in S$, where $S \triangleq \{1, 2, \cdots, s\}$. The transition probability matrix $[p_{m,n}]$ is thus an $s \times s$ matrix, with elements satisfying $p_{m,n} \geq 0$ and $\sum_{n=1}^{s} p_{m,n} = 1$, for each $m \in S$. Denote the initial probability distribution as $p_m \triangleq \Pr\{r_1 = m\}$, for $m \in S$, such that $p_m \geq 0, \forall m \in S$ and $\sum_{m=1}^{s} p_m = 1$. Consider the following JMLS given in (1) and (2), where $x_t \in \mathbb{R}^{n_x}$ is the system state, $y_t \in \mathbb{R}^{n_y}$ is the observation at time t, $u_t \in \mathbb{R}^{n_u}$ is a known deterministic input, $v_t \overset{\text{i.i.d}}{\sim} \mathcal{N}(0, I_{n_v}) \in \mathbb{R}^{n_v}$ and

$\varepsilon_t \overset{\text{i.i.d}}{\sim} \mathcal{N}(0, I_{n_w}) \in \mathbb{R}^{n_w}$ are i.i.d. Gaussian sequences, and $D(i)D^{\mathrm{T}}(i) > 0$ $(\forall i \in S)$. $A(\cdot), B(\cdot), C(\cdot), D(\cdot), F(\cdot)$ and $G(\cdot)$ are functions of the Markov chain state r_t, i.e., $(A(\cdot), B(\cdot), C(\cdot), D(\cdot), F(\cdot), G(\cdot)) \subset \{(A(m), B(m), C(m), D(m), F(m), G(m)); m \in S\}$, and they evolve according to the realization of the finite state Markov chain r_t. We assume that $x_0 \sim \mathcal{N}(\hat{x}_0, P_0)$, where $P_0 > 0$, and let x_0, v_t and ε_t be mutually independent for all t. The model parameters $\lambda \triangleq \{p_m, p_{mn}, A(m), B(m), C(m), D(m), F(m), G(m), \hat{x}_0, P_0; \ m, n \in S\}$ are assumed known.

B. Estimation Objectives

Given at time t the observations $\mathbf{y}_{1:t}$, assuming that the model parameters λ are exactly known, all Bayesian inference for JMLS relies on the joint posterior distribution $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$, where $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) = p(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$. Given $\mathbf{r}_{1:t}$, $p(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})$ is a Gaussian distribution whose parameters can be evaluated using a Kalman filter. $p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$ could be computed exactly, but this discrete distribution has s^t values, and thus, some approximations have to be made as time increases. In this paper, we are interested in the following optimal estimation problems:

- *Filtering objectives:* Obtain the filtering distribution $p(r_t, x_t|\mathbf{y}_{1:t})$ as well as the MMSE estimate of $\varphi_{t|t}(r_t, x_t)$ given by $I(\varphi_{t|t}) \triangleq \mathbb{E}_{p(r_t, x_t|\mathbf{y}_{1:t})}(\varphi_{t|t}(r_t, x_t))$, where $\varphi_{t|t} : S \times \mathbb{R}^{n_\mathbf{x}} \to \mathbb{R}^{n_{\varphi_{t|t}}}$.
- *Fixed-lag smoothing objectives:* Obtain the fixed-lag distribution $p(r_t, x_t|\mathbf{y}_{1:t+L})$, where $L \in \mathbb{N}^*$, as well as the MMSE estimate of $\varphi_{t|t+L}(r_t, x_t)$ given by $I(\varphi_{t|t+L}) \triangleq \mathbb{E}_{p(r_t, x_t|\mathbf{y}_{1:t+L})}(\varphi_{t|t+L}(r_t, x_t))$, where $\varphi_{t|t+L} : S \times \mathbb{R}^{n_\mathbf{x}} \to \mathbb{R}^{n_{\varphi_{t|t+L}}}$.

We restrict ourselves to the common case where $\mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(\varphi_{t|t}(r_t, x_t))$ and $\mathbb{E}_{p(x_t|\mathbf{y}_{1:t+L}, \mathbf{r}_{1:t})}(\varphi_{t|t+L}(r_t, x_t))$ can be computed analytically.

Remark 1: In most filtering applications, we are interested in estimating the MMSE (conditional mean) state estimates $\mathbb{E}(x_t|\mathbf{y}_{1:t})$ and $\text{cov}(x_t|\mathbf{y}_{1:t})$. In these cases, $\mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(\varphi_{t|t}(r_t, x_t))$ can be computed analytically using the Kalman filter for the sequence $\mathbf{r}_{1:t}$.

III. SIMULATION-BASED OPTIMAL FILTER

This section is organized as follows. We describe the standard importance sampling method and then show how variance reduction can be achieved by integrating out the states $\mathbf{x}_{0:t}$ using the Kalman filter. Then, we present a sequential version of importance sampling for optimal filtering, generalizing the current approaches in the literature. We show why it is necessary to introduce a selection scheme, and then, we propose a generic Monte Carlo filter.

A. Monte Carlo Simulation for Optimal Estimation

For any $\varphi_{t|t}$, we will assume subsequently that $|I(\varphi_{t|t})| < +\infty$. If we were able to sample N i.i.d. random samples—called particles—$\{(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}); i = 1, \cdots, N\}$ according to $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$, then an empirical estimate of

this distribution would be given by

$$\overline{p_N}(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) = \frac{1}{N} \sum_{i=1}^{N} \delta_{\left(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}\right)}(d\mathbf{r}_{1:t}, d\mathbf{x}_{0:t})$$

and, as a corollary, an estimate of $p(r_t, x_t|\mathbf{y}_{1:t})$ is $\overline{p_N}(r_t, x_t|\mathbf{y}_{1:t}) = 1/N \sum_{i=1}^{N} \delta_{\left(r_t^{(i)}, x_t^{(i)}\right)}(dr_t, dx_t)$. From this distribution, one can easily obtain an estimate of $I(\varphi_{t|t})$ for any $\varphi_{t|t}$

$$\overline{I_N}(\varphi_{t|t}) = \int \varphi_{t|t}(r_t, x_t)\overline{p_N}(r_t, x_t|\mathbf{y}_{1:t}) \, dr_t dx_t$$

$$= \frac{1}{N} \sum_{i=1}^{N} \varphi_{t|t}\left(r_t^{(i)}, x_t^{(i)}\right).$$

This estimate is unbiased and, from the strong law of large numbers (SLLN), $\overline{I_N}(\varphi_{t|t})$ converges almost surely (a.s.) toward $I(\varphi_{t|t})$ as $N \to +\infty$. If $\sigma^2_{\varphi_{t|t}} \triangleq var_{p(r_t, x_t|\mathbf{y}_{1:t})}(\varphi_{t|t}(r_t, x_t)) < +\infty$, then a central limit theorem (CLT) holds

$$\sqrt{N}\left(\overline{I_N}\left(\varphi_{t|t}\right) - I\left(\varphi_{t|t}\right)\right) \underset{N \to +\infty}{\Rightarrow} \mathcal{N}\left(0, \sigma^2_{\varphi_{t|t}}\right)$$

where "$\Rightarrow$" denotes convergence in distribution. The advantage of the MC method is clear. One can easily estimate $I(\varphi_{t|t})$ for any $\varphi_{t|t}$, and the rate of convergence of this estimate does not depend on t. Unfortunately, it is impossible to sample efficiently from the "target" posterior distribution $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$ at any time t. Therefore, we focus on alternative methods.

A solution to estimate $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$ and $I(\varphi_{t|t})$ consists of using the well-known importance sampling method [5]. This method is based on the following remark. Let us introduce an arbitrary importance distribution $\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$, from which it is easy to obtain samples, and such that $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) > 0$ implies $\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) > 0$. Then

$$I\left(\varphi_{t|t}\right) = \frac{\mathbb{E}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}\left(\varphi_{t|t}(r_t, x_t)w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t})\right)}{\mathbb{E}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))}$$

where the importance weight is equal to

$$w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}) = \frac{p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}.$$

If we have N i.i.d. random samples $\{(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}); i = 1, \cdots, N\}$ distributed according to $\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$, then a Monte Carlo estimate of $I(\varphi_{t|t})$ is given by

$$\widehat{I_N^1} = \frac{\widehat{A_N^1}\left(\varphi_{t|t}\right)}{\widehat{B_N^1}\left(\varphi_{t|t}\right)}$$

$$= \frac{\displaystyle\sum_{i=1}^{N} \varphi_{t|t}\left(r_t^{(i)}, x_t^{(i)}\right) w\left(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}\right)}{\displaystyle\sum_{i=1}^{N} w\left(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}\right)}$$

$$= \sum_{i=1}^{N} \tilde{w}_{1:t}^{(i)} \varphi_{t|t}\left(r_t^{(i)}, x_t^{(i)}\right)$$

where the normalized importance weights $\tilde{w}_{1:t}^{(i)}$ are equal to

$$\tilde{w}_{1:t}^{(i)} = \frac{w\left(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}\right)}{\sum\limits_{j=1}^{N} w\left(\mathbf{r}_{1:t}^{(j)}, \mathbf{x}_{0:t}^{(j)}\right)}.$$

This method is equivalent to the following point mass approximation of $p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$:

$$\widehat{p_N}\left(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}\right) = \sum_{i=1}^{N} \tilde{w}_{1:t}^{(i)} \delta_{\left(\mathbf{r}_{1:t}^{(i)}, \mathbf{x}_{0:t}^{(i)}\right)}(d\mathbf{r}_{1:t}, d\mathbf{x}_{0:t})$$

and thus, $\widehat{p_N}(r_t, x_t|\mathbf{y}_{1:t}) = \sum_{i=1}^{N} \tilde{w}_{1:t}^{(i)} \delta_{\left(r_t^{(i)}, x_t^{(i)}\right)}(dr_t, dx_t)$. The "perfect" simulation case, i.e., $\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}) = p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})$, would correspond to $\tilde{w}_{1:t}^{(i)} = N^{-1}$ for any i. In practice, we will try to select the importance distribution as close as possible to the target distribution in a given sense. For N finite, $\widehat{I_N^1}(\varphi_{t|t})$ is biased (ratio of estimates), but asymptotically according to the SLLN, $\widehat{I_N}(\varphi_{t|t})$ converges a.s. toward $I(\varphi_{t|t})$. Under additional assumptions, a CLT also holds. However, we first show, in the next subsection, how the variance of the estimate $\widehat{I_N^1}(\varphi_{t|t})$ can be decreased.

B. Variance Reduction

It is possible to reduce the problem of estimating $p(r_t, x_t|\mathbf{y}_{1:t})$ and $I(\varphi_{t|t})$ to one of sampling from $p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$. Indeed, $p(\mathbf{r}_{1:t}, x_t|\mathbf{y}_{1:t}) = p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})$, where $p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})$ is a Gaussian distribution whose parameters are given by the Kalman filter. Thus, given an approximation of $p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$, ones gets straightforwardly an approximation of $p(r_t, x_t|\mathbf{y}_{1:t})$. Moreover, if $\mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(\varphi_{t|t}(r_t, x_t))$ can be evaluated in a closed-form expression, then the following alternative Bayesian importance sampling estimate of $I(\varphi_{t|t})$ can be proposed:

$$\widehat{I_N^2}(\varphi_{t|t}) = \frac{\widehat{A_N^2}(\varphi_{t|t})}{\widehat{B_N^2}(\varphi_{t|t})}$$

$$= \frac{\sum\limits_{i=1}^{N} \mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t}^{(i)})}\left(\varphi_{t|t}\left(r_t^{(i)}, x_t\right)\right) w\left(\mathbf{r}_{1:t}^{(i)}\right)}{\sum\limits_{i=1}^{N} w\left(\mathbf{r}_{1:t}^{(i)}\right)}$$

where

$$w(\mathbf{r}_{1:t}) = \frac{p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}$$

and

$$\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t}) = \int \pi\left(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t}\right) d\mathbf{x}_{0:t}.$$

Intuitively, to reach a given precision, $\widehat{I_N^2}(\varphi_{t|t})$ will require a reduced number N of samples over $\widehat{I_N^1}(\varphi_{t|t})$ as we only need to sample from a lower dimensional distribution. This is proven in the following propositions, where it is shown that if one can

integrate $\mathbf{x}_{0:t}$ analytically, then the variances of the resulting estimates are lower than the ones of the "crude" estimates.

Proposition 1: The variances of the importance weights, the numerators, and the denominators satisfy for any N

$$\begin{aligned}
&\mathrm{var}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t})) \\
&\quad - \mathrm{var}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}(w(\mathbf{r}_{1:t})) \\
&\quad = \mathbb{E}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left(\mathrm{var}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))\right) \\
&\mathrm{var}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}\left(\widehat{A_N^1}(\varphi_{t|t})\right) \\
&\quad - \mathrm{var}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left(\widehat{A_N^2}(\varphi_{t|t})\right) \\
&\quad = \mathbb{E}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left(\mathrm{var}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}\left(\widehat{A_N^1}(\varphi_{t|t})\right)\right) \\
&\mathrm{var}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}\left(\widehat{B_N^1}(\varphi_{t|t})\right) \\
&\quad - \mathrm{var}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left(\widehat{B_N^2}(\varphi_{t|t})\right) \\
&\quad = \mathbb{E}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left(\mathrm{var}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}\left(\widehat{B_N^1}(\varphi_{t|t})\right)\right). \quad (3)
\end{aligned}$$

A sufficient condition for $\widehat{I_N^1}(\varphi_{t|t})$ to satisfy a CLT is $\mathrm{var}_{p(r_t, x_t|\mathbf{y}_{1:t})}(\varphi_{t|t}(r_t, x_t)) < +\infty$ and $w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}) < C_t < +\infty$ for any $(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}) \in S^t \times (\mathbb{R}^{n_x})^{t+1}$ [5]. This trivially implies that $\widehat{I_N^2}(\varphi_{t|t})$ also satisfies a CLT. More precisely, we get the following result.

Proposition 2: Under the assumptions given above, $\widehat{I_N^1}(\varphi_{t|t})$ and $\widehat{I_N^2}(\varphi_{t|t})$ satisfy a CLT

$$\sqrt{N}\left(\widehat{I_N^1}(\varphi_{t|t}) - I(\varphi_{t|t})\right) \underset{N\to+\infty}{\Rightarrow} \mathcal{N}(0, \sigma_1^2)$$

$$\sqrt{N}\left(\widehat{I_N^2}(\varphi_{t|t}) - I(\varphi_{t|t})\right) \underset{N\to+\infty}{\Rightarrow} \mathcal{N}(0, \sigma_2^2)$$

where

$$\begin{aligned}
\sigma_1^2 &= \mathbb{E}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}\left((\varphi_{t|t}(r_t, x_t) - I(\varphi_{t|t}))^2 \right. \\
&\qquad \left. \cdot w^2(\mathbf{r}_{1:t}, \mathbf{x}_{0:t})\right) \\
\sigma_2^2 &= \mathbb{E}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left((\mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(\varphi_{t|t}(r_t, x_t)) \right. \\
&\qquad \left. - I(\varphi_{t|t}))^2 w^2(\mathbf{r}_{1:t})\right) \\
\sigma_1^2 - \sigma_2^2 &= \mathbb{E}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}\left(\mathrm{var}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})} \right. \\
&\qquad \left. \cdot ((\varphi_{t|t}(r_t, x_t) - I(\varphi_{t|t}))w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))\right).
\end{aligned}$$

Given these results, we now focus on importance sampling methods to get an approximation of $p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$ and $I(\varphi_{t|t})$ using an importance distribution $\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$. Up to now, the methods we have described are batch methods. We show in the next subsection how to obtain a sequential method.

C. Sequential Importance Sampling

One can always rewrite the importance function at time t as follows:

$$\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t}) = \pi(r_1|\mathbf{y}_{1:t}) \prod_{k=2}^{t} \pi(r_k|\mathbf{y}_{1:t}, \mathbf{r}_{1:k-1})$$

where $\pi(r_k|\mathbf{y}_{1:t}, \mathbf{r}_{1:k-1})$ is the probability distribution of r_k conditional upon $\mathbf{y}_{1:t}$ and $\mathbf{r}_{1:k-1}$. Our aim is to obtain, at any time t, an estimate of the distribution $p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$ and to be able to propagate this estimate in time without modifying, subsequently, the past simulated trajectories $\{\mathbf{r}_{1:t}^{(i)}; i = 1, \cdots, N\}$. This means that $\pi(\mathbf{r}_{1:t+1}|\mathbf{y}_{1:t+1})$ admits $\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$ as marginal distribution at time t. This is possible if we restrict ourselves to importance functions of the following form:

$$\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t}) = \pi(r_1|y_1) \prod_{k=2}^{t} \pi(r_k|\mathbf{y}_{1:k}, \mathbf{r}_{1:k-1}). \quad (4)$$

Such an importance function allows for a recursive evaluation of $w(\mathbf{r}_{1:t}) = w(\mathbf{r}_{1:t-1})w_t$ and, thus, of $\tilde{w}_{1:t}$, where the *incremental weight* w_t is given by

$$w_t = \frac{p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t})p(r_t|r_{t-1})}{p(y_t|\mathbf{y}_{1:t-1})\pi(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})}$$
$$\propto \frac{p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t})p(r_t|r_{t-1})}{\pi(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})}.$$

Further on, $\tilde{w}_t$ denotes the normalized version of w_t, i.e., $\tilde{w}_t^{(i)} = [\sum_{j=1}^{N} w_t^{(j)}]^{-1} w_t^{(i)}$.

1) Choice of the Importance Distribution: There are infinitely many possible choices for $\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$, the only condition being that its support includes the one of $p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$, that is, the support of $p(\mathbf{r}_{1:t})$. A sensible selection criterion is to choose a proposal that minimizes the variance of the importance weights at time t, given $\mathbf{r}_{1:t-1}$ and $\mathbf{y}_{1:t}$. According to this strategy, the following proposition establishes the optimal importance distribution.

Proposition 3: $p(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})$ is the distribution that minimizes the variance of the importance weights conditional on $\mathbf{r}_{1:t-1}$ and $\mathbf{y}_{1:t}$.

The proof is straightforward as one can easily check that the conditional variance is equal to zero in this case. We show how to implement this "optimal" distribution and then describe several suboptimal methods.

- *Optimal Sampling Distribution:* The optimal distribution satisfies

$$p(r_t = m|\mathbf{r}_{1:t-1}, \mathbf{y}_{1:t})$$
$$= \frac{p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t-1}, r_t = m)p(r_t = m|r_{t-1})}{p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t-1})}$$

and the associated importance weight w_t is proportional to

$$p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t-1})$$
$$= \sum_{m=1}^{s} p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t-1}, r_t = m)p(r_t = m|r_{t-1})$$
$$= \sum_{m=1}^{s} \Phi(\widetilde{y}_{t|t-1}(\mathbf{r}_{1:t-1}, r_t = m), S_t(\mathbf{r}_{1:t-1}, r_t = m))$$
$$\times p_{r_{t-1}, m}$$

where $\tilde{y}_{t|t-1}(\mathbf{r}_{1:t})$ and $S_t(\mathbf{r}_{1:t-1}, r_t = m)$ are, respectively, the innovation and the one-step-ahead prediction covariance of the observation conditional on $(\mathbf{r}_{1:t-1}, r_t = m)$. Computing $p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t-1})$ requires the evaluation of s one-step-ahead Kalman filter steps. It is thus computationally intensive if s is large.

- *Prior Distribution:* If we use the prior distribution $p(r_t|r_{t-1})$ as importance distribution, the importance weight is proportional to $p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t}) = \Phi(\tilde{y}_{t|t-1}(\mathbf{r}_{1:t}), S_t(\mathbf{r}_{1:t}))$. It only requires one step of a Kalman filter to be evaluated.
- *Alternative Sampling Distribution:* It is possible to design a variety of alternative sampling distributions. For example, one can use the results of a suboptimal deterministic algorithm to construct an importance sampling distribution.

2) Degeneracy of the Algorithm: The following proposition shows that for importance functions of the form (4), the variance of $w(\mathbf{r}_{1:t})$ can only increase (stochastically) over time. The proof of this proposition is an extension of a Kong–Liu–Wong [23, p. 285] theorem to the case of an importance function of the form (4) and is omitted here.

Proposition 4: The unconditional variance (i.e., with the observations $\mathbf{y}_{1:t}$ being interpreted as random variables) of the importance weights $w(\mathbf{r}_{1:t})$ increases over time.

It is thus impossible to avoid a degeneracy phenomenon. Practically, after a few iterations of the algorithm, all but one of the normalized importance weights are very close to zero, and a large computational burden is devoted to updating trajectories whose contribution to the final estimate is almost zero. That is why it is of crucial importance to introduce a selection step in the algorithm. The aim of this selection step is to discard the particles $\mathbf{r}_{1:t}^{(i)}$ with low normalized importance weights $\tilde{w}(\mathbf{r}_{1:t}^{(i)})$ and to multiply the ones with high $\tilde{w}(\mathbf{r}_{1:t}^{(i)})$ to avoid the degeneracy of the algorithm and to jump into the interesting zones of the space. Each time a selection step is used the weights are reset to N^{-1}.

D. Selection Step

A selection procedure associates with each particle, say $\tilde{\mathbf{r}}_{1:t}^{(i)}$ $(i = 1, \cdots, N)$, a number of "children" $N_i \in \mathbb{N}$, such that $\sum_{i=1}^{N} N_i = N$, to obtain N new particles $\mathbf{r}_{1:t}^{(i)}$. If $N_i = 0$, then $\tilde{\mathbf{r}}_{1:t}^{(i)}$ is discarded; otherwise, it has N_i "children" at time t. If we use a selection scheme *at each time step*, then before the selection scheme, we have a weighted distribution $\widehat{p_N}(\mathbf{r}_{1:t}|\mathbf{y}_{1:t}) = \sum_{i=1}^{N} \tilde{w}_t^{(i)} \delta_{\tilde{\mathbf{r}}_{1:t}^{(i)}}(d\mathbf{r}_{1:t})$, and after the selection step, we have

$$\widehat{p_N}(\mathbf{r}_{1:t}|\mathbf{y}_{1:t}) = N^{-1} \sum_{i=1}^{N} \delta_{\mathbf{r}_{1:t}^{(i)}}(d\mathbf{r}_{1:t}).$$

1) Some Selection Schemes: We describe here some selection schemes and show how to implement them in $O(N)$ iterations.

- *Sampling Importance Resampling (SIR)/Multinomial Sampling Procedure:* This procedure, which was introduced originally by Gordon *et al.* [15], is the most popular

one. One samples N times from $\widetilde{p_N}(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})$ to obtain $(\mathbf{r}_{1:t}^{(i)};\ i = 1, \cdots, N)$. This is equivalent to drawing jointly $(N_i;\ i = 1, \cdots, N)$ according to a multinomial distribution of parameters N and $\tilde{w}_t^{(i)}$. This algorithm has originally been implemented in $O(N \log N)$ operations [15]. In fact, it is possible to implement exactly the SIR procedure in $O(N)$ operations by noticing that it is possible to sample in $O(N)$ operations N i.i.d. variables uniformly distributed in $[0, 1]$ and *ordered*, i.e., $u_1 \leq u_2 \leq \cdots \leq u_N$ using a classical algorithm [9], [32, p. 96]. In this case, we have $\mathbb{E}(N_i) = N\tilde{w}_t^{(i)}$ and $\text{var}(N_i) = N\tilde{w}_t^{(i)}(1 - \tilde{w}_t^{(i)})$. However, as pointed out in [25], it is possible and better to use selection schemes with a reduced variance.

- *Residual Resampling* [25]: This procedure performs as follows. Set $\tilde{N}_i = \lfloor N\tilde{w}_t^{(i)} \rfloor$ and then perform a SIR procedure to select the remaining $\overline{N}_t = N - \sum_{i=1}^{N} \tilde{N}_i$ samples with the new weights $w_t'^{(i)} = (\tilde{w}_t^{(i)}N - \tilde{N}_i)/\overline{N}_t$; finally, add the results to the current $\tilde{N}_i$. In this case, we obtain $\mathbb{E}(N_i) = N\tilde{w}_t^{(i)}$, but $\text{var}(N_i) = \overline{N}_t w_t'^{(i)}(1 - w_t'^{(i)})$.

Recent theoretical results obtained in [8] suggest that it is not necessary to design unbiased selection schemes, i.e., we can have $\mathbb{E}[N_i] \neq N\tilde{w}_t^{(i)}$.

2) On the Use of a Selection Scheme: Two estimates of $I(\varphi_{t|t})$ can be proposed before $(\widetilde{I_N}(\varphi_{t|t}))$ and after $(\widehat{I_N}(\varphi_{t|t}))$ the selection scheme at time t, where

$$\widetilde{I_N}(\varphi_{t|t}) = \int \mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(\varphi_{t|t}(r_t, x_t))$$
$$\cdot \widetilde{p_N}(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})\, d\mathbf{r}_{1:t}$$
$$\widehat{I_N}(\varphi_{t|t}) = \int \mathbb{E}_{p(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(\varphi_{t|t}(r_t, x_t))$$
$$\cdot \widehat{p_N}(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})\, d\mathbf{r}_{1:t}.$$

Using the variance decomposition, it is straightforward to show that if the selection scheme used is unbiased, then

$$\text{var}\left(\widehat{I_N}(\varphi_{t|t})\right)$$
$$= \text{var}\left(\mathbb{E}\left(\widehat{I_N}(\varphi_{t|t})|N_i;\ i = 1, \cdots, N\right)\right)$$
$$\quad + \mathbb{E}\left(\text{var}\left(\widehat{I_N}(\varphi_{t|t})|N_i;\ i = 1, \cdots, N\right)\right)$$
$$= \text{var}\left(\widetilde{I_N}(\varphi_{t|t})\right) + \mathbb{E}\left(\text{var}\left(\widehat{I_N}(\varphi_{t|t})|N_i;\ i = 1, \cdots, N\right)\right)$$
$$\geq \text{var}\left(\widetilde{I_N}(\varphi_{t|t})\right).$$

Therefore, it is better to estimate $I(\varphi_{t|t})$ using $\widetilde{I_N}(\varphi_{t|t})$ as the selection scheme can only increase the variance of the estimate. However, as shown in [24] in a different framework that could be adapted to our case, it is worth resampling as it usually decreases the variance of the following estimates $I(\varphi_{k|k})$, where $k > t$.

E. Implementation Issues

Given at time $t - 1$, $N \in \mathbb{N}^*$ random samples $(\mathbf{r}_{1:t-1}^{(i)};\ i = 1, \cdots, N)$ distributed approximately according to $p(\mathbf{r}_{1:t-1}|\mathbf{y}_{1:t-1})$, the MC filter proceeds as follows at time t.

Particle Filter for JMLS

Sequential Importance Sampling Step
- For $i = 1, \cdots, N$, sample $\tilde{r}_t^{(i)} \sim \pi(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1}^{(i)})$, and set $\tilde{\mathbf{r}}_{1:t}^{(i)} \triangleq (\mathbf{r}_{1:t-1}^{(i)}, \tilde{r}_t^{(i)})$.
- For $i = 1, \cdots, N$, evaluate the importance weights up to a normalizing constant

$$w_t^{(i)} \propto \frac{p\left(y_t|\mathbf{y}_{1:t-1}, \tilde{\mathbf{r}}_{1:t}^{(i)}\right) p\left(\tilde{r}_t^{(i)}|\tilde{r}_{t-1}^{(i)}\right)}{\pi\left(\tilde{r}_t^{(i)}|\mathbf{y}_{1:t}, \tilde{\mathbf{r}}_{1:t-1}^{(i)}\right)}. \tag{5}$$

- For $i = 1, \cdots, N$, normalize the importance weights

$$\tilde{w}_t^{(i)} = \left[\sum_{j=1}^{N} w_t^{(j)}\right]^{-1} w_t^{(i)}. \tag{6}$$

Selection Step
- Multiply/Discard particles $(\tilde{\mathbf{r}}_{1:t}^{(i)};\ i = 1, \cdots, N)$ with respect to high/low normalized importance weights $\tilde{w}_t^{(i)}$ to obtain N particles $(\mathbf{r}_{1:t}^{(i)};\ i = 1, \cdots, N)$.

Clearly, the computational complexity of this algorithm at each iteration is $O(N)$. At first sight, it could appear that one needs to keep in memory the paths of all trajectories, that is, $(\tilde{\mathbf{r}}_{1:t}^{(i)};\ i = 1, \cdots, N)$. In this case, the storage requirements would increase linearly over time. Actually, under the standard assumption that $\pi(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})$ only depends on $\mathbf{r}_{1:t-1}$ via the set of low-dimensional sufficient statistics $m_{t|t-1}(\mathbf{r}_{1:t}^{(i)})$ and $P_{t|t-1}(\mathbf{r}_{1:t}^{(i)})$, this is the case for $p(r_t|r_{t-1})$ and $p(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})$; then, one only needs to keep in memory these statistics. Therefore, the storage requirements are still $O(N)$ and do not increase over time.

Remark 2: In the case where we use $\pi(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1}) = p(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})$, the importance weight $w_t \propto p(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t-1})$ does not depend on r_t. It is thus possible to perform the selection step before the sequential importance sampling step.

IV. SIMULATION-BASED OPTIMAL FIXED-LAG SMOOTHER

In many applications of interest, it is possible to wait for a fixed delay, say L, before performing estimation of the states at time t. Results can then be significantly improved while keeping an on-line estimation scheme. We show in this section that a direct application of the previous methodology to fixed-lag smoothing is not satisfactory if L is large and then propose an alternative method.

A. Some Strategies for Fixed-Lag Smoothing

1) Direct Extension: The proposed simulation-based filter can be theoretically, and straightforwardly, extended to fixed-lag

smoothing. At time $t + L$, we have an MC approximation

$$\widehat{p_N}(\mathbf{r}_{1:t+L}|\mathbf{y}_{1:t+L}) = \frac{1}{N} \sum_{i=1}^{N} \delta_{\mathbf{r}_{1:t+L}^{(i)}}(\mathbf{r}_{1:t+L})$$

of the distribution $p(\mathbf{r}_{1:t+L}|\mathbf{y}_{1:t+L})$. Therefore, an estimate of the marginal distribution is

$$\widehat{p_N}(\mathbf{r}_{1:t}|\mathbf{y}_{1:t+L}) = \frac{1}{N} \sum_{i=1}^{N} \delta_{\mathbf{r}_{1:t}^{(i)}}(\mathbf{r}_{1:t}).$$

However, the trajectories have been resampled L times, from time $t + 1$ to $t + L$. There is consequently a loss of diversity in the trajectories, i.e., one typically only has few distinct trajectories at time $t + L$, and the approximation of the posterior distribution might then be poor; this is the so-called problem of sample depletion [10].

2) Sampling Using $\mathbf{y}_{1:t+L}$*:* One solution to this problem consists of sampling r_t using an importance distribution based on the observations up to time $t + L$, i.e., introducing $r_t \sim \pi(r_t|\mathbf{y}_{1:t+L}, \mathbf{r}_{1:t-1})$. The same developments as in Section III-C apply. If one wants to minimize the variance of importance weights sequentially, conditional on $\mathbf{r}_{1:t-1}$ and $\mathbf{y}_{1:t+L}$, then $p(r_t|\mathbf{y}_{1:t+L}, \mathbf{r}_{1:t-1})$ is the optimal importance function. In doing so, the importance weights at time t are equal to

$$w_t \propto \frac{p(\mathbf{r}_{1:t}|\mathbf{y}_{1:t+L})}{p(\mathbf{r}_{1:t-1}|\mathbf{y}_{1:t+L-1})p(r_t|\mathbf{y}_{1:t+L}, \mathbf{r}_{1:t-1})}$$
$$\propto p(y_{t+L}|\mathbf{y}_{1:t+L-1}, \mathbf{r}_{1:t-1}).$$

Sampling from $p(r_t|\mathbf{y}_{1:t+L}, \mathbf{r}_{1:t-1})$ and evaluating $p(y_{t+L}|\mathbf{y}_{1:t+L-1}, \mathbf{r}_{1:t-1})$ requires s^{L+1} steps of the Kalman filter, each one being associated with all the possible trajectories $\mathbf{r}_{t:t+L}$ from time t to time $t + L$. This is computationally very expensive as $s^L \gg 1$.

3) Using MCMC Methods: We propose here an alternative and computationally cheaper method. This consists of simply adding a step to the simulation-based filter, and it drastically reduces the problem of sample depletion. At time $t + L$, the particles, say $\mathbf{r}_{1:t+L}^{\prime (i)}$, have typically very few distinct values $r_t^{\prime (i)}$ at time t as the trajectories have been resampled $L + 1$ times. Several suboptimal methods have been proposed in the literature to introduce diversity amongst the samples. They are mainly based on kernel density methods [12], [34]. Based on the current set of particles, the probability distribution is estimated using a kernel density estimate and a new set of distinct particles is sampled from it. These methods require the choice of a specific kernel [16] and are not suitable in our case as the distribution to estimate is discrete. Moreover, they introduce an additional MC variation. We propose here a totally different approach that consists of applying MCMC steps to the particles; see [33] for an introduction to MCMC. This idea has been proposed in [27] in a very different context.

Assume that the particles $\mathbf{r}_{1:t+L}^{\prime (i)}$ are distributed marginally according to $p(\mathbf{r}_{1:t+L}^{\prime}|\mathbf{y}_{1:t+L})$. Then, if we apply to each par-

ticle a Markov transition kernel $K(\mathbf{r}_{1:t+L}|\mathbf{r}_{1:t+L}^{\prime})$ of invariant distribution $p(\mathbf{r}_{1:t+L}|\mathbf{y}_{1:t+L})$, i.e., such that

$$\int K(\mathbf{r}_{1:t+L}|\mathbf{r}_{1:t+L}^{\prime})p(\mathbf{r}_{1:t+L}^{\prime}|\mathbf{y}_{1:t+L}) d\mathbf{r}_{1:t+L}^{\prime}$$
$$= p(\mathbf{r}_{1:t+L}|\mathbf{y}_{1:t+L})$$

then the new particles $\mathbf{r}_{1:t+L}^{(i)}$ are still distributed according to the posterior distribution of interest. Therefore, if $K(\mathbf{r}_{1:t+L}|\mathbf{r}_{1:t+L}^{\prime})$ is a kernel that updates stochastically $\mathbf{r}_{1:t+L}^{\prime (i)}$ to obtain $\mathbf{r}_{1:t+L}^{(i)}$, then we have a theoretically valid way of introducing diversity amongst the samples. It is possible to use all the standard MCMC methods such as the Metropolis–Hastings or the Gibbs samplers. However, contrary to classical MCMC methods, it is not necessary to introduce an ergodic transition kernel. This method can only improve results in the sense that it reduces the total variation norm of the current distribution of the particles with respect to the "target" posterior distribution; see [33] for example.

B. Implementation Issues

1) Algorithm: Given at time $t + L - 1$, $N \in \mathbb{N}^*$ random samples $(\mathbf{r}_{1:t+L-1}^{(i)}; i = 1, \cdots, N)$ distributed approximately according to $p(\mathbf{r}_{1:t+L-1}|\mathbf{y}_{1:t+L-1})$, the algorithm proceeds as follows at time $t + L$.

Particle Fixed-Lag Smoother for JMLS

Sequential Importance Sampling Step
- For $i = 1, \cdots, N$, sample $\tilde{r}_{t+L}^{(i)} \sim \pi(r_{t+L}|\mathbf{y}_{1:t+L}, \mathbf{r}_{1:t-1}^{(i)})$ and $\tilde{\mathbf{r}}_{1:t+L}^{(i)} \triangleq (\mathbf{r}_{1:t+L-1}^{(i)}, \tilde{r}_{t+L}^{(i)})$.
- For $i = 1, \cdots, N$, evaluate the normalized importance weights $\tilde{w}_{t+L}^{(i)}$ using (5) and (6).

Selection Step
- Multiply/Discard particles $(\tilde{\mathbf{r}}_{1:t+L}^{(i)}; i = 1, \cdots, N)$ with respect to high/low normalized importance weights $\tilde{w}_{t+L}^{(i)}$ to obtain N particles $(\mathbf{r}_{1:t+L}^{\prime (i)}; i = 1, \cdots, N)$.

MCMC Step
- For $i = 1, \cdots, N$, apply to $\mathbf{r}_{1:t+L}^{\prime (i)}$ a Markov transition kernel $K(\mathbf{r}_{1:t+L}^{(i)}|\mathbf{r}_{1:t+L}^{\prime (i)})$ of invariant distribution $p(\mathbf{r}_{1:t+L}|\mathbf{y}_{1:t+L})$ to obtain N particles $(\mathbf{r}_{1:t+L}^{(i)}; i = 1, \cdots, N)$.

---

2) Implementation of MCMC Steps: There is an infinity of possible choices for the MCMC transition kernel. We propose here to use a Gibbs sampling step that updates at time $t + L$ the values of the Markov chain from time t to $t + L$, i.e., we sample $r_k^{(i)}$ for $k = t, \cdots, t + L$ according to $p(r_k|\mathbf{y}_{1:t+L}, \mathbf{r}_{-k}^{(i)})$, where

$$\mathbf{r}_{1:t-1}^{(i)} = \mathbf{r}_{1:t-1}^{\prime (i)}, \text{ and } r_{-k}^{(i)}$$
$$\triangleq \left(\mathbf{r}_{1:t-1}^{(i)}, r_t^{(i)}, \ldots, r_{k-1}^{(i)}, r_{k+1}^{\prime (i)} \cdots, r_{t+l}^{\prime (i)}\right).$$

It is straightforward to check that this algorithm admits $p(\mathbf{r}_{1:t+L}|\mathbf{y}_{1:t+L})$ as invariant distribution [33]. Sampling from $p(r_k|\mathbf{y}_{1:t+L}, \mathbf{r}_{-k}^{(i)})$ for $k = t, \cdots, t + L$ can be done by the following backward–forward procedure of complexity

$O(L + 1)$. This algorithm has been developed in a batch framework in [9] and [11]; therefore, the proofs are omitted here. It proceeds as follows at time $t + L$ for the particle i.

Backward-Forward Procedure

Backward Step

For $k = t + L, \cdots, t$, compute and store $P'^{-1}_{k|k+1}(\mathbf{r}'^{(i)}_{k+1:t+L})$ and $P'^{-1}_{k|k+1}(\mathbf{r}'^{(i)}_{k+1:t+L}) m'_{k|k+1}(\mathbf{r}'^{(i)}_{k+1:t+L})$ using (8) and (9) given in Appendix B.

Forward Step

For $k = t, \cdots, t + L$

• For $m = 1, \cdots, s$, run one step ahead the Kalman filter with $r_k = m$, store $m_{k|k}(\mathbf{r}^{(i)}_{1:k-1}, r_k = m)$ and $P_{k|k}(\mathbf{r}^{(i)}_{1:k-1}, r_k = m)$, and then compute $p(r_k = m | \mathbf{y}_{1:t+L}, \mathbf{r}^{(i)}_{-k})$ using (7) given below.

• Sample $r^{(i)}_k \sim p(r_k | \mathbf{y}_{1:t+L}, \mathbf{r}^{(i)}_{-k})$, and store $m_{k|k}(\mathbf{r}^{(i)}_{1:k-1}, r^{(i)}_k)$ and $P_{k|k}(\mathbf{r}^{(i)}_{1:k-1}, r^{(i)}_k)$.

The quantities $P'^{-1}_{k|k+1}(\mathbf{r}'^{(i)}_{k+1:t+L})$ and $P'^{-1}_{k|k+1}(\mathbf{r}'^{(i)}_{k+1:t+L}) m'_{k|k+1}(\mathbf{r}'^{(i)}_{k+1:t+L})$ are given by the backward information filter recursion given in Appendix B, and for any $k = t, \cdots, t + L$, we have

$$
\begin{aligned}
& p(r_k | \mathbf{y}_{1:t+L}, \mathbf{r}_{-k}) \\
& \propto p_{r_{k-1} r_k} p_{r_k r_{k+1}} N(\tilde{y}_{k|k-1}(\mathbf{r}_{1:k}), S_k(\mathbf{r}_{1:k})) \\
& \times \left| \tilde{\Pi}_{k|k}(\mathbf{r}_{1:k}) \tilde{Q}^{\mathsf{T}}_{k|k}(\mathbf{r}_{1:k}) \right. \\
& \quad \left. P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L}) \tilde{Q}_{k|k}(\mathbf{r}_{1:k}) + I_{n_k} \right|^{-1/2} \\
& \times \exp\left(-\frac{1}{2} \left[m^{\mathsf{T}}_{k|k}(\mathbf{r}_{1:k}) P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L}) m_{k|k}(\mathbf{r}_{1:k}) \right. \right. \\
& \quad - 2 m^{\mathsf{T}}_{k|k}(\mathbf{r}_{1:k}) P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L}) m'_{k|k+1}(\mathbf{r}_{k+1:t+L}) \\
& \quad - (m'_{k|k+1}(\mathbf{r}_{k+1:t+L}) - m_{k|k}(\mathbf{r}_{1:k}))^{\mathsf{T}} \\
& \quad \times P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L}) R_{k|k}(\mathbf{r}_{1:t+L}) P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L}) \\
& \quad \left. \left. \times \left(m'_{k|k+1}(\mathbf{r}_{k+1:t+L}) - m_{k|k}(\mathbf{r}_{1:k}) \right) \right] \right) \quad (7)
\end{aligned}
$$

where $P_{k|k}(\mathbf{r}_{1:k}) = \tilde{Q}_{k|k}(\mathbf{r}_{1:k}) \tilde{\Pi}_{k|k}(\mathbf{r}_{1:k}) \tilde{Q}^{\mathsf{T}}_{k|k}(\mathbf{r}_{1:k})$,[1] $\tilde{\Pi}_{k|k}(\mathbf{r}_{1:k})$ being a $n_k \times n_k$, $1 \leq n_k \leq n_x$ diagonal matrix with the nonzero eigenvalues of $P_{k|k}(\mathbf{r}_{1:k})$ as elements and

$$
\begin{aligned}
R_{k|k}(\mathbf{r}_{1:t+L}) = \tilde{Q}_{k|k}(\mathbf{r}_{1:k}) & \\
\times \Big[\tilde{\Pi}^{-1}_{k|k}(\mathbf{r}_{1:k}) & + \tilde{Q}^{\mathsf{T}}_{k|k}(\mathbf{r}_{1:k}) \\
& \times P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L}) \tilde{Q}_{k|k}(\mathbf{r}_{1:k}) \Big]^{-1} \\
& \times \tilde{Q}^{\mathsf{T}}_{k|k}(\mathbf{r}_{1:k}).
\end{aligned}
$$

The computational complexity of the resulting MC fixed-lag smoother algorithm at each iteration is $O((L + 1)N)$, and one needs to keep in memory the paths of all trajectories $(\mathbf{r}^{(i)}_{t:t+L}; i = 1, \cdots, N)$ as well as $(m_{k|k}(\mathbf{r}^{(i)}_{1:k}), P_{k|k}(\mathbf{r}^{(i)}_{1:k}); i = 1, \cdots, N)$ for $k = t, \cdots, t + L$.

---

[1] $\tilde{Q}_{k|k}(\mathbf{r}_{1:k})$ and $\tilde{\Pi}_{k|k}(\mathbf{r}_{1:k})$ are straightforwardly obtained using the singular value decomposition of $P_{k|k}(\mathbf{r}_{1:k})$.

V. DISCUSSION

Numerous methods have been proposed earlier in the literature to address the problems posed in this paper. In this section, we discuss these approaches and compare them with ours. To obtain MMSE filtered estimates, most algorithms are based on deterministic finite Gaussian mixtures approximations with a few components such as the popular IMM algorithm [4], [29] or the detection/estimation algorithm (DEA) [37]; see, for example, [18] and [20] for related methods. These algorithms are computationally cheaper than the ones we present here, but, for example, the IMM can fail in difficult situations such as when the likelihood at time t is multimodal, and the posterior distribution at time $t - 1$ is vague. Taking advantage of the increasing computational power available and of the intrinsic parallelizability of particle filters, several recent papers have proposed to use the bootstrap filter to perform optimal filtering of JMLS; see [3], [21], and [22] and an improved version of it based on kernel density estimation [16]. It has been shown by Monte Carlo simulations that these methods outperform the classical suboptimal methods. However, the bootstrap filter consists of sampling, at time t, the states (x_t, r_t) according to $p(x_t | r_t, x_{t-1}) p(r_t | r_{t-1})$. r_t and x_t are simulated independently of y_t; therefore, this strategy is sensitive to outliers, and the distribution of the importance weights can be highly skewed. In our scheme, the continuous states x_t are integrated out, which leads to estimates with a reduced variance, and the importance distribution can be chosen on a case-by-case basis. In the bootstrap filter, the particles are selected using the SIR algorithm. We showed that it is possible to use a quicker selection scheme having a reduced variance. Some previous works have, however, already introduced some of these ideas. The closest filtering algorithm to the one presented in this paper is the random sampling algorithm (RSA) of Akashi and Kumamoto [1], [37]. This algorithm corresponds to the sequential importance sampling method using $p(r_t | \mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})$ as importance distribution. This distribution was introduced following a different criterion. Similar work was developed later in [6] and [35]. However, these authors neither presented the general sequential importance sampling framework allowing the use of general importance distributions nor discussed the improvements brought by variance reduction. Moreover, the key selection step is omitted in these papers, and thus, after a few time steps, the algorithm collapses. It is worth noticing that Tugnait [37] compared, via Monte Carlo simulations, the RSA algorithm and some computationally cheaper alternative algorithms like the DEA. It appeared that the results of the RSA algorithm were less precise than those of the alternative methods. Nevertheless, in this early work, only a small number N of particles and no selection scheme were used. In our Monte Carlo simulations, the obtained results are totally different.

Finally, we are not aware of any paper combining particle filters and MCMC for JMLS. In a batch framework, some algorithms have been proposed to estimate $p(r_k | \mathbf{y}_{1:T}, \mathbf{r}_{-k})$ for $k = 1, \cdots, T$ to maximize $p(\mathbf{r}_{1:T} | \mathbf{y}_{1:T})$ using a coordinate ascent method. However, the popular single most likely replacement (SMLR) algorithm [30] has a complexity $O(T^2)$, whereas our method applied to this case has a complexity $O(T)$.
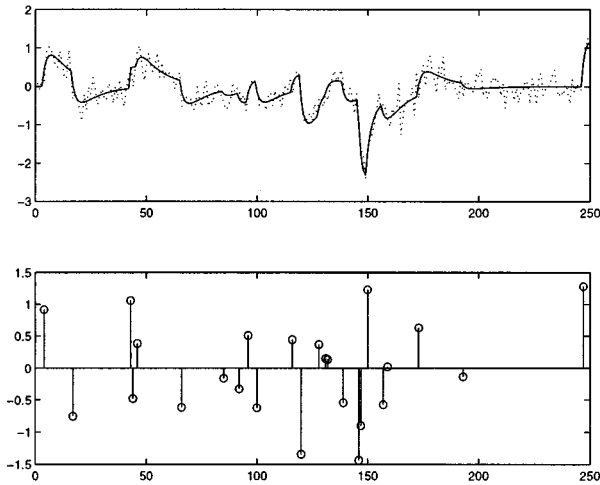
Fig. 1. Top: simulated signal x_t (solid line) and observations y_t (dotted line). Bottom: simulated sequence v'_t.

Fig. 2. Top: estimated probabilities $p(r_t = 1|\mathbf{y}_{1:t})$ (Particle filter). Middle: estimated probabilities $p(r_t = 1|\mathbf{y}_{1:t+L})$ (for $t = 1, \cdots, 235$) (Particle fixed-lag smoother). Bottom: estimated probabilities $p(r_t = 1|\mathbf{y}_{1:T})$ (Gibbs sampler).

VI. SIMULATIONS

Computer simulations were carried out to evaluate the performance of our algorithms. Section VI-A considers the problem of estimating a sparse signal based on a set of noisy data. Section VI-B considers the problem of tracking a maneuvering target.

A. Detection of Bernoulli–Gaussian Processes

In several problems related to seismic signal processing and nuclear science [7], [30], the signal of interest can be modeled as the output of a linear filter excited by a BG process and observed in white Gaussian noise. The input sequence is $v'_t \overset{\text{i.i.d}}{\sim} \lambda \mathcal{N}(0, \sigma_v^2) + (1 - \lambda)\delta_0$, $0 < \lambda < 1$, and the observation noise is $\varepsilon'_t \overset{\text{i.i.d}}{\sim} \mathcal{N}(0, \sigma_w^2)$. v'_t and ε'_t are mutually independent sequences. The linear filter is modeled by an AR(2) model. Thus, we have $S = \{1, 2\}$, and the signal admits the following state-space model:

$$A = \begin{pmatrix} a_1 & a_2 \\ 1 & 0 \end{pmatrix}, \quad C = \begin{pmatrix} 1 & 0 \end{pmatrix}$$
$$D = \sigma_w, \quad F = 0, \quad G = 0$$
$$B(1) = \begin{pmatrix} \sigma_v & 0 \end{pmatrix}^{\text{T}}, \quad B(2) = \begin{pmatrix} 0 & 0 \end{pmatrix}^{\text{T}}.$$

In the following simulations, we set the parameters to $a_1 = 1.51$, $a_2 = -0.55$, $\sigma_v = 0.50$, and $\sigma_w = 0.25$. $T = 250$ observations are generated and are depicted in Fig. 1. In Fig. 2, we present the results obtained using the particle filtering algorithms with $N = 500$ particles. We sample from the optimal distribution $p(r_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t-1})$ and perform fixed-lag smoothing with $L = 15$. The results are compared with an off-line iterative Gibbs sampling method to compute $p(r_t|\mathbf{y}_{1:T})$ [11]. Fixed-lag smoothing significantly improves the detection of occurrences with respect to filtering. Moreover, as shown in Fig. 2, the performance of the fixed-lag smoothing and batch methods appears to be very similar. Our studies show that the results obtained using a higher number of particles are indistinguishable.
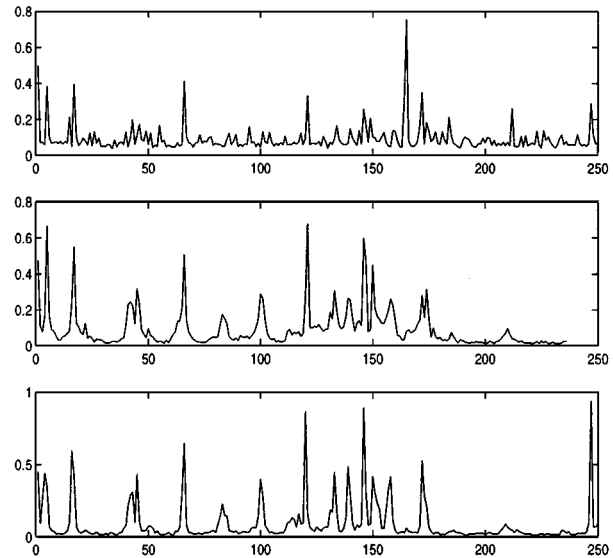
B. Tracking of a Maneuvering Target

We address the problem of tracking a maneuvering target in noise. The difficulty in this problem arises from the uncertainty in the maneuvering command driving the target. The state of the target at time t is denoted as $x_t \overset{\Delta}{=} (l_{x,t}, s_{x,t}, l_{y,t}, s_{y,t})^{\text{T}}$, where $l_{x,t}$ ($l_{y,t}$) and $s_{x,t}$ ($s_{y,t}$) represent the position and velocity of the target in the x (resp. in the y) direction. It evolves according to a JMLS model of parameters [4]

$$A = \begin{pmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad B = 0.1I_4,$$
$$C = I_4, \quad G = 0_{4 \times n_u}$$

and $D = \sqrt{3}\,\text{diag}(20, 1, 20, 1)$. The switching term is $F(r_t)u_t$, where r_t is a three-state Markov chain corresponding to the three possible maneuver commands:

- straight;
- left turn;
- right turn.

It has the following transition probabilities: $p_{m,m} = 0.9$ and $p_{m,n} = 0.05$ for $m \neq n$. We have for any t $F(1)u_t = (0, 0, 0, 0)^{\text{T}}$, $F(2)u_t = (-1.225, -0.35, 1.225, 0.35)^{\text{T}}$, and $F(3)u_t = (1.225, 0.35, -1.225, -0.35)^{\text{T}}$. We sample according to the optimal distribution. In Fig. 3, we display a realization of the signal and the MMSE estimate $(\mathbb{E}(l_{x,t}|y_{1:t}), \mathbb{E}(l_{y,t}|y_{1:t}))$ computed using $N = 500$ particles.

We also perform $M = 100$ different measurement realizations and compare our results with the IMM algorithm [29] and the standard bootstrap filter [15]. The performance measure is the root mean square (RMS) position error computed as follows
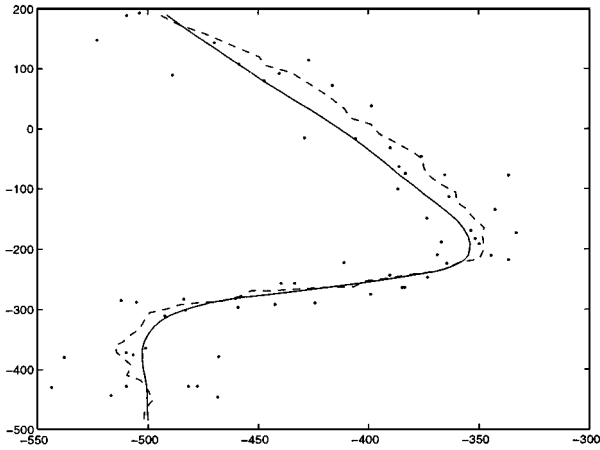
Fig. 3. Simulated path of the target $(l_{x,t}, l_{y,t})$ (solid line), observations (dots), and MMSE estimate (dashed line).

TABLE I
RMS FOR IMM, BOOTSTRAP FILTER, AND MC FILTER

Algorithm/N	50	100	250	500	1000	2500	5000
IMM Filter	24.69	-	-	-	-	-	-
Bootstrap Filter	26.22	24.76	24.02	23.88	23.45	23.11	22.79
MC Filter	22.95	22.74	22.69	22.64	22.62	22.63	22.62

from the MMSE estimates with respect to the true simulated trajectories

$$(\text{RMS})^2 = \frac{1}{MT} \sum_{m=1}^{M} \sum_{t=1}^{T} \left((l_{x,t} - l_{x,t}^{\text{MMSE}}(m))^2 + (l_{y,t} - l_{y,t}^{\text{MMSE}}(m))^2 \right)$$

where $l_{x,t}^{MMSE}(m)$ is the MMSE target position estimate in the x direction at time t of the mth Monte Carlo simulation. We present in Table I the performance of the IMM filter, the bootstrap filter, and our MC filter for various number of particles; note that the IMM filter is an algorithm with a computational complexity independent of N. The MC filter is more precise than the other methods. Its performance stabilizes quickly as N increases.

VII. CONCLUSION

In this paper, we presented on-line simulation-based algorithms to perform optimal filtering and fixed-lag smoothing of JMLS. These Monte Carlo algorithms are based on several efficient variance reduction methods. Although these algorithms are computationally intensive, they can be straightforwardly implemented on parallel computers allowing for real-time applications. Two applications were presented to illustrate the performance of these algorithms for online deconvolution of Bernoulli–Gaussian processes and tracking of a maneuvering target. We showed in simulations that these methods outper-

form not only the standard suboptimal methods but recent simulation-based methods as well.

Throughout this paper, the model parameters λ are assumed known. It is, however, possible to perform batch and on-line maximum likelihood estimation of these parameters combining the particle filtering methods developed here and standard expectation-maximization or gradient-type algorithms; see [19] for details.

APPENDIX A
KALMAN FILTER AND BACKWARD INFORMATION FILTER

Conditional on $\mathbf{r}_{1:t}$, the system (1) and (2) is linear Gaussian until t; it is thus possible to compute, using the Kalman filter, the one-step-ahead prediction and covariance of x_t $(m_{t|t-1}(\mathbf{r}_{1:t}) \triangleq \mathbb{E}(x_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t})$, and $P_{t|t-1}(\mathbf{r}_{1:t}) \triangleq cov(x_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t}))$, the filtered estimate and covariance of x_t $(m_{t|t}(\mathbf{r}_{1:t}) \triangleq \mathbb{E}(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})$, and $P_{t|t}(\mathbf{r}_{1:t}) \triangleq cov(x_t|\mathbf{y}_{1:t}, \mathbf{r}_{1:t}))$, the innovation at time t, and the covariance of this innovation $(\tilde{y}_{t|t-1}(\mathbf{r}_{1:t}) = y_t - \mathbb{E}(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t})$, and $S_t(\mathbf{r}_{1:t}) \triangleq cov(y_t|\mathbf{y}_{1:t-1}, \mathbf{r}_{1:t}))$.

The backward information filter proceeds as follows from time $t + L$ to t:

$$P'^{-1}_{t+L|t+L}(r_{t+L})$$
$$= C^{\mathrm{T}}(r_{t+L})(D(r_{t+L})D^{\mathrm{T}}(r_{t+L}))^{-1}C(r_{t+L})$$
$$P'^{-1}_{t+L|t+L}(r_{t+L})m'_{t+L|t+L}(r_{t+L})$$
$$= C^{\mathrm{T}}(r_{t+L})(D(r_{t+L})D^{\mathrm{T}}(r_{t+L}))^{-1}$$
$$\times (y_{t+L} - G(r_{t+L})u_{t+L}) \qquad (8)$$

and for $k = t + L - 1, \cdots, t$

$$\mathbf{\Delta}_{k+1}(\mathbf{r}_{k+1:t+L})$$
$$= \left[I_{n_v} + B^{\mathrm{T}}(r_{k+1})P'^{-1}_{k+1|k+1}(\mathbf{r}_{k+1:t+L})B(r_{k+1}) \right]^{-1}$$
$$P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L})$$
$$= A^{\mathrm{T}}(r_{k+1})P'^{-1}_{k+1|k+1}(\mathbf{r}_{k+1:t+L})$$
$$\times (I_{n_x} - B(r_{k+1})\mathbf{\Delta}_{k+1}(\mathbf{r}_{k+1:t+L})$$
$$\times B^{\mathrm{T}}(r_{k+1})P'^{-1}_{k+1|k+1}(\mathbf{r}_{k+1:t+L}))A(r_{k+1})$$
$$P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L})m'_{k|k+1}(\mathbf{r}_{k+1:t+L})$$
$$= A^{\mathrm{T}}(r_{k+1})(I_{n_x} - P'^{-1}_{k+1|k+1}(\mathbf{r}_{k+1:t+L})$$
$$\times B(r_{k+1})\mathbf{\Delta}_{k+1}(\mathbf{r}_{k+1:t+L})B^{\mathrm{T}}(r_{k+1}))$$
$$\times P'^{-1}_{k+1|k+1}(\mathbf{r}_{k+1:t+L})$$
$$\times (m'_{k+1|k+1}(\mathbf{r}_{k+1:t+L}) - F(r_{k+1})u_{k+1})$$
$$P'^{-1}_{k|k}(\mathbf{r}_{k:t+L})$$
$$= P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L})$$
$$+ C^{\mathrm{T}}(r_k)(D(r_k)D^{\mathrm{T}}(r_k))^{-1}C(r_k)$$
$$P'^{-1}_{k|k}(\mathbf{r}_{k:t+L})m'_{k|k}(\mathbf{r}_{k:t+L})$$
$$= P'^{-1}_{k|k+1}(\mathbf{r}_{k+1:t+L})m'_{k|k+1}(\mathbf{r}_{k+1:t+L})$$
$$+ C^{\mathrm{T}}(r_k)(D(r_k)D^{\mathrm{T}}(r_k))^{-1}(y_k - G(r_k)u_k). \qquad (9)$$

$$\mathrm{var}_\pi\left(\frac{\widehat{A_N^1}(\varphi_{t|t})}{\widehat{B_N^1}(\varphi_{t|t})}\right) = \frac{I(\varphi_{t|t})^2\mathrm{var}_\pi(\widehat{B_N^1}(\varphi_{t|t})) + \mathrm{var}_\pi(\widehat{A_N^1}(\varphi_{t|t})) - 2I(\varphi_{t|t})\mathrm{cov}_\pi(\widehat{A_N^1}(\varphi_{t|t}), \widehat{B_N^1}(\varphi_{t|t}))}{N^2} + O(N^{-3/2})$$

$$= \frac{\mathrm{var}_\pi((\varphi_{t|t}(r_t, x) - I(\varphi_{t|t}))w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))}{N} + O(N^{-3/2})$$

APPENDIX B
PROOFS OF PROPOSITIONS

Proof of Proposition 1: By the variance decomposition, one has

$$\mathrm{var}_{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))$$
$$= \mathrm{var}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}(\mathbb{E}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t})))$$
$$+ \mathbb{E}_{\pi(\mathbf{r}_{1:t}|\mathbf{y}_{1:t})}(\mathrm{var}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))).$$

However

$$\mathbb{E}_{\pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})}(w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))$$
$$= \int \frac{p(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}{\pi(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}|\mathbf{y}_{1:t})}$$
$$\times \pi(\mathbf{x}_{0:t}|\mathbf{y}_{1:t}, \mathbf{r}_{1:t})d\mathbf{x}_{0:t} = w(\mathbf{r}_{1:t}).$$

We obtain similar results for $\widehat{A_N^1}(\varphi_{t|t})$, $\widehat{A_N^2}(\varphi_{t|t})$, $\widehat{B_N^1}(\varphi_{t|t})$, and $\widehat{B_N^2}(\varphi_{t|t})$, and (3) follows.

Proof of Proposition 2: As $\widehat{A_N^1}(\varphi_{t|t})$ and $\widehat{B_N^1}(\varphi_{t|t})$ are sums of N i.i.d. random variables, we obtain, using the delta method

$$\mathrm{var}_\pi\left(\frac{\widehat{A_N^1}(\varphi_{t|t})}{\widehat{B_N^1}(\varphi_{t|t})}\right)$$
$$= \frac{\mathbb{E}_\pi(\widehat{A_N^1}(\varphi_{t|t}))^2\mathrm{var}_\pi(\widehat{B_N^1}(\varphi_{t|t}))}{\mathbb{E}_\pi(\widehat{B_N^1}(\varphi_{t|t}))^4}$$
$$+ \frac{\mathrm{var}_\pi(\widehat{A_N^1}(\varphi_{t|t}))}{\mathbb{E}_\pi(\widehat{B_N^1}(\varphi_{t|t}))^2}$$
$$- 2\frac{\mathbb{E}_\pi(\widehat{A_N^1}(\varphi_{t|t}))\mathrm{cov}_\pi(\widehat{A_N^1}(\varphi_{t|t}), \widehat{B_N^1}(\varphi_{t|t}))}{\mathbb{E}_\pi(\widehat{B_N^1}(\varphi_{t|t}))^3}$$
$$+ O(N^{-3/2})$$

where $\mathbb{E}_\pi(\widehat{A_N^1}(\varphi_{t|t})) = N\mathbb{E}_p(\varphi_{t|t}(r_t, x_t)) = NI(\varphi_{t|t})$, and $\mathbb{E}_\pi(\widehat{B_N^1}(\varphi_{t|t})) = N$. Thus, as it is shown in the equation at the top of the page, one obtains $\mathbb{E}_\pi((\varphi_{t|t}(r_t, x_t) - I(\varphi_{t|t}))w(\mathbf{r}_{1:t}, \mathbf{x}_{0:t})) = 0$; therefore

$$\mathrm{var}_\pi\left(\frac{\widehat{A_N^1}(\varphi_{t|t})}{\widehat{B_N^1}(\varphi_{t|t})}\right)$$
$$= \frac{\mathbb{E}_\pi((\varphi_{t|t}(r_t, x_t) - I(\varphi_{t|t}))^2 w^2(\mathbf{r}_{1:t}, \mathbf{x}_{0:t}))}{N}$$
$$+ O(N^{-3/2}).$$

The expression of σ_1^2 follows, similarly we obtain σ_2^2. Finally, one obtains $\sigma_1^2 - \sigma_2^2$ using the variance decomposition.

REFERENCES

[1] H. Akashi and H. Kumamoto, "Random sampling approach to state estimation in switching environments," *Automatica*, vol. 13, pp. 429–434, 1977.

[2] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall, 1979.

[3] D. Avitzour, "A stochastic simulation Bayesian approach to multitarget tracking," *Proc. Inst. Elect. Eng., Radar, Sonar, Navig.*, vol. 142, no. 2, pp. 41–44, 1995.

[4] Y. Bar-Shalom and X. R. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*: YBS, 1995.

[5] J. M. Bernardo and A. F. M. Smith, *Bayesian Theory*. New York: Wiley, 1994.

[6] M. Billio and A. Monfort, "Switching state-space models: Likelihood function, filtering and smoothing," *J. Stat. Planning Inform.*, vol. 68, pp. 65–103, 1998.

[7] O. Cappé, É. Moulines, and M. Lavielle, "Simulation-based methods for blind maximum-likelihood filter identification," *Signal Process.*, vol. 73, no. 1, pp. 3–25, 1999.

[8] D. Crisan, P. Del Moral, and T. Lyons, "Discrete filtering using branching and interacting particle systems," *Markov Processes Related Fields*, vol. 5, no. 3, pp. 293–318, 1999.

[9] A. Doucet, "Monte Carlo Methods for Bayesian Estimation of Hidden Markov Models," Ph.D. dissertation (in French), Univ. Paris-Sud, Orsay, France, 1997.

[10] A. Doucet, A. Logothetis, and V. Krishnamurthy, "Stochastic sampling algorithms for state estimation of jump Markov linear systems," *IEEE Trans. Automat. Contr.*, vol. 45, pp. 188–201, Feb. 2000.

[11] A. Doucet and C. Andrieu, "Iterative algorithms for optimal state estimation of jump Markov linear systems," *Proc. Conf. IEEE ICASSP*, 1999.

[12] A. Doucet, J. F. G. de Freitas, and N. J. Gordon, Eds., *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag, 2001.

[13] A. Doucet, S. J. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statist. Comput.*, vol. 10, no. 3, pp. 197–208, 2000.

[14] A. Doucet, N. J. Gordon, and V. Krishnamurthy, "Particle filtering algorithms for state estimation of jump Markov linear systems," Cambridge University, http://www-sigproc.eng.cam.ac.uk/~ad2/tr359.ps, 1999.

[15] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *Proc. Inst. Elect. Eng. F*, vol. 1400, no. 2, pp. 107–113, 1993.

[16] N. J. Gordon, "A hybrid bootstrap filter for target tracking in clutter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, pp. 353–358, Jan. 1997.

[17] J. E. Handschin and D. Q. Mayne, "Monte Carlo techniques to estimate the conditional expectation in multi-stage nonlinear filtering," *Int. J. Contr.*, vol. 9, no. 5, pp. 547–559, 1969.

[18] R. E. Helmick, W. D. Blair, and S. A. Hoffman, "Fixed-interval smoothing for Markovian switching systems," *IEEE Trans. Inform. Theory*, vol. 41, pp. 1845–1855, Nov. 1995.

[19] M. Hürzeler and H. R. Künsch, "Approximating and maximizing the likelihood for a general state space model," in *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag, 2001.

[20] G. Kitagawa, "The two-filter formula for smoothing and an implementation of the Gaussian-sum smoother," *Ann. Inst. Stat. Math.*, vol. 46, pp. 605–623, 1994.

[21] ——, "Monte Carlo filter and smoother for non-Gaussian nonlinear state space models," *J. Comput. Graph. Stat.*, vol. 5, no. 1, pp. 1–25, 1996.

[22] G. Kitagawa and W. Gersch, *Smoothness Priors Analysis of Time Series*. Berlin, Germany: Springer, 1996, vol. 116.

[23] A. Kong, J. S. Liu, and W. H. Wong, "Sequential imputations and Bayesian missing data problems," *J. Amer. Stat. Assoc.*, vol. 89, pp. 278–288, 1994.

[24] J. S. Liu and R. Chen, "Blind deconvolution via sequential imputation," *J. Amer. Stat. Assoc.*, vol. 90, pp. 567–576, 1995.

[25] ——, "Sequential Monte Carlo methods for dynamic systems," *J. Amer. Stat. Assoc.*, vol. 93, pp. 1032–1044, 1998.

[26] A. Logothetis and V. Krishnamurthy, "Expectation-maximization algorithms for MAP estimation of jump Markov linear systems," *IEEE Trans. Signal Processing*, vol. 47, pp. 2139–2156, Aug. 1999.

[27] S. N. MacEachern, M. Clyde, and J. S. Liu, "Sequential importance sampling for nonparametric Bayes models: The next generation," *Can. J. Statist.*, vol. 27, no. 2, pp. 251–267, 1999.

[28] D. Q. Mayne, "A solution of the smoothing problem for linear dynamic systems," *Automatica*, vol. 4, pp. 73–92, 1966.

[29] E. Mazor, A. Averbuck, Y. Bar-Shalom, and J. Dayan, "Interacting multiple model methods in target tracking: A survey," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 34, pp. 103–123, 1998.

[30] J. M. Mendel, *Maximum-Likelihood Deconvolution: A Journey into Model-Based Signal Processing*. New York: Springer-Verlag, 1990.

[31] J. B. Moore and V. Krishnamurthy, "De-interleaving pulse trains using discrete-time stochastic dynamic-linear models," *IEEE Trans. Signal Processing*, vol. 42, pp. 3092–3103, Nov. 1994.

[32] B. D. Ripley, *Stochastic Simulation*, New York: Wiley, 1987.

[33] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. New York: Springer-Verlag, 1999.

[34] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. London, U.K.: Chapman & Hall, 1986.

[35] V. B. Svetnik, "Applying the Monte Carlo method for optimum estimation in systems with random disturbances," *Automat. Remote Contr.*, vol. 47, no. 6, pp. 818–825, 1986.

[36] D. M. Titterington, A. F. M. Smith, and U. E. Makov, *Statistical Analysis of Finite Mixture Distributions*. New York: Wiley, 1985.

[37] J. K. Tugnait, "Adaptive estimation and identification for discrete systems with Markov jump parameters," *IEEE Trans. Automat. Contr.*, vol. AC–25, pp. 1054–1065, 1982.

[38] ——, "Detection and estimation for abruptly changing systems," *Automatica*, vol. 18, no. 5, pp. 607–615, 1982.

Arnaud Doucet was born in Melle, France, on November 2, 1970. He received the M.S. degree from the Institut National des Télécommunications, Paris, France, in 1993 and the Ph.D. degree from University Paris XI, Orsay, France, in 1997.

During 1998, he was a visiting scholar with the Signal Processing Group, Cambridge University, Cambridge, U.K. From 1999 to 2000, he was a research associate with the same group. Since January 2001, he has been a Lecturer with the Department of Electrical and Electronic Engineering, University of Melbourne, Parkville, Australia. His research interests include Markov chain Monte Carlo methods, sequential Monte Carlo methods, Bayesian statistics, and hidden Markov models. He has co-edited, with J. F. G. de Freitas and N. J. Gordon, *Sequential Monte Carlo Methods in Practice* (New York: Springer-Verlag, 2001).

Neil J. Gordon received the B.Sc. degree in mathematics and physics from Nottingham University, Nottingham, U.K., in 1988 and the Ph.D. degree in statistics from Imperial College, University of London, London, U.K., in 1993.

Since 1988, he has been with the Defence Evaluation and Research Agency (DERA), Malvern, U.K., where he is currently with the Pattern and Information Processing Group. His research interests include Bayesian dynamic estimation, target tracking, and missile guidance.

Vikram Krishnamurthy (SM'99) was born in 1966. He received the Ph.D. degree from the Australian National University, Canberra, in 1992.

He is currently a Professor with the Department of Electrical Engineering, University of Melbourne, Parkville, Australia. His research interests span several areas, including stochastic processing. He is an Associate Editor for *Systems and Control Letters*.

Dr. Krishnamurthy is currently an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING. He has served on the Technical Program Committee of several conferences, including the 37th IEEE Conference on Decision and Control, Tampa, FL, in 1998 and the IFAC Symposium on System Identification, Santa Barbara, CA, in 2000.