16th European Signal Processing Conference (EUSIPCO 2008), Lausanne, Switzerland, August 25-29, 2008, copyright by EURASIP

ANALYSIS OF A SEQUENTIAL MONTE CARLO OPTIMIZATION METHODOLOGY

Joaquín Míguez

Department of Signal Theory & Communications, Universidad Carlos III de Madrid
Avenida de la Universidad 30, 28911 Leganés, Madrid (Spain).
Phone: +34916248736. E-mail: joaquin.miguez@uc3m.es. Web: www.tsc.uc3m.es/∼jmiguez

ABSTRACT

We investigate a family of stochastic exploration methods that has been recently proposed to carry out estimation and prediction in discrete-time random dynamical systems. The key of the novel approach is to identify a cost function whose minima provide valid estimates of the system state at successive time instants. This function is recursively optimized using a sequential Monte Carlo minimization (SMCM) procedure which is similar to standard particle filtering algorithms but does not require a explicit probabilistic model to be imposed on the system. In this paper, we analyze the asymptotic convergence of SMCM methods and show that a properly designed algorithm produces a sequence of system-state estimates with individually minimal contributions to the cost function. We apply the SMCM method to a target tracking problem in order to illustrate how convergence is achieved in the way predicted by the theory.

1. INTRODUCTION

Let us consider the problem of tracking the unobserved state of a discrete-time, random dynamical system. Assuming first-order Markov dynamics, a state-space model of the system consists of the pair of equations

$$\mathbf{x}_t = f(\mathbf{x}_{t-1}, \mathbf{u}_t) \quad \text{(state eq.)} \tag{1}$$

$$\mathbf{y}_t = g(\mathbf{x}_t, \mathbf{m}_t) \quad \text{(observation eq.)} \tag{2}$$

where $t \in \mathbb{N}$ denotes discrete time, $\mathbf{x}_t \in \mathbb{R}^{n_x}$ is the system state at time t, $\mathbf{y}_t \in \mathbb{R}^{n_y}$ is the associated observation and $\mathbf{u}_t \in \mathbb{R}^{n_u}$ and $\mathbf{m}_t \in \mathbb{R}^{n_m}$ are noise terms. Functions $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \to \mathbb{R}^{n_x}$ and $g : \mathbb{R}^{n_x} \times \mathbb{R}^{n_m} \to \mathbb{R}^{n_y}$ describe the state dynamics and the measurement of observations, respectively.

It is often of interest to approximate the filtering probability density function (pdf), i.e., the pdf of $\mathbf{x}_t$ given the sequence of observations $\mathbf{y}_{1:t} \triangleq \{\mathbf{y}_1, ..., \mathbf{y}_t\}$. When both f and g are linear and the noise terms are Gaussian, the Kalman filter [7] yields an exact solution. When the system is nonlinear and/or non-Gaussian, particle filters (PFs) [4, 14] provide a point-mass approximation that converges to the filtering pdf under weak assumptions on model (1)-(2) [12].

However, the practical accuracy of PFs is strongly dependent on the validity of the model. Discrepancies between the assumed model and the actual processes may lead to poor estimation performance. For example, if function f does not sufficiently account for the dynamical features of the state process, $\mathbf{x}_t$, a PF derived from such model is not likely to consistently track the sequence of states $\{\mathbf{x}_t\}_{t \in \mathbb{N}}$. E.g., in a target tracking problem we may choose a linear motion model, but the target exhibit maneuvering dynamics [14]. Mismatches between the assumed noise pdf's, either in the state or the observation equations, and the actual statistics of the sequences $\{\mathbf{x}_t\}_{t \in \mathbb{N}}$ and $\{\mathbf{y}_t\}_{t \in \mathbb{N}}$ may also lead to a performance degradation. The latter situation is not unusual. In communication receivers, for instance, observational noise is often assumed Gaussian, but impulsive processes appear in many environments [9].

In this paper, we investigate a family of stochastic exploration methods that has been recently proposed for the on-line estimation of the state sequence, $\mathbf{x}_t$, without explicit assumptions on the pdf's of the noise processes, $\mathbf{u}_t$ and $\mathbf{m}_t$ [11]. The key of the novel approach is to identify a cost function whose minima provide valid estimates of the system state at successive time instants. This function is recursively optimized using a sequential Monte Carlo minimization (SMCM) procedure which is similar to standard PFs, including sampling, weighting and selection (resampling) steps [3]. We analyze the asymptotic convergence of SMCM methods and show that a properly designed algorithm produces a sequence of system-state estimates with individually minimal contributions to the cost function. Moreover, this result can be attained only with an approximate knowledge of the system dynamics (i.e., function f does not have to be known exactly). Finally, we apply the SMCM method to a target tracking problem that yields a graphical depiction of how convergence is achieved in the way predicted by the theory.

The remaining of the paper is organized as follows. In Section 2, we introduce a description of the SMCM which is both more general and more concise than the original statement of the method in [11]. Asymptotic convergence results are stated and discussed in Section 3. Section 4 is devoted to simulations and, finally, we briefly present some conclusions in Section 5.

2. SEQUENTIAL MONTE CARLO MINIMIZATION

Let us assume that the minimization of the cost function $C(\mathbf{x}_{0:t}, \mathbf{y}_{1:t})$ is a valid, although possibly suboptimal, criterion for the estimation of $\mathbf{x}_{0:t}$. We further assume that C can be recursively decomposed as

$$C_{t-1} \triangleq C(\mathbf{x}_{0:t-1}, \mathbf{y}_{1:t-1}) \quad (t > 1), \tag{3}$$

$$C(\mathbf{x}_{0:t}, \mathbf{y}_{1:t}) \triangleq h(c(\mathbf{x}_t, \mathbf{y}_t), C_{t-1}), \tag{4}$$

where $C_0 \geq 0$ is an arbitrary constant, $c : \mathbb{R}^{n_x} \times \mathbb{R}^{n_y} \to [0, \infty)$ is a marginal cost function and $h : [0, \infty) \times [0, \infty) \to [0, \infty)$ incorporates the marginal cost into the overall cost. A typical example is an additive cost of the form

$$C(\mathbf{x}_{0:t}, \mathbf{y}_{1:t}) = \sum_{n=1}^{t} |\mathbf{y}_n - \ell(\mathbf{x}_n)| \tag{5}$$

where ℓ is some nonlinearity related to the observation equation, $c(\mathbf{x}_t, \mathbf{y}_t) = |\mathbf{y}_t - \ell(\mathbf{x}_t)|$ and $h(a, b) = a + b$.

A SMCM algorithm applies the algorithmic steps of a PF to the minimization of $C(\mathbf{x}_{0:t}, \mathbf{y}_{1:t})$. In order to describe the generic methodology we need to introduce some notation. Specifically: let M_t denote the number of samples (also called *particles*) drawn at time t; let $C_t^{(k)} = h(c(\mathbf{x}_t^{(k)}, \mathbf{y}_t), C_{t-1}^{(k)})$ be shorthand for the cost of the k-th particle; let $\mathscr{X}_0 \subset \mathbb{R}^{n_x}$ denote a set with finite Lebesgue measure, termed the *prior set*; let σ_t denote a probability mass function (pmf) with support in the set $\{1, ..., M_{t-1}\}$, referred to as the *selection pmf*; let $s : [0, \infty) \to [0, \infty)$ be termed the *cost-selection* function; and let $\tilde{p}_t(\mathbf{x}_t | \mathbf{x}_{t-1})$ denote a Markov transition pdf. The recursive SMCM algorithm, based on the technique in [11], can be outlined as follows.

1. **Initialization.** Draw M_0 random samples from the prior set $\mathscr{X}_0$. The initial costs take a non-negative constant value, $C_0^{(k)} = C_0 \geq 0$ for all k.

2. **Recursive step.** Let $\Omega_t = \{\mathbf{x}_t^{(k)}, C_t^{(k)}\}_{k=1}^{M_t}$ be the collection of particles at time t.
   (a) **Selection.** Draw indices from $\{1,...,M_t\}$ according to σ_{t+1},

$$i^{(k)} \sim \sigma_{t+1}(i), \quad k = 1,...,M_{t+1}, \quad i \in \{1,...,M_t\}. \quad (6)$$

   Set $\tilde{\mathbf{x}}_t^{(k)} = \mathbf{x}_t^{(i^{(k)})}$ and $\tilde{C}_t^{(k)} = s(C_t^{(i^{(k)})})$ in order to build $\tilde{\Omega}_t = \{\tilde{\mathbf{x}}_t^{(k)}, \tilde{C}_t^{(k)}\}_{k=1}^{M_t}$.
   (b) **Propagation.** Draw new particles and compute new costs

$$\mathbf{x}_{t+1}^{(k)} \sim \tilde{\rho}_{t+1}(\mathbf{x}_{t+1}|\tilde{\mathbf{x}}_t^{(k)}), \quad (7)$$

$$C_{t+1}^{(k)} = h(c(\mathbf{x}_{t+1}^{(k)}, \mathbf{y}_{t+1}), \tilde{C}_t^{(k)}), \quad (8)$$

   in order to build $\Omega_{t+1} = \{\mathbf{x}_{t+1}^{(k)}, C_{t+1}^{(k)}\}_{k=1}^{M_{t+1}}$.

Both the Bayesian bootstrap filter and the generic sequential importance resampling algorithm can be derived from the SMCM framework by adequately selecting functions σ_t, s, c, h and $\tilde{\rho}_t$ [10]. In general, however, the latter can be chosen quite freely in order to obtain a broader range of algorithms.

Estimation using the particles in $\Omega_t = \{\mathbf{x}_t^{(k)}, C_t^{(k)}\}_{k=1}^{M_t}$ can be as simple as choosing the sample with the lowest cost. Alternatively, given a monotonically non-increasing function $b : [0,\infty) \to [0,\infty)$, we can build an *ad hoc* discrete probability measure $\beta_t^{(k)} \propto b(C_t^{(k)})$, $k = 1,\ldots,M_t$, and use it to obtain the averaged state-estimate

$$\hat{\mathbf{x}}_t^{\beta} = \sum_{k=1}^{M} \beta_t^{(k)} \mathbf{x}_t^{(k)}. \quad (9)$$

A typical choice of function b is $b(z) = 1/\left|z - \min_k\{C_t^{(k)}\} + 1/M\right|^q$, with fixed $q \geq 1$, in order to emphasize low-cost particles [10].

3. CONVERGENCE OF THE MARGINAL COSTS

3.1 Preliminary definitions and results

The straightforward question regarding the convergence of the SMCM methods as described in Section 2 is whether the algorithm can produce a sequence of state estimates with a cost which is arbitrarily close to the minimum one. This type of convergence, however, is hard to analyze without further assumptions on the structure of function $C(\mathbf{x}_{0:t}, \mathbf{y}_{1:t})$. Instead, we are going to consider the simpler, but still meaningful, problem of the convergence of the marginal costs given by function $c(\mathbf{x}_t, \mathbf{y}_t)$. In particular, given deterministic observations $\mathbf{y}_{1:T}$, where $T < \infty$ but arbitrarily large, let

$$\mathbf{x}_{t,o} \triangleq \arg\min_{\mathbf{x}\in\mathbb{R}^{n_x}} c(\mathbf{x}, \mathbf{y}_t), \quad (10)$$

$$\hat{\mathbf{x}}_{t,o} \triangleq \arg\min_{\mathbf{x}\in\{\mathbf{x}_t^{(k)}\}_{k=1}^{M_t}} c(\mathbf{x}, \mathbf{y}_t) \quad (11)$$

be a marginal-cost minimizer in the state-space (not necessarily unique) and the marginal-cost minimizer in the discrete set $\{\mathbf{x}_t^{(k)}\}_{k=1}^{M_t}$ generated by the SMCM algorithm at time t. We tackle the question of whether $\hat{\mathbf{x}}_{t,o}$ can be made arbitrarily close to $\mathbf{x}_{t,o}$, possibly with arbitrarily large numbers $\{M_t\}_{t=0}^T$ and for sufficiently large t (yet $t \leq T$). When the answer is positive, we say that the SMCM algorithm attains *marginal-cost convergence*. For many cost functions in practical systems the minimization of the sequence of marginal costs will lead to an adequate steady-state performance of the SMCM algorithm, although possibly after some transient period for which accurate estimation may not be attained.

We need to establish some assumptions and specific notation to address this problem. We start by assuming that the selection pmf, σ_t, and the conditional propagation pdf, $\tilde{\rho}_t$, can be joined into a single propagation density. In particular, we define the joint pdf

$$\rho_t(\mathbf{x}_t, k) \triangleq \tilde{\rho}_t(\mathbf{x}_t|\mathbf{x}_{t-1}^{(k)})\sigma_t(k), \quad (12)$$

hence the selection and propagation steps can be jointly written as $\left(\mathbf{x}_t^{(i)}, k^{(i)}\right) \sim \rho_t(\mathbf{x}_t, k)$. The (auxiliary) particle index can be summed out, to yield

$$\mathbf{x}_t^{(i)} \sim \rho_t(\mathbf{x}_t) \triangleq \sum_{k=1}^{M_{t-1}} \rho_t(\mathbf{x}_t, k). \quad (13)$$

Note that the pdf's $\rho_t(\mathbf{x}_t, k)$ and $\rho_t(\mathbf{x}_t)$ are implicitly conditional on the particle set $\{\mathbf{x}_{t-1}^{(k)}\}_{k=1}^{M_{t-1}}$.

Since $\mathbf{y}_{1:T}$ is deterministic, so they are the sequences of marginal cost minimizers, $\{\mathbf{x}_{t,o}\}_{t=1}^T$ (note that multiple sequences of minimizers may exist). Let us also assume the availability of functions $a_t : \mathbb{R}^{n_x} \to \mathbb{R}^{n_x}$ and a distance $d(\cdot,\cdot)$, properly defined in $\mathbb{R}^{n_x}$, for which:

Assumption 1 *There exists a finite constant $A \in \mathbb{R}$ such that, for all t, $A_t \triangleq d\left(a_t(\mathbf{x}_{t,o}), \mathbf{x}_{t+1,o}\right) < A$.*

Assumption 2 *For any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^{n_x}$, if $d(\mathbf{x}, \mathbf{x}') < \infty$ then $d\left(a_t(\mathbf{x}), a_t(\mathbf{x}')\right) < \infty$, $\forall t$.*

Given Assumptions 1 and 2, the dynamics of the minimizers can be related to the dynamics of sequences of sets with certain properties. To be specific, let $\{\mathscr{X}_t\}_{t\in\mathbb{N}^*}$, where $\mathbb{N}^* = \mathbb{N} \cup \{0\}$, be a sequence of sets in $\mathbb{R}^{n_x}$. We impose the following assumptions.

Assumption 3 *For all $t \in \mathbb{N}^*$, $\mathscr{X}_t = \sum_{i=1}^{n_t} \mathscr{O}_{i,t}$, where $n_t < \infty$ and $\mathscr{O}_{1,t}, \ldots, \mathscr{O}_{n,t}$ are bounded open cells [2] in $\mathbb{R}^{n_x}$.*

Assumption 4 *Let*

$$a_t(\mathscr{X}_t) \triangleq \{\mathbf{x} \in \mathbb{R}^{n_x} : \mathbf{x} = a(\mathbf{y}) \quad \textit{for some} \quad \mathbf{y} \in \mathscr{X}_t\} \quad (14)$$

and let $\mathscr{A}, \mathscr{B} \subset \mathbb{R}^{n_x}$. We define the distance between the sets $\mathscr{A}$ and $\mathscr{B}$ as $d(\mathscr{A}, \mathscr{B}) \triangleq \inf_{\mathbf{x}\in\mathscr{A}, \mathbf{y}\in\mathscr{B}} d(\mathbf{y}, \mathbf{x})$, and the complement of $\mathscr{A}$ as $\overline{\mathscr{A}} \triangleq \{\mathbf{y} \in \mathbb{R}^{n_x} : \mathbf{y} \notin \mathscr{A}\}$. Then, for all $t \in \mathbb{N}$, $d\left(\overline{\mathscr{X}}_t, a(\mathscr{X}_{t-1})\right) > 0$.

Let $m(\cdot)$ denote the Lebesgue measure. Assumption 3 constrains the class of sets that can appear in the sequence $\{\mathscr{X}_t\}_{t\in\mathbb{N}^*}$. However, for any measurable set $\mathscr{E} \in \mathbb{R}^{n_x}$, such that $m(\mathscr{E}) < \infty$, and for any $\varepsilon > 0$, there exists a sequence of bounded open cells $\{\mathscr{O}\}_{i=1}^n$, with $n < \infty$, such that that the measure of the symmetric difference between the set $\mathscr{E}$ and its approximation $\mathscr{X}_n = \cup_{i=1}^n \mathscr{O}_i$ is less than ε [2], i.e., $m(\mathscr{E} \triangle \mathscr{X}_n) = m\left((\mathscr{E} - \mathscr{X}_n) \cup (\mathscr{X}_n - \mathscr{E})\right) < \varepsilon$.

It is of interest to upper bound the separation between a minimizer $\mathbf{x}_{t+1,o}$ and the set $\mathscr{X}_{t+1}$ given the distance $d(\mathscr{X}_t, \mathbf{x}_{t,o})$. Let us define

$$D_{t+1}^{sup} \triangleq \max_{\mathbf{x},\mathbf{x}'\in\mathbb{R}^{n_x}} d(a_t(\mathbf{x}), a_t(\mathbf{x}')), \quad \text{subject to } d(\mathbf{x}, \mathbf{x}') \leq d(\mathscr{X}_t, \mathbf{x}_{t,o}), \quad (15)$$

for $t < T$, and the associated bound for the distance increment

$$\triangle_{t+1}^{sup} \triangleq D_{t+1}^{sup} - d(\mathscr{X}_t, \mathbf{x}_{t,o}). \quad (16)$$

Because of Assumption 2, if $d(\mathscr{X}_t, \mathbf{x}_{t,o}) < \infty$ then $D_{t+1}^{sup} < \infty$. From Assumption 4, we can also define the sequence of positive numbers

$$d_{t+1}^{out} \triangleq d(\overline{\mathscr{X}}_{t+1}, a_t(\mathscr{X}_t)) > 0, \quad t \in \mathbb{N}^* \quad (17)$$

and, taking together (16), (17) and Assumption 1, we introduce the sequence

$$K_{t+1} \triangleq d_{t+1}^{out} - A_t - \triangle_{t+1}^{sup}, \quad t < T. \quad (18)$$

The following Proposition will be a key tool to prove the marginal-cost convergence of SMCM methods.

Proposition 1 *If* $K_t \geq 0$ *for all* $t \in \{t_1+1,\dots,t_2\}$ *and* $d(\mathscr{X}_{t_1}, \mathbf{x}_{t_1,o}) < \infty$*, then for all* $t \in \{t_1+1,\dots,t_2\}$*,*

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(\mathscr{X}_{t_1}, \mathbf{x}_{t_1,o}) - \sum_{n=t_1+1}^{t} K_n\right\}. \quad (19)$$

See Appendix A for a proof. Intuitively, Proposition 1 states that if the sequence K_t is non-negative for $t_1 < t \leq t_2$, then the sets $\mathscr{X}_{t_1<t\leq t_2}$ *converge* towards the minimizers $\mathbf{x}_{t,o}$. If $t_2 - t_1$ can be made arbitrarily large and $K_t = 0$ only a finite number of times, then $d(\mathscr{X}_{t_2}, \mathbf{x}_{t_2,o}) = 0$.

3.2 Convergence

We will show that, for sufficiently large natural numbers $\{M_t\}_{t=1}^T$, a SMCM algorithm can be designed in such a way that the state-space sample

$$\mathbf{x}_{t,*} \triangleq \arg\min_{\mathbf{x}\in\{\mathbf{x}_t^{(i)}\}_{i=1}^{M_t}} d(\mathbf{x}, \mathbf{x}_{t,o}). \quad (20)$$

converges to $\mathbf{x}_{t,o}$, in probability, when t is sufficiently large. To express this result, we will use notation $\text{Prob}\{E\}$ to denote the probability of a random event E.

Proposition 2 *Consider a sequence of deterministic observations* $\mathbf{y}_{t_1:t_2}$ *such that* $d(\mathbf{x}_{t_1,*}, \mathbf{x}_{t_1,o}) < \infty$*. For any fixed real numbers* $\varepsilon, \delta > 0$ *there exist pdf's* $\{\rho_t\}_{t_1<t\leq t_2}$ *and natural numbers* $\{M_{t,\varepsilon,\delta}\}_{t_1<t\leq t_2}$ *such that, for all* $\{M_t > M_{t,\varepsilon,\delta}\}_{t_1<t\leq t_2}$*,*

$$Prob\{d(\mathbf{x}_{t,*}, \mathbf{x}_{t,o}) \leq \varepsilon + \max\{0, d(\mathbf{x}_{t_1,*}, \mathbf{x}_{t_1,o}) - J_t\}\} > 1-\delta, \quad (21)$$

with $0 < J_t < J_{t+1}$*, for all* $t \in \{t_1+1,\dots,t_2\}$*.*

The proof, given in Appendix B, consists of two steps. First, it is shown that the samples generated by a properly defined SMCM algorithm are contained in a sequence of sets that comply with Assumptions 3 and 4. Then, Proposition 1 is applied to obtain the desired result, with $J_t = \sum_{n=t_1+1}^{t} K_n$ and $K_n > 0$ for $n = t_1+1,...,t_2$. Intuitively, Proposition 2 says that a soundly designed SMCM algorithm produces clouds of particles closer and closer (as time evolves) to the marginal-cost minimizers. Two remarks are relevant:

- If, for all t, $|J_{t+1} - J_t| \geq \varepsilon_J > 0$, with ε_J independent of t, and we allow t_2 to be arbitrarily large, then Proposition 2 implies that $\mathbf{x}_{t,*} \to \mathbf{x}_{t,o}$ in probability.
- If the marginal cost c is uniformly continuous at the minimizing points $\mathbf{x}_{t,o}$, then the convergence of $\mathbf{x}_{t,*}$ implies the convergence of $\hat{\mathbf{x}}_{t,o}$ (i.e., the convergence of the minimum marginal-cost estimates produced by the SMCM algorithm).

4. SIMULATION EXAMPLE

The problem of target tracking in a 2-dimensional space can be well represented using a random dynamical model [6]. The system state at time $t = 0,1,2,...$ consists of the target position, $\mathbf{r}_t = [r_{1,t}, r_{2,t}]^\top \in \mathbb{R}^2$, and velocity, $\mathbf{v}_t = [v_{1,t}, v_{2,t}]^\top \in \mathbb{R}^2$ and its evolution can be modeled as

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{Q}\mathbf{u}_t, \quad (22)$$

where $\mathbf{x}_t = [\mathbf{r}_t^\top, \mathbf{v}_t^\top]^\top \in \mathbb{R}^4$ is the state vector, $\mathbf{u}_t \in \mathbb{R}^4$ is a zero-mean Gaussian process. Matrices $\mathbf{A}$ and $\mathbf{Q}$ are defined as

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & T_o & 0 \\ 0 & 1 & 0 & T_o \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{Q} = \begin{bmatrix} \frac{1}{2}T_o^2 & 0 \\ 0 & \frac{1}{2}T_o^2 \\ T_o & 0 \\ 0 & T_o \end{bmatrix}, \quad (23)$$

where $T_o > 0$ is the observation period (i.e., the time between consecutive observations). The *a priori* pdf of the state is also Gaussian, namely $\mathbf{r}_0 \sim N(\mathbf{0}, \frac{5}{2}\mathbf{I}_2)$ and $\mathbf{v}_0 \sim N(\mathbf{0}, \frac{1}{8}\mathbf{I}_2)$, where $\mathbf{I}_2$ is the 2×2 identity matrix.

Observations are collected through a set of $N = 10$ sensors that measure the power of a radio signal transmitted from the target. The sensor positions are known, denoted as $\mathbf{b}_i \in \mathbb{R}^2$, $i = 1,\dots,N$, and they lie in a square region of side $L = 2000$ meters. Assuming the log-normal model widely used in cellular communications [13], the observation at time t in the i-th sensor is

$$y_{i,t} = g_i(\mathbf{r}_t) \triangleq 10\log_{10}\left(\frac{P_0}{\|\mathbf{r}_t - \mathbf{b}_i\|_2^\gamma}\right) + m_{i,t} \quad \text{(dB)}, \quad i = 1,\dots,N, \quad (24)$$

where $P_0 = 1$ is the transmitted power, $\gamma = 2$ determines the rate of the (exponential) power decay and $m_{i,t}$ is zero-mean Gaussian observational noise with variance $\sigma_m^2 = 1$. We use $\mathbf{y}_t = [y_{1,t},\dots,y_{N,t}]^\top \in \mathbb{R}^N$ to denote the collection of observations at time t.

In order to estimate the sequence of positions $\mathbf{r}_{0:t}$ from the observations $\mathbf{y}_{1:t}$, we have applied a SMCM method and a mixture Kalman filter (MKF) algorithm that approximates the pdf $p(\mathbf{r}_{0:t}|\mathbf{y}_{1:t})$ by integrating out the sequence of velocities $\mathbf{v}_{0:t}$ [6]. The SMCM technique is constructed as follows.

1. Cost function: We define

$$C(\mathbf{r}_{1:t}, \mathbf{y}_{1:t}) \triangleq \lambda C(\mathbf{r}_{1:t-1}, \mathbf{y}_{1:t-1}) + c(\mathbf{r}_t, \mathbf{y}_t), \quad (25)$$

   where $c(\mathbf{r}_t, \mathbf{y}_t) = \sum_{i=1}^{N} g_i(\mathbf{r}_t)$ and $\lambda = 0.99$ is a forgetting factor. This definition implies, when compared with (4), that $H(a,b) = \lambda b + a$. As a reference, we approximate the minima of the marginal costs, $c(\mathbf{r}_t, \mathbf{y}_t)$, using an accelerated random search (ARS) algorithm [1] with a large number of iterations.
2. Initialization: Let $M_t = M = 3000$ for all t. At $t = 0$, draw M samples $\mathbf{r}_0^{(i)} \sim N([100, 200]^\top, \mathbf{I}_2)$ and set $C_0^{(i)} = 0$, $i = 1,...,M$. Note that the initial state is actually distributed as $\mathbf{r}_0 \sim N(\mathbf{0}, \frac{5}{2}\mathbf{I}_2)$. We generate the initial samples of the SMCM algorithm far from the $\mathbf{0}$ point in order to clearly show how convergence is achieved in a few time steps.
3. Recursive step: The selection pdf at time $t+1$ is

$$\sigma_{t+1}(i) \propto \left|C_t^{(i)} - \min_k C_t^{(k)} + \frac{1}{M}\right|^{-1}, \quad (26)$$

   hence indices are drawn as $i^{(k)} \sim \sigma_{t+1}$, $k = 1,...,M$, and the selected particles are $\left\{\tilde{\mathbf{r}}_t^{(k)} = \mathbf{r}_t^{(i^{(k)})}, \tilde{C}_t^{(k)} = C_t^{(i^{(k)})}\right\}_{k=1}^{M}$. The propagation pdf is Gaussian,

$$\mathbf{r}_{t+1}^{(k)} \sim \tilde{\rho}(\mathbf{r}_{t+1}|\tilde{\mathbf{r}}_t^{(k)}) = N\left(a_t(\tilde{\mathbf{r}}_t^{(k)}), 40^2\mathbf{I}_2\right), \quad (27)$$

   where $a_t(\tilde{\mathbf{r}}_t^{(k)}) = \tilde{\mathbf{r}}_t^{(k)} - \tilde{\mathbf{r}}_{t-1}^{(k)}$. Note that we do not generate samples from the velocity vector, but instead use a rough approximation of it (by subtracting two successive positions) in order to generate new particles.

Figure 1 shows an example of application of the MKF and SMCM algorithms to track the target from $t = 0$ to $t = 266$ with $T_o = 1$ second. The SMCM algorithm recovers from its poor initialization in a few time steps and then stays locked to the true trajectory for the complete simulation. However, from this example it is apparent that the MKF algorithm attains a much better estimation accuracy. This result could be easily expected, since the MKF method fully exploits the statistical structure of the dynamic model.

Figure 2 shows the clouds of particles generated by the SMCM algorithm at time instants $t = 6, 20, 40, 60, 80$ for another simulation run. Again because of its poor initialization, the particle sets up to time $t = 6$ do not *wrap* around the marginal-cost minimizers approximated using the ARS algorithm (shown as black squares). For

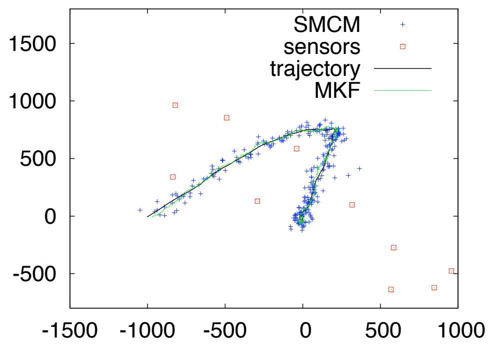

Figure 1: A sample target trajectory (black solid line) and its estimates using the MKF algorithm (green solid line) and the SMCM technique (blue crosses). Sensor locations are depicted by red squares.

$t = 20, 40, 60, 80$, however, we see that the minimizers lie clearly inside the particle clouds. This implies that, for sufficiently large M, we can obtain particles as close as we wish to these minimizers, as predicted by Proposition 2.

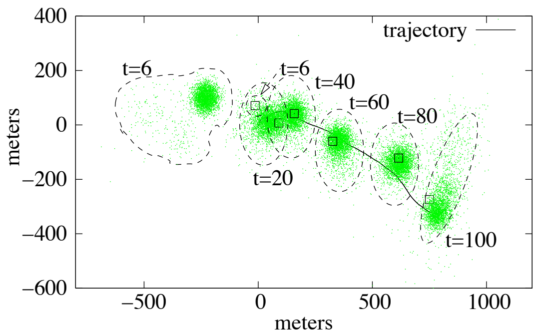

Figure 2: Sets of samples generated by the SMCM method at different time instants. It can be seen that, for $t \geq 20$, the marginal-cost minimizers (black squares) lie within the particle sets (dotted).

Figure 3 shows the smallest marginal costs attained by the SMCM algorithm (solid line) from $t = 1$ to $t = 100$ for the same simulation as in Figure 2. Here, we can clearly observe the convergence period from $t = 1$ to $t = 10$. The specific values of the smallest marginal costs at times $t = 2, 4, 6, \ldots, 20, 40, 60, 80, 100$ for the SMCM technique and the ARS algorithm are also shown. After convergence (which occurs for $t = 10$) both methods yield practically identical marginal costs. We note that the number of iterations of the ARS used in these simulations is one order of magnitude higher than the number of particles, M.

5. CONCLUSIONS

We have revisited a recently proposed Monte Carlo methodology aimed at the sequential minimization of cost functions that can be recursively written in terms of the sequences of states and observations of a discrete-time Markovian dynamical system. The technique, coined sequential Monte Carlo minimization (SMCM), is more flexible than conventional sequential Monte Carlo (SMC) methods, in the sense that algorithms can be derived with very weak assumptions on the state and observation noise processes. The main contribution of the paper is the analysis of the asymptotic convergence of the state estimates produced by SMCM algorithms. We have found sufficient conditions to ensure that, for sufficiently large numbers of particles and conditional on an arbitrary but fixed sequence of observations, the novel method yields approximations of

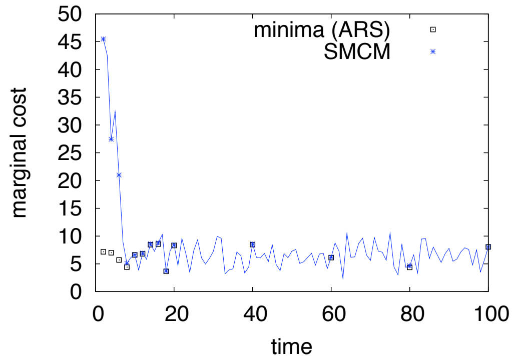

Figure 3: Minimum marginal costs (approximated via the ARS algorithm) and lowest marginal costs attained by the SMCM method. For $t > 10$, both approximations match.

the minima of the sequence of marginal cost functions that converge, in probability, to the true sequence of minima. Moreover, our analysis shows that convergent SMCM algorithms can be designed without a precise knowledge of the system dynamics. Instead, any function of the system state that can be used to predict the minimum of the marginal cost function at time t from the minimum at time $t-1$ with a finite error can, in theory, be used.

A. PROOF OF PROPOSITION 1

Choose any $t \in \{t_1+1, \ldots, t_2\}$. From Assumption 4 we deduce an initial upper bound of $d(\mathscr{X}_t, \mathbf{x}_{t,o})$,

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(a_{t-1}(\mathscr{X}_{t-1}), \mathbf{x}_{t,o}) - d_t^{out}\right\}, \qquad (28)$$

which, using the triangular inequality,

$$\begin{aligned} d(a_{t-1}(\mathscr{X}_{t-1}), \mathbf{x}_{t,o}) \quad &\leq \quad d(a_{t-1}(\mathscr{X}_{t-1}), a_{t-1}(\mathbf{x}_{t-1,o})) \\ &\quad + d(a_{t-1}(\mathbf{x}_{t-1,o}), \mathbf{x}_{t,o}), \end{aligned} \qquad (29)$$

and Assumption 1 yields

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(a_{t-1}(\mathscr{X}_{t-1}), a_{t-1}(\mathbf{x}_{t-1,o})) + A_{t-1} - d_t^{out}\right\}. \qquad (30)$$

From the definition of D_t^{sup} and $\triangle_t^{sup}$, we have

$$d(a_{t-1}(\mathscr{X}_{t-1}), a_{t-1}(\mathbf{x}_{t-1,o})) \leq D_t^{sup} = \triangle_t^{sup} + d(\mathscr{X}_{t-1}, \mathbf{x}_{t-1,o}), \qquad (31)$$

hence, substituting (31) into (30), we arrive at the inequality

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(\mathscr{X}_{t-1}, \mathbf{x}_{t-1,o}) + \triangle_t^{sup} + A_{t-1} - d_t^{out}\right\} \qquad (32)$$

and, since we have previously defined $K_t = d_t^{out} - \triangle_t^{sup} - A_{t-1}$, we readily obtain

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(\mathscr{X}_{t-1}, \mathbf{x}_{t-1,o}) - K_t\right\}. \qquad (33)$$

The assumption $K_t \geq 0$, $t > t_1$, enables us to recursively apply (33) to find the relationship

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(\mathscr{X}_{t_1}, \mathbf{x}_{t_1,o}) - \sum_{n=t_1+1}^{t} K_n\right\}, \qquad (34)$$

valid for any $t > t_1$. $\square$

B. PROOF OF PROPOSITION 2

Part 1: Let $\varepsilon, \delta > 0$ be arbitrarily small but fixed positive real numbers. Let $\mathscr{X}_0 = \cup_{i=1}^{n_0} \mathscr{O}_{i,0}$, with arbitrary $n_0 < \infty$ and bounded open cells $\mathscr{O}_{i,0} \in \mathbb{R}^{n_x}$. Since $m(\mathscr{O}_{i,0}) < \infty$ for all $1 \leq i \leq n_0$ and, as a consequence, $m(\mathscr{X}_0) < \infty$, we can build a proper uniform pdf in $\mathscr{X}_0$, denoted $\rho_0 = U(\mathscr{X}_0)$, i.e., $\rho_0(\mathbf{x}) = \frac{1}{m(\mathscr{X}_0)} > 0$ for all $\mathbf{x} \in \mathscr{X}_0$. If we draw a set of M_0 independent and identically distributed samples from ρ_0, denoted $\Omega_0 = \{\mathbf{x}_0^{(i)}\}_{i=1}^{M_0}$, we can invoke the weak law of large numbers (see, e.g., [5, Chapter 7]) to claim that, for any $\mathbf{x} \in \mathscr{X}_0$, there exists $M_{0,\varepsilon,\delta} \in \mathbb{N}$ such that, for all $M_0 > M_{0,\varepsilon,\delta}$,

$$\text{Prob}\left\{d(\mathbf{x}, \mathbf{x}_0^{(k)}) < \varepsilon \quad \text{for some } k \in \{1, \ldots, M_0\}\right\} > 1 - \delta. \quad (35)$$

Next, we build the sequence $\{\mathscr{X}_t\}_{t=0}^T$ recursively. Let us assume that $\mathscr{X}_{t-1}$ has already been constructed, together with the associated set of discrete points $\{\mathbf{x}_{t-1}^{(i)}\}_{i=1}^{M_{t-1}}$ drawn from a pdf ρ_{t-1}. We define

$$\eta \triangleq \sup_{\mathbf{x} \in \mathscr{X}_{t-1}, k \in \{1,\ldots,M_{t-1}\}} d(\mathbf{x}, \mathbf{x}_{t-1}^{(k)}) < \infty \quad (36)$$

and note that $\{\mathbf{x}_{t-1}^{(i)}\}_{i=1}^{M_{t-1}}$ is, then, a finite η-net (note that M_{t-1} can be arbitrarily large but finite) of $\mathscr{X}_{t-1}$ [8, Chapter 2]. As a consequence, it turns out that

$$a_{t-1}(\mathscr{X}_{t-1}) \subseteq \cup_{i=1}^{M_{t-1}} a_{t-1}\left(\mathscr{B}(\mathbf{x}_{t-1}^{(i)}, \eta)\right), \quad (37)$$

where $\mathscr{B}(\mathbf{x}, \eta) \triangleq \{\mathbf{x}' \in \mathbb{R}^{n_x} : d(\mathbf{x}, \mathbf{x}') < \eta\}$ denotes an open ball centered at $\mathbf{x}$. Moreover, if we construct

$$\mathscr{X}_t \triangleq \cup_{i=1}^{M_{t-1}} \mathscr{B}\left(a_{t-1}(\mathbf{x}_{t-1}^{(i)}), r_t\right), \quad (38)$$

with sufficiently large r_t, then we can write

$$a_{t-1}(\mathscr{X}_{t-1}) \subseteq \cup_{i=1}^{M_{t-1}} a_{t-1}\left(\mathscr{B}(\mathbf{x}_{t-1}^{(i)}, \eta)\right) \subset \mathscr{X}_t. \quad (39)$$

In particular, (39) holds if we take

$$\begin{aligned} r_t &> \sup_{\mathbf{x}, \mathbf{x}' \in \mathbb{R}^{n_x}, d(\mathbf{x}, \mathbf{x}') < 2\eta} d(a_{t-1}(\mathbf{x}), a_{t-1}(\mathbf{x}')) \\ &\geq \sup_{\mathbf{x}, \mathbf{x}' \in \mathscr{X}_{t-1}, d(\mathbf{x}, \mathbf{x}') < 2\eta} d(a_{t-1}(\mathbf{x}), a_{t-1}(\mathbf{x}')), \end{aligned} \quad (40)$$

(note that the suprema are finite because of Assumption 2). Given $M_{t-1} < \infty$, (38) and (39) ensure that $\mathscr{X}_t$ complies with Assumptions 3 and 4.

One possible way to build the pdf ρ_t associated to $\mathscr{X}_t$ (not unique) is to define $\sigma_t(k) > 0$ for all $k \in \{1, \ldots, M_{t-1}\}$ (this is possible, since $M_{t-1} < \infty$) and $\tilde{\rho}_t(\mathbf{x}_t|\mathbf{x}_{t-1}^{(k)}) = U\left(\mathscr{B}\left(a_{t-1}(\mathbf{x}_{t-1}^{(k)}), r_t\right)\right)$. Let $\mathbf{x}_t^{(k)} \sim \rho_t(\mathbf{x}_t)$, $k = 1, ..., M_t$. Since M_{t-1} is arbitrarily large but finite, $m(\mathscr{X}_t) < \infty$ and the weak law of large numbers can be applied to ensure that, for any $\mathbf{x} \in \mathscr{X}_t$, there exists $M_{t,\varepsilon,\delta}$ such that, for all $M_t > M_{t,\varepsilon,\delta}$

$$\text{Prob}\left\{d(\mathbf{x}, \mathbf{x}_t^{(k)}) < \varepsilon \quad \text{for some } k \in \{1, \ldots, M_t\}\right\} > 1 - \delta. \quad (41)$$

Part 2: In order to apply Proposition 1 to the sequence $\{\mathscr{X}_t\}_{t=t_1}^{t_2}$ we need to guarantee that $d_t^{out} = d\left(\overline{\mathscr{X}}_t, a_{t-1}(\mathscr{X}_{t-1})\right)$ be large enough. Specifically, the inequality

$$d_t^{out} > A_{t-1} + \triangle_t^{sup} \quad (42)$$

must hold for all $t \in \{t_1, \ldots, t_2\}$. From (38) and (40) we infer that

$$d_t^{out} \geq r_t - \sup_{\mathbf{x}, \mathbf{x}' \in \mathbb{R}^{n_x}, d(\mathbf{x}, \mathbf{x}') < 2\eta} d(a_{t-1}(\mathbf{x}), a_{t-1}(\mathbf{x}')). \quad (43)$$

Since it is always possible to choose r_t such that

$$r_t > \sup_{\mathbf{x}, \mathbf{x}' \in \mathbb{R}^{n_x}, d(\mathbf{x}, \mathbf{x}') < 2\eta} d(a_{t-1}(\mathbf{x}), a_{t-1}(\mathbf{x}')) + \triangle_t^{sup} + A_{t-1} \quad (44)$$

for all $t \in \{t_1 + 1, \ldots, t_2\}$ then it is also always possible to make the inequality (42) hold for $t \in \{t_1 + 1, \ldots, t_2\}$ and, as a consequence, $K_t = d_t^{out} - A_{t-1} - \triangle_t^{sup} > 0$ whenever $t_1 < t \leq t_2$.

Therefore, we can apply Proposition 1 to obtain

$$d(\mathscr{X}_t, \mathbf{x}_{t,o}) \leq \max\left\{0, d(\mathscr{X}_{t_1}, \mathbf{x}_{t_1,o}) - \sum_{n=t_1+1}^{t} K_n\right\}. \quad (45)$$

Also, from the definition of sequences $\{\mathscr{X}_t\}_{t=0}^T$ and $\{\rho_t\}_{t \in \mathbb{N}}$, there exists $M_{t,\varepsilon,\delta}$ such that, for all $M_t > M_{t,\varepsilon,\delta}$ and all $\mathbf{x} \in \mathscr{X}_t$,

$$\text{Prob}\{d(\mathbf{x}_t^{(k)}, \mathbf{x}) < \varepsilon, \quad \text{for some } k \in \{1, \ldots, M_t\}\} > 1 - \delta. \quad (46)$$

Taking together (45) and (46), and setting $J_t = \sum_{n=t_1}^{t} K_n$, we arrive at

$$\begin{aligned} \text{Prob}\Big\{d(\mathbf{x}_t^{(k)}, \mathbf{x}_{t,o}) &< \varepsilon + \max\{0, d(\mathscr{X}_{t_1}, \mathbf{x}_{t_1,o}) - J_t\} \\ &\text{for some } k \in \{1, \ldots, M_t\}\Big\} > 1 - \delta. \end{aligned} \quad (47)$$

Since (42) holds when $t_1 < t \leq t_2$, it is apparent that $J_t > J_{t-1}$. □

REFERENCES

[1] M. J. Appel, R. Labarre, and D. Radulovic. On accelerated random search. *SIAM Journal of Optimization*, 14(3):708–730, 2003.

[2] R. G. Bartle. *The Elements of Integration and Lebesgue Measure*. Wiley Classics Library. Wiley & sons, New York, 1995.

[3] P. M. Djurić, J. H. Kotecha, J. Zhang, Y. Huang, T. Ghirmai, M. F. Bugallo, and J. Míguez. Particle filtering. *IEEE Signal Processing Magazine*, 20(5):19–38, September 2003.

[4] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York (USA), 2001.

[5] B.V. Gnedenko. *The Theory of Probability, 6th ed.* Gordon and Breach Science Publishers, 1997.

[6] F. Gustafsson, F. Gunnarsson, N. Bergman, U. Forssell, J. Jansson, R. Karlsson, and P.-J. Nordlund. Particle filters for positioning, navigation and tracking. *IEEE Transactions Signal Processing*, 50(2):425–437, February 2002.

[7] R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82:35–45, 1960.

[8] A. N. Kolmogorov and S. V. Fomin. *Introductory Real Analysis*. Dover Publications, New York, 1975.

[9] D. Middleton. Non-gaussian noise models in signal processing for telecommunications: new methods an results for class A and class B noise models. *IEEE Transactions Information Theory*, 45(4):1129–1149, May 1999.

[10] J. Míguez. Analysis of selection methods for cost-reference particle filtering with applications to maneuvering target tracking and dynamic optimization. *Digital Signal Processing*, 17:787–807, 2007.

[11] J. Míguez, M. F. Bugallo, and P. M. Djurić. A new class of particle filters for random dynamical systems with unknown statistics. *EURASIP Journal on Advances in Signal Processing*, 2004(15):2278–2294, November 2004.

[12] P. Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer, 2004.

[13] T. S. Rappaport. *Wireless Communications: Principles and Practice (2nd edition)*. Prentice-Hall, Upper Saddle River, NJ (USA), 2001.

[14] B. Ristic, S. Arulampalam, and N. Gordon. *Beyond the Kalman Filter*. Artech House, Boston, 2004.