Online Sequential Monte Carlo smoother for partially observed stochastic differential equations

Pierre Gloaguen, Marie-Pierre Etienne, Sylvain Le Corff

Online Sequential Monte Carlo smoother for partially observed stochastic differential equations

Pierre Gloaguen* Marie-Pierre Etienne* Sylvain Le Corff†

Abstract

This paper introduces a new algorithm to approximate smoothed additive functionals for partially observed stochastic differential equations. This method relies on the recent procedure introduced in [24] which allows to compute such approximations online, i.e. as the observations are received, and with a computational complexity growing linearly with the number of Monte Carlo samples. The algorithm of [24] cannot be used in the case of partially observed stochastic differential equations since the transition density of the latent data is usually unknown. We prove that a similar algorithm may still be defined for partially observed continuous processes by replacing this unknown quantity by an unbiased estimator obtained for instance using general Poisson estimators. We prove that this estimator is consistent and its performance are illustrated using data from two models.

Keywords: Stochastic differential equations, Smoothing, Sequential Monte Carlo Methods.

1 Introduction

This paper introduces a new algorithm to solve the smoothing problem for partially observed continuous time stochastic processes. In this setting, the hidden state process $(X_t)_{t\geq 0}$ is assumed to be a solution to a stochastic differential equation (SDE) and the only information available is given by noisy observations $(Y_k)_{0\leq k\leq n}$ of the states $(X_k)_{0\leq k\leq n}$ at some discrete time points $(t_k)_{0\leq k\leq n}$. The bivariate stochastic process $\{(X_k, Y_k)\}_{0\leq k\leq n}$ is a state space model such that conditional on the state sequence $(X_k)_{0\leq k\leq n}$ the observations $(Y_k)_{0\leq k\leq n}$ are independent and for all $0 \leq \ell \leq n$ the conditional distribution of Y_ℓ given $\{X_k\}_{0\leq k\leq n}$ depends on X_ℓ only.

[1]AgroParistech, UMR MIA 518, F-75231 Paris, France.
[2]Laboratoire de Mathématiques d'Orsay, Univ. Paris-Sud, CNRS, Université Paris-Saclay.

Statistical inference for partially observed state sequences often requires to solve bayesian filtering and smoothing problems, i.e. the computation of the posterior distributions of sequences of hidden states given observations. The filtering problem refers to the estimation, for each $0 \leq k \leq n$, of the distributions of the hidden state X_k given the observations $(Y_0, \ldots, Y_k)$. Smoothing stands for the estimation of the distributions of the sequence of states $(X_k, \ldots, X_p)$ given observations $(Y_0, \ldots, Y_\ell)$ with $0 \leq k \leq p \leq \ell \leq n$. These posterior distributions are crucial to compute maximum likelihood estimators of unknown parameters using the observations $(Y_0, \ldots, Y_n)$ only. For instance, the E-step of the EM algorithm introduced in [7] boils down to the computation of a conditional expectation of an additive functional of the hidden states given all the observations up to time n. Similarly, by Fisher's identity, recursive maximum likelihood estimates may be computed using the gradient of the loglikelihood which can be written as the conditional expectation of an additive functional of the hidden states. See [5, Chapter 10 and 11], [15, 19, 20, 27] for further references on the use of these smoothed expectations of additive functionals applied to maximum likelihood parameter inference in latent data models.

The exact computation of these expectations is usually not possible in the case of partially observed diffusions. In this paper, we propose to use Sequential Monte Carlo (SMC) methods to approximate smoothing distributions with random particles associated with importance weights. [13, 18] introduced the first particle filters and smoothers for state space models by combining importance sampling steps to propagate particles with resampling steps to duplicate or discard particles according to their importance weights. Unfortunately, these methods cannot be applied directly to partially observed stochastic differential equations since some elementary quantities, such as transition densities of the hidden states, are not available explicitly. Discretization procedures may be used to approximate transition densities, for instance the Euler-Maruyama method, the Ozaki discretization which proposes a linear approximation of the drift coefficient between two observations [25, 28], or Gaussian based approximations using Taylor expansions of the posterior mean and variance of an observation given the observation at the previous time step, [16, 17, 29]. Other approaches based on Hermite polynomials expansion were also introduced by [1, 2, 3] and extended in several directions recently, see [21] and all the references on the approximation of transition densities therein. However, even the most recent discretization based approximations of the transition densities induce a systematic bias of particle based approximations of posterior distributions, see for instance [6]. To overcome this difficulty, [11] proposed to solve the filtering problem by combining SMC methods with an unbiased estimate of the transition densities based on the generalized Poisson estimator (GPE). In this case, only the Monte Carlo error has to be controlled as there is no Taylor expansion to approximate unknown transition densities.

The only solution to solve the smoothing problem for partially observed SDE using SMC methods has been proposed in [23] and extends the fixed-lag smoother of [22]. Using forgetting properties of the hidden chain, the algorithm improves the performance of [11] to approximate smoothing distributions but at the cost of a bias that does not vanish as the number of particles grows to infinity. In the case of discrete time state space models, approximations of the smoothing distributions may also be obtained using the Forward Filtering Backward Smoothing algorithm (FFBS) and the Forward Filtering Backward Simulation algorithm (FFBSi) developed respectively in [18, 14, 9] and [12]. Both algorithms require first a forward pass which produces a set of particles and weights approximating the sequence of filtering distributions up to time n. Then, a backward pass is performed to compute new weights (FFBS) or sample trajectories (FFBSi) in order to approximate the smoothing distributions. Recently, [24] proposed a new SMC algorithm, the particle-based rapid incremental smoother (PaRIS), to approximate on-the-fly (i.e. using the observations as they are received) smoothed expectations of additive functionals. Unlike the FFBS algorithm, the complexity of this algorithm grows only linearly with the number of particles N and contrary to the FFBSi algorithm, no backward pass is required.

In this paper, we extend the use of PaRIS algorithm to partially observed SDE. The proposed algorithm allows to approximate smoothed expectations of additive functionals online and with a complexity growing only linearly with the number of particles. The crucial and simple result (Lemma 1) of the application of PaRIS algorithm to SDE is that the accept reject mechanism introduced in [8] ensuring the linear complexity of the procedure is still correct when the transition densities are replaced by unbiased estimates. The usual FFBS and FFBSi algorithms may not be extended this easily since they both require the computation of weights defined as ratios involving the transition densities, thus replacing these unknown quantities by unbiased estimates does not lead to unbiased estimators of the weights. The proposed Generalized Random version of PaRIS algorithm, hereafter named GRand PaRIS algorithm, may be applied to general hidden Markov models whose Markovian dynamics is ruled by a stochastic differential equation (one of the first two domains defined in [4]) but also to any general state space model where the transition density of the hidden chain may be estimated unbiasedly.

Section 2 describes the proposed algorithm to approximate smoothed additive functionals using unbiased estimates of the transition density of the hidden states and details the application of this algorithm when the transition density may be approximated using a GPE. In Section 3, classical convergence results for SMC smoothers are extended to the setting of this paper and illustrated with numerical experiments in Section 4. All proofs are postponed to Appendix A.

2 The Generalized Random PaRIS algorithm

$(X_t)_{t\geq 0}$ is defined as a weak solution to the following SDE in $\mathbb{R}^d$:

$$X_0 = x_0 \quad \text{and} \quad \mathrm{d}X_t = \alpha(X_t)\mathrm{d}t + \mathrm{d}W_t \ , \tag{1}$$

where $(W_t)_{t\geq 0}$ is a standard Brownian motion. It is assumed that α is of the form $\alpha(x) = \nabla_x A(x)$ where $A : \mathbb{R}^d \to \mathbb{R}$ is a twice continuously differentiable function. The solution to (1) is supposed to be partially observed at times $t_0 = 0, \ldots, t_n$ through an observation process $(Y_k)_{0\leq k\leq n}$ in $(\mathbb{R}^m)^{n+1}$. For all $0 \leq k \leq n$, the distribution of Y_k given $X_k := X_{t_k}$ has a density with respect to a reference measure λ on $\mathbb{R}^m$ given by $g(X_k, \cdot) = g_k(X_k)$. The distribution of X_0 has a density with respect to a reference measure μ on $\mathbb{R}^d$ given by χ. For all $0 \leq k \leq n-1$, the conditional distribution of X_{k+1} given X_k has a density $q_k(X_k, \cdot)$ with respect to μ.

Let $0 \leq k \leq k' \leq n$, the joint smoothing distributions of the hidden states are defined, for all measurable function h on $(\mathbb{R}^d)^{k'-k+1}$, by:

$$\phi_{k:k'|n}[h] = \mathbb{E}\left[h(X_k, \ldots, X_{k'})|Y_{0:n}\right] \ .$$

For all $0 \leq k \leq n$, $\phi_k = \phi_{k:k|k}$ denote the filtering distributions. The aim of this section is to detail the extension of PaRIS algorithm to approximate expectations of the form

$$\phi_{0:n|n}[H_n] = \mathbb{E}\left[H_n(X_{0:n})|Y_{0:n}\right] \text{ where } H_n = \sum_{k=0}^{n-1} h_k(X_k, X_{k+1}) \ , \tag{2}$$

when the transition density of the hidden states is not available explicitly and where $\{h_k\}_{k=0}^{n-1}$ are given functions on $\mathbb{R}^d \times \mathbb{R}^d$. The algorithm is based on the following link between the filtering and smoothing distributions for additive functionals, see [24]:

$$\phi_{0:n|n}[h] = \phi_n[T_n[h]] \ , \text{ where } T_n[h](X_n) = \mathbb{E}\left[h(X_{0:n})|X_n, Y_{0:n}\right] \ . \tag{3}$$

The approximation of (3) requires first to approximate the sequence of filtering distributions. Sequential Monte Carlo methods provide an efficient and simple solution to obtain these approximations using sets of particles $\{\xi_k^\ell\}_{\ell=1}^N$ associated with weights $\{\omega_k^\ell\}_{\ell=1}^N$, $0 \leq k \leq n$.

At time $k = 0$, N particles $\{\xi_0^\ell\}_{\ell=1}^N$ are sampled independently according to $\xi_0^\ell \sim \eta_0$, where η_0 is a probability density with respect to μ. Then, ξ_0^ℓ is associated with the importance weights $\omega_0^\ell = \chi(\xi_0^\ell) g_0(\xi_0^\ell)/\eta_0(\xi_0^\ell)$. For any bounded and measurable function h defined on $\mathbb{R}^d$, the expectation $\phi_0[h]$ is approximated by

$$\phi_0^N[h] = \frac{1}{\Omega_0^N} \sum_{\ell=1}^N \omega_0^\ell h\left(\xi_0^\ell\right) \ , \quad \Omega_0^N := \sum_{\ell=1}^N \omega_0^\ell \ .$$

Then, for $1 \leq k \leq n$, using $\{(\xi_{k-1}^\ell, \omega_{k-1}^\ell)\}_{\ell=1}^N$, the auxiliary particle filter of [26] samples pairs $\{(I_k^\ell, \xi_k^\ell)\}_{\ell=1}^N$ of indices and particles using an instrumental transition density p_k on $\mathbb{R}^d \times \mathbb{R}^d$ and an adjustment multiplier function ϑ_k on $\mathbb{R}^d$. Each new particle ξ_k^ℓ and weight ω_k^ℓ at time k are computing following these steps:

- choose a particle index I_k^ℓ at time $k-1$ in $\{1, \ldots, N\}$ with probabilities proportional to $\omega_{k-1}^j \vartheta_k(\xi_{k-1}^j)$, for j in $\{1, \ldots, N\}$;
- sample ξ_k^ℓ using this chosen particle according to $\xi_k^\ell \sim p_k(\xi_{k-1}^{I_k^\ell}, \cdot)$;
- associate the particle ξ_k^ℓ with the importance weight:

$$\omega_k^\ell := \frac{q_k(\xi_{k-1}^{I_k^\ell}, \xi_k^\ell) g_k(\xi_k^\ell)}{\vartheta_k(\xi_{k-1}^{I_k^\ell}) p_k(\xi_{k-1}^{I_k^\ell}, \xi_k^\ell)} \ . \tag{4}$$

The expectation $\phi_k[h]$ is approximated by

$$\phi_k^N[h] := \frac{1}{\Omega_k^N} \sum_{\ell=1}^N \omega_k^\ell h\left(\xi_k^\ell\right) \ , \quad \Omega_k^N := \sum_{\ell=1}^N \omega_k^\ell \ .$$

PaRIS algorithm uses the same decomposition as the FFBS algorithm introduced in [10] and the FFBSi algorithm proposed by [12] to approximate smoothing distributions. It combines both the forward only version of the FFBS algorithm with the sampling mechanism of the FFBSi algorithm. It does not produce an approximation of the smoothing distributions but of the smoothed expectation of a fixed additive functional and thus may be used to approximate (2). Its crucial property is that it does not require a backward pass, the smoothed expectation is computed on-the-fly with the particle filter and no storage of the particles or weights is needed.

PaRIS algorithm relies on the following fundamental property of $T_k[H_k]$ when H_k is as in (2):

$$\begin{aligned} T_k[H_k](X_k) &= \mathbb{E}\left[T_{k-1}[H_{k-1}](X_{k-1}) + h_{k-1}(X_{k-1}, X_k) | X_k, Y_{0:k-1}\right] \ , \\ &= \frac{\int \phi_{k-1}(\mathrm{d}x_{k-1}) q_{k-1}(x_{k-1}, X_k) \left\{T_{k-1}[H_{k-1}](x_{k-1}) + h_{k-1}(x_{k-1}, X_k)\right\}}{\int \phi_{k-1}(\mathrm{d}x_{k-1}) q_{k-1}(x_{k-1}, X_k)} \ . \end{aligned}$$

Therefore, [24] introduces sufficient statistics τ_k^i (starting with $\tau_0^i = 0$, $1 \leq i \leq N$), approximating $T_k[H_k](\xi_k^i)$, for $1 \leq i \leq N$ and $0 \leq k \leq n$. First, replacing ϕ_{k-1} by ϕ_{k-1}^N in the last equation leads to the following approximation of $T_k[H_k](\xi_k^i)$:

$$T_k^N[H_k](\xi_k^i) = \sum_{j=1}^N \Lambda_{k-1}^N(i,j) \left\{T_{k-1}[H_{k-1}](\xi_{k-1}^j) + h_{k-1}(\xi_{k-1}^j, \xi_k^i)\right\} \ , \tag{5}$$

where

$$\Lambda_k^N(i,\ell) = \frac{\omega_k^\ell q_k(\xi_k^\ell, \xi_{k+1}^i)}{\sum_{\ell=1}^N \omega_k^\ell q_k(\xi_k^\ell, \xi_{k+1}^i)} , \quad 1 \le \ell \le N . \tag{6}$$

Computing exactly these approximations would lead to a complexity growing quadratically with N because of the normalizing constant in (6). Therefore, PaRIS algorithm samples particles in the set $\{\xi_{k-1}^j\}_{j=1}^N$ with probabilities $\Lambda_k^N(i,\cdot)$ to approximate the expectation (5) and produce τ_k^i. Choosing $\tilde{N} \ge 1$, at each time step $0 \le k \le n-1$ these statistics are updated according to the following steps.

(i) Run one step of a particle filter to produce $\{(\xi_k^\ell, \omega_k^\ell)\}$ for $1 \le \ell \le N$.

(ii) For all $1 \le i \le N$, sample independently $J_k^{i,\ell}$ in $\{1,\dots,N\}$ for $1 \le \ell \le \widetilde{N}$ with probabilities $\Lambda_k^N(i,\cdot)$, given by (6).

(iii) Set

$$\tau_{k+1}^i := \frac{1}{\widetilde{N}} \sum_{\ell=1}^{\widetilde{N}} \left\{ \tau_k^{J_k^{i,\ell}} + h_k\left(\xi_k^{J_k^{i,\ell}}, \xi_{k+1}^i\right)\right\} .$$

Then, (2) is approximated by

$$\phi_{0:n|n}^N[\tau_n] = \frac{1}{\Omega_n^N} \sum_{i=1}^N \omega_n^i \tau_n^i .$$

As proved in [24], the algorithm is asymptotically consistent (as N goes to infinity) for any precision parameter $\tilde{N}$. However, there is a significant qualitative difference between the cases $\tilde{N} = 1$ and $\tilde{N} \ge 2$. As for the FFBSi algorithm, when there exists σ_+ such that $0 < q_k < \sigma_+$, PaRIS algorithm may be implemented with $\mathcal{O}(N)$ complexity using the accept-reject mechanism of [8].

In general situations, PaRIS algorithm cannot be used for stochastic differential equations as q_k is unknown. Therefore, the computation of the importance weights ω_k^ℓ and of the acceptance ratio of [8] is not tractable. Following [11, 23], filtering weights can be approximated by replacing $q_k(\xi_k^\ell, \xi_{k+1}^i)$ by an unbiased estimator $\widehat{q}_k(\xi_k^\ell, \xi_{k+1}^i; \zeta_k)$, where ζ_k is a random variable in $\mathbb{R}^q$ such that:

$$\widehat{q}_k(\xi_k^\ell, \xi_{k+1}^i; \zeta_k) > 0 \ \text{ a.s} \quad \text{and} \quad \mathbb{E}\left[\widehat{q}_k(\xi_k^\ell, \xi_{k+1}^i; \zeta_k) \middle| \mathcal{G}_{k+1}^N\right] = q_k(\xi_k^\ell, \xi_{k+1}^i) ,$$

where, for all $0 \le k \le n$,

$$\mathcal{F}_k^N = \sigma\left\{ Y_{0:k}; (\xi_u^\ell, \omega_u^\ell, \tau_u^\ell); J_v^{\ell,j};\ 1 \le \ell \le N,\ 0 \le u \le k, 1 \le j \le \widetilde{N}, 0 \le v < k \right\} ,$$
$$\mathcal{G}_{k+1}^N = \mathcal{F}_k^N \vee \sigma\left\{ Y_{k+1}; (\xi_{k+1}^\ell, \omega_{k+1}^\ell);\ 1 \le \ell \le N \right\} .$$

Practical choices for ζ_k are discussed below, see for instance (9) which presents the choice made for the implementation of such estimators in our context. In the case where q_k is unknown, the filtering weights in (4) then become:

$$\widehat{\omega}_k^\ell := \frac{\widehat{q}_k(\xi_{k-1}^{I_k^\ell}, \xi_k^\ell; \zeta_k) g_k(\xi_k^\ell)}{\vartheta_k(\xi_{k-1}^{I_k^\ell}) p_k(\xi_{k-1}^{I_k^\ell}, \xi_k^\ell)} \ . \tag{7}$$

Therefore, to obtain a generalized random version of PaRIS algorithm, we only need to be able to sample from the discrete probability distribution $\Lambda_k^N(i, \cdot)$ in the case of SDE based HMM. Consider the following assumption: for all $0 \leq k \leq n$, there exists a random variable $\hat{\sigma}_+^k$ measurable with respect to $\mathcal{G}_{k+1}^N$ such that,

$$\sup_{x,y,\zeta} \widehat{q}_k(x, y; \zeta) \leq \hat{\sigma}_+^k \ . \tag{A1}$$

Lemma 1. *Assume that A_1 holds for some $0 \leq k \leq n-1$. For all $1 \leq i \leq N$, define the random variable J_k^i as follows:*

repeat
Sample independently ζ, $U \sim \mathcal{U}[0,1]$ and $J \in \{1, \ldots, N\}$ with probabilities proportional to $\{\widehat{\omega}_k^1, \ldots, \widehat{\omega}_k^N\}$.
until $U \leq \widehat{q}_k(\xi_k^J, \xi_{k+1}^i, \zeta)/\hat{\sigma}_+^k$.
Set $J_k^i = J$.

Then, the conditional probability distribution given $\mathcal{G}_{k+1}^N$ of J_k^i is $\Lambda_k^N(i, \cdot)$.

Proof. See Appendix A. □

Note that Lemma 1 still holds if assumption (**A**1) is relaxed and replaced by one of the two following assumptions:

$$\sup_{j,y,\zeta} \widehat{q}_k(\xi_k^j, y, \zeta) \leq \hat{\sigma}_+^k \ . \tag{A2}$$

$$\sup_{i,j,\zeta} \widehat{q}_k(\xi_k^j, \xi_{k+1}^j, \zeta) \leq \hat{\sigma}_+^k \ . \tag{A3}$$

It is worth noting that under assumptions (**A**1) or (**A**2), the linear complexity property of PaRIS algorithm still holds, whereas if only assumption (**A**3) holds, the algorithm has a quadratic complexity.

Bounded estimator of q_k

For $x, y \in \mathbb{R}^d$, by Girsanov and Ito's formulas, the transition density $q_k(x, y)$ of (1) satisfies, with $\Delta_k = t_{k+1} - t_k$,

$$q_k(x, y) = \varphi_{\Delta_k}(x, y) \exp\{A(y) - A(x)\} \, \mathbb{E}_{\mathbb{W}^{x,y,\Delta_k}} \left[\exp\left\{ - \int_0^{\Delta_k} \phi(\mathsf{w}_s) \mathrm{d}s \right\} \right] \ ,$$

where $\mathbb{W}^{x,y,\Delta_k}$ is the law of Brownian bridge starting at x at 0 and hitting y at Δ_k, $(\mathsf{w}_t)_{0\leq t\leq\Delta_k}$ is such a Brownian bridge, $\varphi_{\Delta_k}(x,y)$ is the p.d.f. of a normal distribution with mean x and variance Δ_k, evaluated at y and $\phi : \mathbb{R}^d \to \mathbb{R}$ is defined as:

$$\phi(x) = \left(\|\alpha(x)\|^2 + \triangle A(x)\right)/2 \,,$$

with $\triangle$ the Laplace operator. Assume that there exist random variables $\mathsf{L_w}$ and $\mathsf{U_w}$ such that for all $0 \leq s \leq \Delta_k$, $\mathsf{L_w} \leq \phi(\mathsf{w}_s) \leq \mathsf{U_w}$. Let κ be a random variable taking values in $\mathbb{N}$ with distribution μ and $(U_j)_{1\leq j\leq\kappa}$ be independent uniform random variables on $[0,\Delta_k]$, and $\zeta_k = \{\kappa, \mathsf{w}, U_1, \ldots, U_\kappa\}$. As shown in [11], a positive unbiased estimator is given by

$$\widehat{q}_k(x,y;\zeta_k) = \varphi_{\Delta_k}(x,y)\exp\{A(y)-A(x)\} \times \exp\{-\mathsf{U_w}\Delta\}\frac{\Delta_k^\kappa}{\mu(\kappa)\kappa!}\prod_{j=1}^{\kappa}\left(\mathsf{U_w}-\phi(\mathsf{w}_{U_j})\right) \,. \quad (8)$$

Interesting choices of μ are discussed in [11] and we focus here on the so called GPE-1, where μ is a Poisson distribution with intensity $(\mathsf{U_w}-\mathsf{L_w})\Delta_k$. In that case, the estimator (8) becomes:

$$\widehat{q}_k(x,y;\zeta_k) = \varphi_{\Delta_k}(x,y)\exp\{A(y)-A(x)-\mathsf{L_w}\Delta_k\}\prod_{j=1}^{\kappa}\frac{\mathsf{U_w}-\phi(\mathsf{w}_{U_j})}{\mathsf{U_w}-\mathsf{L_w}} \,. \quad (9)$$

On the r.h.s. of (9), the product over κ elements is bounded by 1, therefore, a sufficient condition to satisfy of the assumptions (**A**1)-(**A**3) is that the function:

$$\begin{aligned}\rho_{\Delta_k} :\ & \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R} \\ & (x,y) \mapsto \varphi_{\Delta_k}(x,y)\exp\{A(y)-A(x)-\mathsf{L_w}\Delta_k\}\end{aligned} \quad (10)$$

is upper bounded almost surely by $\hat{\sigma}_+^k$. In particular, if $\mathsf{L_w}$ is bounded almost surely, (10) always satisfies assumption (**A**3) and Algorithm 1 can be used. This condition is always satisfied for models in the domains $\mathcal{D}_1$ and $\mathcal{D}_2$ defined in [4], i.e. domains for which the exact algorithms EA1 and EA2 can be used.

When (**A**1) or (**A**2) holds, it can be nonetheless of practical interest to choose the bound $\hat{\sigma}_+^k$ corresponding to (**A**3). Indeed, this might increase significantly the acceptance rate of the algorithm, and therefore reduce the number of drawings of the random variable ζ, which has a much higher cost than the computation of ρ, as it requires simulations of Brownian Bridges. Moreover, this latter option can also avoid numerical optimization if no analytical expression of $\hat{\sigma}_+^k$ is available. In practice, we found this option more efficient in terms of computation time when N has moderate values.

Algorithm 1 GRand PaRIS algorithm

for all $i \in 1, \dots, N$ **do**
  Sample $\xi_0^i \sim \eta_0$, $\tau_0^i = 0$ and $\widehat{\omega}_0^i = g_0(\xi_0^i)\chi_0(\xi_0^i)/\eta_0(\xi_0^i)$.
end for
for $k \in 0, \dots, n-1$ **do**
  for all $i \in 1, \dots, N$ **do**
    Set $\tau_{k+1}^i = 0$;
    Sample I_{k+1}^i in $\{1, \dots, N\}$ with probabilities proportional to $\{\widehat{\omega}_k^1 \vartheta_{k+1}(\xi_k^1), \dots, \widehat{\omega}_k^N \vartheta_{k+1}(\xi_k^N)\}$.
    Sample $\xi_{k+1}^i \sim p_k(\xi_k^{I_{k+1}^i}, \cdot)$.
    For all $1 \leq m \leq M$, sample independently $\zeta_k^m = (\kappa_m, \mathsf{w}_m, (U_j^m)_{1 \leq j \leq \kappa_m})$ with $\kappa_m \sim \mu$, $\mathsf{w}_m \sim \mathbb{W}_k^{X_{k+1}}$ and $(U_j^m)_{1 \leq j \leq \kappa_m} \sim \mathcal{U}[0, \Delta_k]^{\otimes \kappa_m}$.
    Compute $\widehat{\omega}_{k+1}^i$ using equation (7).
    for all $\ell \in 1, \dots, \widetilde{N}$ **do**
      Sample $J_k^{i,\ell}$ as in Lemma 1.
      Update $\tau_{k+1}^i = \tau_{k+1}^i + (\tau_k^{J_k^{i,\ell}} + h_k(\xi_k^{J_k^{i,\ell}}, \xi_{k+1}^i))/\tilde{N}$.
    end for
  end for
end for

3 Convergence results

Consider the following assumptions.

H1 (i) For all $k \geq 0$ and all $x \in \mathbb{R}^d$, $g_k(x) > 0$.
(ii) $\sup_{k \geq 0} |g_k|_\infty < \infty$.

H2 $\sup_{k \geq 1} |\vartheta_k|_\infty < \infty$, $\sup_{k \geq 1} |p_k|_\infty < \infty$ and $\sup_{k \geq 1} |\widehat{\omega}_k|_\infty < \infty$, where

$$\widehat{\omega}_0(x) = \frac{\chi(x) g_0(x)}{\eta_0(x)} \quad \text{and for } k \geq 1 \quad \widehat{\omega}_k(x, x'; z) = \frac{\widehat{q}_k(x, x'; z) g_{k+1}(x')}{\vartheta_{k+1}(x) p_k(x, x')} \, .$$

Lemma 2. *For all $0 \leq k \leq n-1$, the random variables $\{\widehat{\omega}_{k+1}^i \tau_{k+1}^i\}_{i=1}^N$ are independent conditionally on $\mathcal{F}_k^N$ and*

$$\mathbb{E}\left[\widehat{\omega}_{k+1}^1 \tau_{k+1}^1 \middle| \mathcal{F}_k^N\right] = \left(\phi_k^N[\vartheta_{k+1}]\right)^{-1} \phi_k^N \left[\int q_k(\cdot, x) g_{k+1}(x) \left\{\tau_k(\cdot) + h_{k+1}(\cdot, x)\right\} \mathrm{d}x\right] \, .$$

Proof. See appendix A □

Proposition 1. *Assume that H1 and H2 hold and that for all* $1 \leq k \leq n$, $\mathrm{osc}(h_k) < +\infty$. *For all* $0 \leq k \leq n$ *and all* $\widetilde{N} \geq 1$, *there exist* $b_k, c_k > 0$ *such that for all* $N \geq 1$ *and all* $\varepsilon \in \mathbb{R}_+^\star$,

$$\mathbb{P}\left(\left|\phi_k^N[\tau_k] - \phi_k\left[T_k h_k\right]\right| \geq \varepsilon\right) \leq b_k \exp\left(-c_k N \varepsilon^2\right) \ .$$

Proof. See appendix A □

4 Numerical experiments

This section investigates the performance of the proposed algorithm with the sine and log-growth models. In both cases, the proposal distribution p_k is chosen as the following approximation of the optimal filter (or the fully adapted particle filter in the terminology of [26]):

$$p_k(x_{k-1}, x_k) \propto \tilde{q}_k(x_{k-1}, x_k) g_k(x_k) \ ,$$

where $\tilde{q}_k(x_{k-1}, x_k)$ is the p.d.f. of Gaussian distibution with mean $\alpha(x_{k-1})\Delta_k$ and variance $\Delta_k I_d$, i.e. the Euler approximation of equation (1). As the observation model is linear and Gaussian, the proposal distribution is therefore Gaussian with explicit mean and variance.

In order to evaluate the performance of the proposed algorithm, the following strategy has been chosen. We compare the estimation of the EM intermediate quantity with the one obtained by the fixed lag method of [23], for different values of the lag (namely, 1,2,5,10,50). The particle approximation of $\mathcal{Q}(\theta, \theta)$ for each model is computed using each algorithm, see Figure 1 for the SINE model (and respectively Figure 3 for the log-growth model). This estimation is performed 200 times to obtain the estimates $\widehat{Q}_1, \ldots, \widehat{Q}_{200}$, using $\tilde{N} = 2$ particles for PaRIS algorithm, and $M = 30$ replications for the Monte Carlo approximation $\widehat{q}_k$ of each q_k. Moreover, the E step requires the computation of a quantity such as (2) with $h_k = \log g_k + \log q_k$. $\log q_k$ is not available explicitly and is approximated using the unbiased estimator proposed in [23, Appendix B] based on 30 independent Monte Carlo simulations. The intermediate quantity of the EM algorithm is also estimated with our algorithm 30 times using $N = 5000$ particles, the reference value is then computed as the arithmetic mean of these 30 estimations, and denoted by $\widehat{Q}_\star$. Figure 1 (resp. 3) shows this estimate for an example on one simulated data set. The GRand Paris algorithm is performed using $N = 400$ particles in both cases, the fixed lag technique using $N = 1600$ so that both estimations require similar computational times.

The SINE model

The performance of the GRand PaRIS algorithm are first highlighted using the SINE model, where $(X_t)_{t\geq 0}$ is supposed to be the solution to:

$$\mathrm{d}X_t = \sin\left(X_t - \theta\right)\mathrm{d}t + \mathrm{d}W_t, \quad X_0 = x_0 \ . \tag{11}$$

This simple model has no explicit transition density, however GPE estimators may be computed by simulating Brownian bridges. The process solution to (11) is observed regularly at times $t_0 = 0, \ldots, t_{100} = 50$ through the observation process $(Y_k)_{0\leq k\leq 100}$:

$$Y_k = X_k + \varepsilon_k \ , \tag{12}$$

where the $(\varepsilon)_{0\leq k\leq 100}$ are i.i.d. $\mathcal{N}(0,1)$. In the example displayed on Figure 1, we set $\theta = 0$. In that case, the function ρ_{Δ_k} defined in (10) can be upper bounded either on (x, y) or only on y, the GRand PaRIS algorithm has therefore a linear complexity.

This same experiment was reproduced on 100 different simulated data sets. For each simulation s, the empirical absolute relative bias arb_s and the empirical absolute coefficient of variation acv_s are computed as

$$\mathsf{arb}_s = \frac{|m(\widehat{Q}^s) - \widehat{Q}^s_\star|}{|\widehat{Q}^s_\star|} \tag{13}$$

$$\mathsf{acv}_s = \frac{\sigma(\widehat{Q}^s)}{|m(\widehat{Q}^s)|} \tag{14}$$

where $m(\widehat{Q}^s)$ and $\sigma(\widehat{Q}^s)$ are the empirical mean and standard deviation of the sample $Q^s_1, \ldots, Q^s_{200}$. For each estimation method, the resulting distributions of $\mathsf{arb}_1, \ldots, \mathsf{arb}_{100}$ and $\mathsf{acv}_1, \ldots, \mathsf{acv}_{100}$ are shown on Figure 2.

The GRand PaRIS algorithm outperforms the fixed lag methods for any value of the lag as the bias is the lowest (it is already negligible for $N = 400$) and with a lower variance than fixed lag estimates with negligible bias (i.e., in this case, lags larger than 10). Small lags lead to strongly biased estimates for the fixed lag method, and unbiased estimates are at the cost of a large variance. It is worth noting here that the lag for which the bias is small is model dependent.

Log-growth model

Following [4] and [24], the performance of the proposed algorithm are also illustrated with the log-growth model defined by:

$$\mathrm{d}Z_t = \kappa Z_t \left(1 - \frac{Z_t}{\gamma}\right)\mathrm{d}t + \sigma Z_t \mathrm{d}W_t, \quad Z_0 = z_0. \tag{15}$$

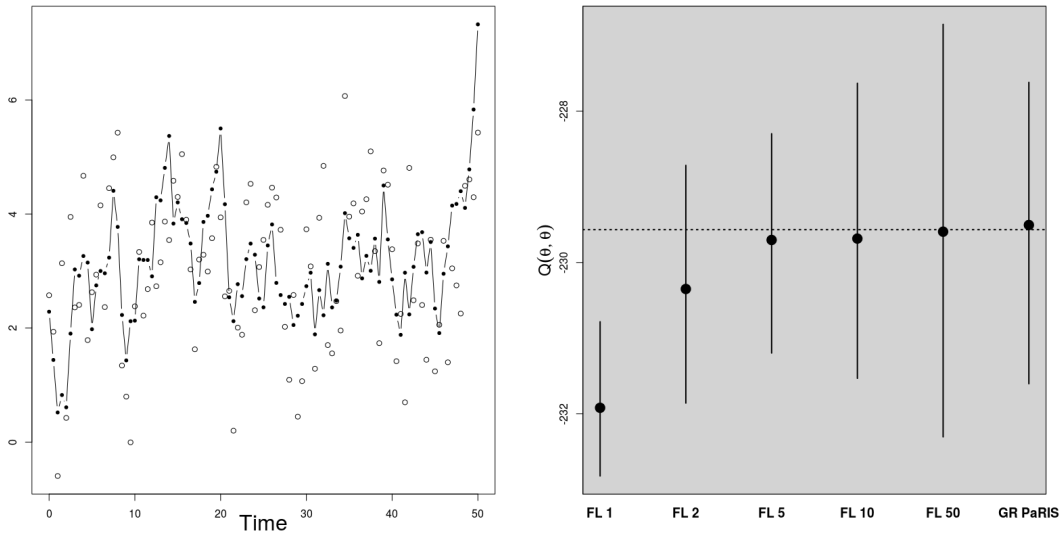

Figure 1: *SINE model*. Process X solution to the SDE (balls) and observations Y (circles) at times $t_0 = 0, \ldots, t_{100} = 50$ [left]. Estimation of the EM intermediate quantity $\mathcal{Q}(\theta, \theta)$ using the fixed lag (FL) technique for 5 different lags, and the GRand PaRIS algorithm using 200 replicates [right]. The whiskers represent the extent of the 95% central values. The dot represents the empirical mean over the 200 replicates. The dotted line shows the reference value, computed using the GRand PaRIS algorithm with $N = 5000$ particles.

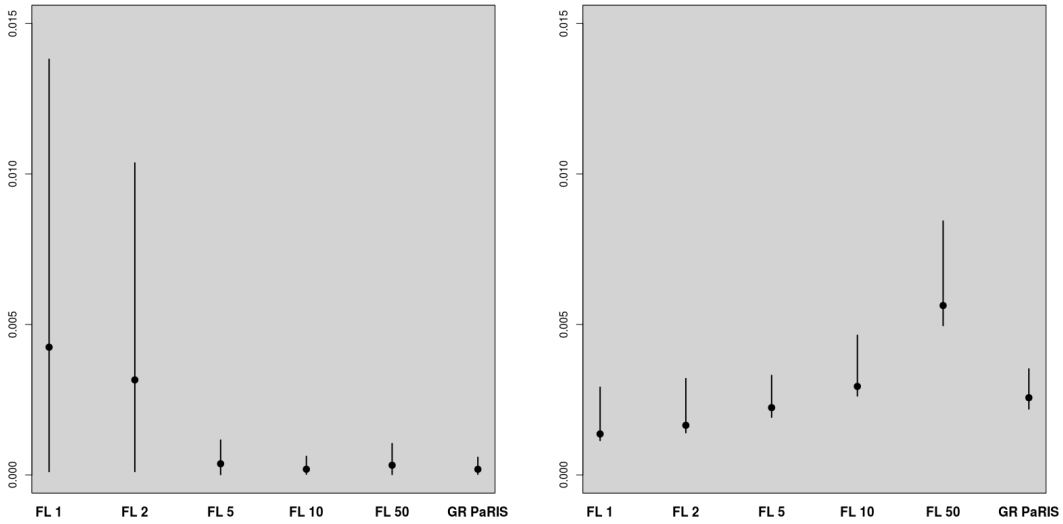

Figure 2: *SINE model*. Distribution of the empirical absolute relative bias [left] and of the empirical absolute coefficient of variation [right] for each method.

In order to use the exact algorithms of [4] and the GPE of [11], we consider (15) after the Lamperti transform, i.e., the process defined by $X_t = \eta(Z_t)$, with $\eta(z) := -\log(z)/\sigma$, which satisfies the following SDE:

$$\mathrm{d}X_t = \overbrace{\left(\frac{\sigma}{2} - \frac{\kappa}{\sigma} + \frac{\kappa}{\gamma\sigma}\exp\left(-\sigma X_t\right)\right)}^{:=\alpha(X_t)}\mathrm{d}t + \mathrm{d}W_t, \quad X_0 = x_0 = \eta(z_0). \tag{16}$$

In this case, the conditions of the Exact Algorithm 2 defined in [4] are satisfied, as for any $m \in \mathbb{R}$ there exists U_m such that for all $x \geq m$, $\psi(x) := \alpha^2(x) + \alpha'(x) \leq \mathsf{U}_m$. Moreover, ψ is lower bounded uniformly by L. Then, GPE estimators may be computed by simulating the minimum of a Brownian bridge, and simulating Bessel bridges conditionally to this minimum, as proposed by [4].

The process solution to (16) is observed regularly at times $t_0 = 0, \ldots, t_{50} = 100$ through the observation process $(Y_k)_{0 \leq k \leq 50}$ defined as:

$$Y_k = X_k + \varepsilon_k \ , \tag{17}$$

where the $(\varepsilon_k)_{0 \leq k \leq 50}$ are i.i.d. $\mathcal{N}(0, \sigma^2_{obs})$. The parameters are given by

$$\theta = (\kappa = 0.1, \sigma = 0.1, \gamma = 1000, \sigma^2_{obs} = 4) \ .$$

In that case, the ρ_{Δ_k} function defined in (10) can be upper bounded as a function of y when $x \in \{\xi_k^1, \ldots, \xi_k^N\}$, the GRand PaRIS algorithm has therefore a linear complexity. The intermediate quantity of the EM algorithm is evaluated as for the SINE model, see Figures 3 and 4.

The results for the fixed lag technique are similar to the ones presented in [23, Figure 1] on the same model. For small lags, the variance of the estimates is small, but the estimation is highly biased. The bias rapidly decreases as the lag increases, together with a great increase of variance. Again, the GRand PaRIS algorithm outperforms the fixed lag smoother as it shows a similar (vanishing) bias as the fixed lag for the largest lag and a smaller variance than the fixed lags estimates with negligible bias.

5 Conclusions

This paper presents a new online SMC smoother for partially observed differential equations. This algorithm relies on an accept-reject procedure inspired from the recent PaRIS algorithm. The main result of the article for practical applications is that the mechanism of this procedure remains valid when the transition density is approximated by a an unbiased positive estimator. The proposed procedure outperforms the existing fixed lag smoother for SDE of [23], as it does not introduce

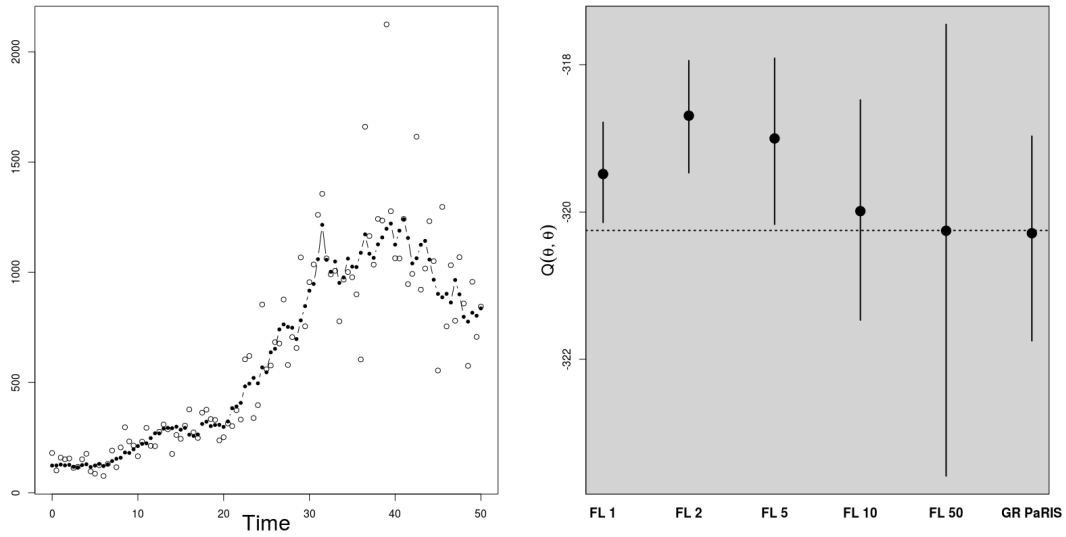

Figure 3: *Log-growth model.* Process X solution to the SDE (balls) and observations Y (circles) at times $t_0 = 0, \ldots, t_{100} = 50$ [left]. Estimation of the EM intermediate quantity $\mathcal{Q}(\theta, \theta)$ using the fixed lag (FL) technique for 5 different lags, and the GRand PaRIS algorithm using 200 replicates [right]. The whiskers represent the extent of the 95% central values. The dot represents the empirical mean over the 200 replicates. The dotted line shows the reference value, computed using the GRand PaRIS algorithm with $N = 5000$ particles.

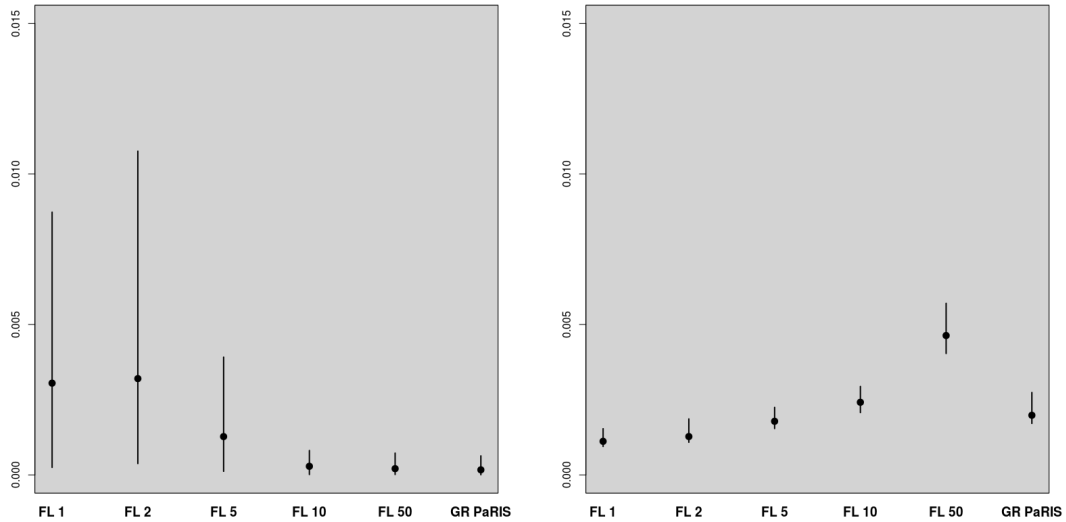

Figure 4: *Log-growth model.* Distribution of the empirical absolute relative bias [left] and of the empirical absolute coefficient of variation [right] for each method.

an intrinsic and non vanishing bias. In addition, numerical simulations highlight a better variance using data from two different models. It can be implemented for the class of models $\mathcal{D}_1$ and $\mathcal{D}_2$ defined in [4] with a linear complexity in N.

A Proofs

Proof of Lemma 1. Let τ be the first time draws are accepted in the accept-reject mechanism. For all $\ell \geq 1$, write

$$\mathcal{A}_\ell^k = \left\{ U_\ell < \widehat{q}_k(\xi_k^{J_\ell}, \xi_{k+1}^i, \zeta_k^\ell)/\hat{\sigma}_+^k \right\} .$$

Let h be a function defined on $\{1, \ldots, N\}$,

$$\begin{aligned}
\mathbb{E}\left[h(J_k^{i,j})\middle|\mathcal{G}_{k+1}^N\right] &= \sum_{m\geq 1} \mathbb{E}\left[h(J_m)1_{\tau=m}\middle|\mathcal{G}_{k+1}^N\right] , \\
&= \sum_{m\geq 1} h(m) \left(\prod_{\ell=1}^{m-1} \mathbb{E}\left[1_{(\mathcal{A}_\ell^k)^c}\middle|\mathcal{G}_{k+1}^N\right] \right) \mathbb{E}\left[h(J_m)1_{\mathcal{A}_m^k}\middle|\mathcal{G}_{k+1}^N\right] , \\
&= \sum_{m\geq 1} \left(\prod_{\ell=1}^{m-1} \mathbb{E}\left[1 - \frac{\widehat{q}_k(\xi_k^{J_\ell}, \xi_{k+1}^i; \zeta_k^\ell)}{\hat{\sigma}_+^k}\middle|\mathcal{G}_{k+1}^N\right] \right) \\
&\qquad \times \mathbb{E}\left[h(J_m)\frac{\widehat{q}_k(\xi_k^{J_m}, \xi_{k+1}^i; \zeta_k^m)}{\hat{\sigma}_+^k}\middle|\mathcal{G}_{k+1}^N\right] , \\
&= \sum_{m\geq 1} \left(\mathbb{E}\left[1 - \frac{q_k(\xi_k^{J_1}, \xi_{k+1}^1)}{\hat{\sigma}_+^k}\middle|\mathcal{G}_{k+1}^N\right] \right)^{m-1} \mathbb{E}\left[h(J_1)\frac{q_k(\xi_k^{J_1}, \xi_{k+1}^1)}{\hat{\sigma}_+^k}\middle|\mathcal{G}_{k+1}^N\right] , \\
&= \mathbb{E}\left[h(J_1)q_k(\xi_k^{J_1}, \xi_{k+1}^i)\middle|\mathcal{G}_{k+1}^N\right] / \mathbb{E}\left[q_k(\xi_k^{J_1}, \xi_{k+1}^i)\middle|\mathcal{G}_{k+1}^N\right] , \\
&= \sum_{\ell=1}^N \frac{h(\ell)\omega_{k-1}^\ell q_k(\xi_k^\ell, \xi_{k+1}^i)}{\sum_{m=1}^N \omega_{k-1}^m q_k(\xi_k^m, \xi_{k+1}^i)} , \\
&= \sum_{\ell=1}^N \Lambda_{k-1}^N(i,\ell)h(\ell) ,
\end{aligned}$$

which concludes the proof. □

Proof of Lemma 2. The independence is ensured by the mechanism of SMC methods. By (7),

$$\mathbb{E}\left[\widehat{\omega}_{k+1}^i \tau_{k+1}^i \middle| \mathcal{F}_k^N\right] = \mathbb{E}\left[\frac{\widehat{q}_k(\xi_k^{I_{k+1}^i}, \xi_{k+1}^i; \zeta_k) g_{k+1}(\xi_{k+1}^i)}{\vartheta_{k+1}(\xi_k^{I_{k+1}^i}) p_{k+1}(\xi_k^{I_{k+1}^i}, \xi_{k+1}^i)} \tau_{k+1}^i \middle| \mathcal{F}_k^N \right] .$$

Note that by Lemma 1,

$$\mathbb{E}\left[\tau_{k+1}^i \middle| \mathcal{G}_{k+1}^N\right] = \sum_{\ell=1}^N \frac{\omega_k^\ell q_k(\xi_k^\ell, \xi_{k+1}^i)\left(\tau_k^\ell + h_k(\xi_k^\ell, \xi_{k+1}^i)\right)}{\sum_{\ell'=1}^N \omega_k^{\ell'} q_k(\xi_k^{\ell'}, \xi_{k+1}^i)} ,$$
$$\mathbb{E}\left[\widehat{q}_k(\xi_k^{I_{k+1}^i}, \xi_{k+1}^i; \zeta_k) \middle| \mathcal{G}_{k+1}^N\right] = q_k(\xi_k^{I_{k+1}^i}, \xi_{k+1}^i) .$$

Since τ_{k+1}^i and ζ_k are independent conditionally to $\mathcal{G}_{k+1}^N$:

$$\mathbb{E}\left[\tau_{k+1}^i \widehat{q}_k(\xi_k^{I_{k+1}^i}, \xi_{k+1}^i; \zeta_k) \middle| \mathcal{G}_{k+1}^N\right] = q_k(\xi_k^{I_{k+1}^i}, \xi_{k+1}^i) \sum_{\ell=1}^N \frac{\omega_k^\ell q_k(\xi_k^\ell, \xi_{k+1}^i)\left(\tau_k^\ell + h_k(\xi_k^\ell, \xi_{k+1}^i)\right)}{\sum_{\ell'=1}^N \omega_k^{\ell'} q_k(\xi_k^{\ell'}, \xi_{k+1}^i)} .$$

Moreover, conditionally to $\mathcal{F}_k^N$, the probability density function of (ξ_{k+1}^i, I_{k+1}^i) is given by

$$(x, j) \mapsto \frac{\omega_k^j \vartheta_{k+1}(\xi_k^j) p_k(\xi_k^j, x)}{\Omega_k \phi_k^N[\vartheta_{k+1}]} .$$

Therefore, this yields:

$$\begin{aligned}
\mathbb{E}\left[\widehat{\omega}_{k+1}^i \tau_{k+1}^i \middle| \mathcal{F}_k^N\right] &= \left(\phi_k^N[\vartheta_{k+1}]\right)^{-1} \sum_{j=1}^N \frac{\omega_k^j}{\Omega_k} \int \vartheta_{k+1}(\xi_k^j) \frac{q_k(\xi_k^j, x) g_{k+1}(x)}{\vartheta_{k+1}(\xi_k^j) p_k(\xi_k^j, x)} \\
&\quad \times \sum_{\ell=1}^N \frac{\omega_k^\ell q_k(\xi_k^\ell, x)\left(\tau_k^\ell + h_k(\xi_k^\ell, x)\right)}{\sum_{\ell'=1}^N \omega_k^{\ell'} q_k(\xi_k^{\ell'}, x)} p_k(\xi_k^j, x) \mathrm{d}x , \\
&= \left(\phi_k^N[\vartheta_{k+1}]\right)^{-1} \\
&\quad \times \sum_{\ell=1}^N \frac{\omega_k^\ell}{\Omega_k} \left[\int \frac{\sum_{j=1}^N \omega_k^j q_k(\xi_k^j, x)}{\sum_{\ell'=1}^N \omega_k^{\ell'} q_k(\xi_k^{\ell'}, x)} g_{k+1}(x) q_k(\xi_k^\ell, x)\left(\tau_k^\ell + h_k(\xi_k^\ell, x)\right) \mathrm{d}x\right] \\
&= \left(\phi_k^N[\vartheta_{k+1}]\right)^{-1} \phi_k^N\left[\int q_k(\cdot, x) g_{k+1}(x)\left\{\tau_k(\cdot) + h_k(\cdot, x)\right\} \mathrm{d}x\right] ,
\end{aligned}$$

which concludes the proof. □

Proof of Proposition 1. The results is proved by induction. At time $k = 0$, the result holds using that for all $1 \leq i \leq N$, $\rho_0^i = 0$ and the convention $T_0[h_0] = 0$. In addition, ϕ_0^N is a standard importance sampler estimator of ϕ_0 with $\widehat{\omega}_0^i \leq |\widehat{\omega}_0|_\infty$ so that for any bounded function h on X,

$$\mathbb{P}\left(\left|\phi_0^N[h] - \phi_0[h]\right| \geq \varepsilon\right) \leq b_0 \exp\left(-c_0 N \varepsilon^2\right) .$$

Assume the results holds for $k \geq 1$ and that $\vartheta_{k+1} = 1$ for simplicity. Write

$$\phi_{k+1}^N[\tau_{k+1}] - \phi_{k+1}\left[T_{k+1}[h_{k+1}]\right] = a_N/b_N \ ,$$

where $a_N = N^{-1}\sum_{i=1}^N \widehat{\omega}_{k+1}^i \left(\tau_{k+1}^i - \phi_{k+1}\left[T_{k+1}[h_{k+1}]\right]\right)$ and $b_N = N^{-1}\sum_{i=1}^N \widehat{\omega}_{k+1}^i$. By Lemma 2, the random variables $\{\widehat{\omega}_{k+1}^i \tau_{k+1}^i\}_{i=1}^N$ are independent conditionally on $\mathcal{F}_k^N$ and by H2,

$$\left|\widehat{\omega}_{k+1}^i \left(\tau_{k+1}^i - \phi_{k+1}\left[T_{k+1}[h_{k+1}]\right]\right)\right| \leq 2|\widehat{\omega}_{k+1}|_\infty |H_{k+1}|_\infty \ .$$

Therefore, by Hoeffding inequality,

$$\mathbb{P}\left(\left|a_N - \mathbb{E}\left[a_N \middle| \mathcal{F}_k^N\right]\right| \geq \varepsilon\right) = \mathbb{E}\left[\mathbb{P}\left(\left|a_N - \mathbb{E}\left[a_N \middle| \mathcal{F}_k^N\right]\right| \geq \varepsilon \middle| \mathcal{F}_k^N\right)\right] \leq 2\exp\left(-c_k N\varepsilon^2\right) \ .$$

On the other hand,

$$\mathbb{E}\left[a_N \middle| \mathcal{F}_k^N\right] = \phi_k^N\left[\Upsilon_k\right] \ ,$$

where

$$\Upsilon_k(x_k) = \int q_k(\cdot, x) g_{k+1}(x) \left(\tau_k(x_k) + h_{k+1}(x_k, x) - \phi_{k+1}\left[T_{k+1}[h_{k+1}]\right]\right) \mathrm{d}x \ .$$

By [24, Lemma 11], $\phi_k\left[\Upsilon_k\right] = 0$ which implies by the induction assumption that

$$\mathbb{P}\left(\left|\mathbb{E}\left[a_N \middle| \mathcal{F}_k^N\right]\right| \geq \varepsilon\right) \leq b_k \exp\left(-c_k N\varepsilon^2\right) \ .$$

Then,

$$\mathbb{P}\left(|a_N| \geq \varepsilon\right) \leq b_k \exp\left(-c_k N\varepsilon^2\right) \ .$$

Similarly, as $b_N \leq |\widehat{\omega}_k|_\infty$, by Hoeffding inequality,

$$\begin{aligned}\mathbb{P}\left(\left|b_N - \mathbb{E}\left[b_N \middle| \mathcal{F}_k^N\right]\right| \geq \varepsilon\right) \\ = \mathbb{E}\left[\mathbb{P}\left(\left|b_N - \mathbb{E}\left[b_N \middle| \mathcal{F}_k^N\right]\right| \geq \varepsilon \middle| \mathcal{F}_k^N\right)\right] \leq 2\exp\left(-c_k N\varepsilon^2\right) \ .\end{aligned}$$

Note that

$$\mathbb{E}\left[b_N \middle| \mathcal{F}_k^N\right] = \phi_k^N\left[\int q_k(\cdot, x) g_{k+1}(x) \mathrm{d}x\right] \ .$$

By the induction assumption,

$$\mathbb{P}\left(\left|\mathbb{E}\left[b_N \middle| \mathcal{F}_k^N\right] - \phi_k\left[\int q_k(\cdot, x) g_{k+1}(x) \mathrm{d}x\right]\right| \geq \varepsilon\right) \leq b_k \exp\left(-c_k N\varepsilon^2\right) \ .$$

The proof is completed using Lemma 3. □

Lemma 3. *Assume that a_N, b_N, and b are random variables defined on the same probability space such that there exist positive constants β, B, C, and M satisfying*

(i) $|a_N/b_N| \leq M$, $\mathbb{P}$-*a.s. and* $b \geq \beta$, $\mathbb{P}$-*a.s.*,

(ii) For all $\epsilon > 0$ *and all* $N \geq 1$, $\mathbb{P}\left[|b_N - b| > \epsilon\right] \leq B \exp\left(-CN\epsilon^2\right)$,

(iii) For all $\epsilon > 0$ *and all* $N \geq 1$, $\mathbb{P}\left[|a_N| > \epsilon\right] \leq B \exp\left(-CN\left(\epsilon/M\right)^2\right)$.

Then,

$$\mathbb{P}\left\{\left|\frac{a_N}{b_N}\right| > \epsilon\right\} \leq B \exp\left(-CN\left(\frac{\epsilon\beta}{2M}\right)^2\right).$$

Proof. See [8]. □

References

[1] Y. Ait-Sahalia. Transition densities for interest rate and other nonlinear diffusions. *Journal of Finance*, 54:1361–1395, 1999.

[2] Y. Ait-Sahalia. Maximum likelihood estimation of discretely sampled diffusions: a closed-form approximation approach. *Econometrica*, 70:223–262, 2002.

[3] Y. Ait-Sahalia. Closed-form likelihood expansions for multivariate diffusions. *The Annals of Statistics*, 36:906–937, 2008.

[4] A. Beskos, O. Papaspiliopoulos, G. Roberts, and P. Fearnhead. Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discusion). *J. Roy. Statist. Soc. Ser. B*, 68(3):333–382, 2006.

[5] O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. Springer, 2005.

[6] P. Del Moral, J. Jacod, and P. Protter. The Monte Carlo method for filtering with discrete-time observations. *Probability Theory and Related Fields*, 120:346 – 368, 2001.

[7] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Statist. Soc. B*, 39(1):1–38 (with discussion), 1977.

[8] R. Douc, A. Garivier, E. Moulines, and J. Olsson. Sequential Monte Carlo smoothing for general state space hidden Markov models. *Ann. Appl. Probab.*, 21(6):2109–2145, 2011.

[9] A. Doucet, S. Godsill, and C. Andrieu. On sequential Monte-Carlo sampling methods for Bayesian filtering. *Stat. Comput.*, 10:197–208, 2000.

[10] A. Doucet, S. Godsill, and C. Andrieu. On sequential monte-carlo sampling methods for bayesian filtering. *Statistics and Computing*, 10:197 – 208, 2000.

[11] P. Fearnhead, O. Papaspiliopoulos, and G Roberts. Particle filters for partially observed diffusions. *J. Roy. Statist. Soc. Ser. B*, 70(4):755–777, 2008.

[12] S. J. Godsill, A. Doucet, and M. West. Monte Carlo smoothing for non-linear time series. *J. Am. Statist. Assoc.*, 50:438–449, 2004.

[13] N. Gordon, D. Salmond, and A.F. Smith. Novel approach to nonlinear/non-Gaussian bayesian state estimation. *IEE Proc. F, Radar Signal Process*, 140:107–113, 1993.

[14] M. Hürzeler and H. R. Künsch. Monte Carlo approximations for general state-space models. *J. Comput. Graph. Statist.*, 7:175–193, 1998.

[15] N. Kantas, A. Doucet, S.S. Singh, J. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statist. Sci.*, 30(3):328–351, 2015.

[16] M. Kessler. Estimation of an ergodic diffusion from discrete observations. *Scandinavian Journal of Statistics*, 24(2):211–229, 1997.

[17] M. Kessler, A. Lindner, and M. Sorensen. *Statistical methods for stochastic differential equations*. CRC Press, 2012.

[18] G. Kitagawa. Monte-Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Statist.*, 1:1–25, 1996.

[19] S. Le Corff and G. Fort. Convergence of a particle-based approximation of the block online Expectation Maximization algorithm. *ACM Transactions on Modeling and Computer Simulation*, 23(1):2, 2013.

[20] S. Le Corff and G. Fort. Online Expectation Maximization based algorithms for inference in hidden Markov models. *Electronic Journal of Statistics*, 7:763–792, 2013.

[21] C. Li. Maximum-likelihood estimation for diffusion processes via closed-form density expansions. *The Annals of Statistics*, 41(3):1350–1380, 2013.

[22] J. Olsson, O. Cappe, R. Douc, and E. Moulines. Sequential monte carlo smoothing with application to parameter estimation in nonlinear state space models. *Bernoulli*, 14(1):155–179, 2008.

[23] J. Olsson and J. Strojby. Particle-based likelihood inference in partially observed diffusion processes using generalised Poisson estimators. *Electron. J. Statist.*, 5:1090–1122, 2011.

[24] J. Olsson and J. Westerborn. Efficient particle-based online smoothing in general hidden Markov models: the PaRIS algorithm. *ArXiv:1412.7550*, 2016.

[25] T. Ozaki. A bridge between nonlinear time series models and nonlinear stochastic dynamical systems: a local linearization approach. *Statistica Sinica*, 2:1130–135, 1992.

[26] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *J. Am. Statist. Assoc.*, 94(446):590–599, 1999.

[27] G. Poyiadjis, A. Doucet, and S.S Singh. Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, 98:65–80, 2011.

[28] I. Shoji and T. Ozaki. Estimation for nonlinear stochastic differential equations by a local linearization method 1. *Stochastic Analysis and Applications*, 16(4):733–752, 1998.

[29] Masayuki Uchida and Nakahiro Yoshida. Adaptive estimation of an ergodic diffusion process based on sampled data. *Stochastic Processes and their Applications*, 122(8):2885 – 2924, 2012.